

UNIVERSITY OF LEICESTER

DOCTORAL THESIS

**Precision Photometry for the Next
Generation Transit Survey**

Author:

Alexander CHAUSHEV

Supervisors:

Dr. M. GOAD
& Dr M. BURLEIGH

*A thesis submitted in fulfilment of the requirements
for the degree of Doctor of Philosophy*

in the

X-Ray and Observational Astronomy Group
Department of Physics

September 20, 2018

Declaration of Authorship

I, Alexander CHAUSHEV, declare that this thesis titled, "Precision Photometry for the Next Generation Transit Survey" and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed:

Date:

“In the beginning the Universe was created. This had made many people very angry and has been widely regarded as a bad move.”

Douglas Adams, *The Restaurant at the End of the Universe*

UNIVERSITY OF LEICESTER

Abstract

College of Science and Engineering

Department of Physics

Doctor of Philosophy

Precision Photometry for the Next Generation Transit Survey

by Alexander CHAUSHEV

Finding Neptune and Super-Earth sized planets from the ground poses a serious technical challenge. The Next Generation Transit Survey (NGTS) is a ground-based transit survey designed to produce precision photometry down to 13th magnitude to enable these kinds of discoveries. The production of high signal to noise calibration frames is discussed as part of the NGTS reduction pipeline. Next, the red noise properties of the data are looked at in detail in order to understand if any of the pipeline components are contributing correlated noise to the lightcurves. The NGTS data is search for single transit events using change point analysis, producing one strong single transit candidate. Finally, photometric follow-up observations from the South Africa Astronomical Observatory are undertaken with the aim of characterising and vetting NGTS candidates.

Acknowledgements

I would like to thank my supervisors Mike Goad and Matt Burleigh for their guidance these last four years. To Mike for his patience in teaching me, the helpful discussions and tireless support he has provided me. To Matt for teaching me the ins and outs of observing. Many thanks are due to my office mates Ian and Sarah for their helpful advice and friendly chats. I would also like to thank the numerous members of the NGTS consortium who have helped me throughout the PhD. Thank you to the STFC for providing me with a studentship to complete this research and to the members of the XROA research group and the Department of Physics for providing a welcoming and supportive environment.

I would also like to thank the many people in my life who have provided encouragement and support, without which this thesis would not have been possible. Thank you to my parents, who have sacrificed so much so that I may follow my passion. To my father who inspired my passion for science, and to my mother who encouraged me to take on the challenges I needed to get here. Thank you to my sister Linda, for all the GBK dinners and late nights playing video games. Thank you to Liam for endless NGTS chats and climbing, Tom for the 2am Asda trips, Mark for good food, Hossam for coffee and Nishad for being an all round great friend. Thanks to Jamie, Arthur and Alex for being fantastic house mates for three years. There are many other people who have been a part of my life in Leicester, who I'd like to thank as well.

Finally, I would like to thank Saavi for her steadfast support, which got me through the toughest parts of the PhD. My life is better for having met you.

Contents

Declaration of Authorship	i
Abstract	iii
Acknowledgements	iv
1 Introduction	1
1.1 Background	2
1.1.1 Transiting Exoplanets	2
1.1.2 Radial Velocity Measurements	7
1.2 The Next Generation Transit Survey	9
1.2.1 Science Aims	9
1.3 Thesis Overview	11
2 Instrumentation	13
2.1 Instrument Design	13
2.1.1 Telescopes and Mount	15
2.1.2 CCDs	15
2.1.3 Auto-guiding	15
2.2 Operations	17
2.2.1 Field Selection	17
2.2.2 Observations	18
2.3 Photometric Pipeline	19
2.3.1 Input catalogue	21
2.3.2 Reduction	21
2.3.3 Astrometric Solution	21
2.3.4 Photometry	22
2.3.5 SYSREM	23
2.3.6 BLS Search	23
2.4 Summary	23
3 NGTS Reduction with On-Sky Calibration Frames	24

3.1	Introduction	24
3.1.1	Calibration process	25
3.1.2	Calibration Data	26
3.1.3	Reduction Pipeline Overview	27
3.2	Dark and Bias Pipe	28
3.3	Flat Pipe	30
3.3.1	Sky Background	30
3.3.2	Shutter Map	31
	Clouds	35
3.4	Decoupling the Pipeline	37
3.4.1	Combining bias actions	38
3.4.2	Subtracting the bias	39
3.4.3	Subtracting the Sky Background	40
3.4.4	Decoupling the Flat Field	40
3.5	Reduction Data Quality	42
3.5.1	Results	42
3.5.2	Lab Flats	51
3.5.3	Night Flats	51
3.5.4	Simulating the Flat Field	56
3.6	Conclusion	57
4	Correlated noise in NGTS	58
4.1	Introduction	58
4.2	Measuring Correlated Noise	63
4.2.1	Correlated noise and binning	63
4.3	Specific Contributions to the Correlated Noise Budget	68
4.3.1	Flat-Fielding	69
4.3.2	Aperture Placement	70
4.3.3	Airmass	73
4.3.4	Background Correction	80
4.4	Optimisation	82
4.4.1	Background correction	82
4.4.2	Optimisation results	83
4.5	Conclusions	84
5	Bayesian Change Point Detection and Single Transits	85
5.1	Introduction	85
5.2	Change Point Detection	88

5.3	Simulation	90
5.3.1	White Noise Test	90
5.3.2	Transit Injection into Real Data	91
5.4	Data Cleaning and Processing	93
5.4.1	Data Cleaning	93
5.4.2	Removing Correlated Change Points	93
5.5	Results	95
5.5.1	Single Transit Search	95
5.5.2	One-Day Period Search	99
5.5.3	Variable and Transient Event Search	101
5.6	Conclusion	108
6	Photometric Follow Up of NGTS Candidates	110
6.1	Introduction	110
6.2	Candidate Selection	111
6.3	Observations	113
6.4	Reduction and Photometry	116
6.5	Results	121
6.5.1	Transit Confirmations	121
6.5.2	NG0612-2518.44284	123
6.5.3	Performance of the SAAO 1-m	123
6.6	Conclusion	126
7	Conclusion	128
	Bibliography	131

List of Figures

1.1	Mass and radius of well-characterised planets	4
1.2	Mass vs period plot	5
1.3	Artist's impression of a exoplanet planet transit and lightcurve. Photo credit: NASA/Ames.	6
1.4	Comparison of NGTS and SuperWASP photometry	7
1.5	Artist's impression of the radial velocity method. Photo credit: ESO.	8
1.6	NGTS parameter space	11
2.1	Telescopes in the dome. Photo Source: ESO/R. Wesson	14
2.2	NGTS enclosure in Paranal, Chile. Photo source: ESO/R. West	14
2.3	Close up of NGTS telescope, during commissioning in Geneva. Photo source: M. Burleigh	16
2.4	Flow chart of NGTS photometric pipeline.	20
3.1	Reduced and unreduced science frames	27
3.2	Histogram of values for a specific pixel in the flat	28
3.3	Bias, dark and flat calibration masters	29
3.4	Sky background comparison between flat field frames	32
3.5	Sky removal and comparison from flat field frames	33
3.6	Two flat field frames with different exposure times	34
3.7	Standard deviation as a function of exposure time for NGTS flat field frames	34
3.8	An example shutter map for camera 802	35
3.9	Cloud rejection in flat fields.	37
3.10	Example master flat frame.	44
3.11	Flat field error as a function of binning.	45
3.12	Example master bias frame.	46
3.13	Example master dark frame.	47
3.14	Noise versus number of combined dark actions.	48
3.15	Noise versus number of combined bias actions.	48
3.16	Comparison of lab flat fields taken at different wavelengths.	49
3.17	Comparison between a combined lab flat field and a sky flat field.	50

3.18	Sky flat and night flat close-up comparison	52
3.19	Ratio of sky flat to night flat for 801	54
3.20	Absoluted difference between night and sky flat for 801	55
3.21	Numerical simulation of flat-field noise	56
4.1	Transit depth against radius for currently known planets	59
4.2	Transit depth against magnitude for currently known planets	60
4.3	Example NGTS image with an aperture and cross-section	61
4.4	Standard deviations as a function of binning example	64
4.5	Correlated noise and median correlated noise for two field	66
4.6	Correlated noise against star movement	69
4.7	Position of a low-movement and high-movement group of stars on the CCD	71
4.8	Aperture placement error as a function of radius size	72
4.9	Correlated noise as a function of aperture radius	74
4.10	Airmass in NGTS data for a single night	75
4.11	Correlated noise versus airmass	77
4.12	Nightly median values for faint stars	78
4.13	Nightly median values for bright stars	79
4.14	Correlated noise after a median normalisation	79
4.15	Sky and model background comparison	81
5.1	Example of change point methods.	86
5.2	Single transit injection tests in white noise.	90
5.3	Single transit injection in NGTS data.	92
5.4	Posterior probability of a change-point for a test field.	94
5.5	Removing time-correlated change-points from a test field.	96
5.6	Single transit candidate lightcurve	98
5.7	Full lightcurve of a single transit candidate.	99
5.8	Phase folding of NGTS data on a 1-day period	100
5.9	One day transit candidate lightcurves.	102
5.10	Posterior probability cut for variables	103
5.11	Comparison of detected variables	104
5.12	Histogram to select threshold for transient search	106
5.13	Example detected transient events	107
6.1	Candidate vetting example	112
6.2	Master calibration frames for follow-up	119
6.3	Images taken with the 1m at the SAAO	121

6.4	Lightcurves of vetted candidates	122
6.5	SAAO and NGTS photometry comparison	124
6.6	Transit depth versus aperture radius for NG44284	125
6.7	Noise properties of follow-up observations	126

List of Tables

4.1	Correlated noise as a function of how it is measured	65
4.2	Correlated noise properties of NGTS data for different magnitudes	67
4.3	Measured correlated noise for different cameras	68
4.4	Correlated noise for different background subtraction parameters	83
4.5	White noise for different background subtraction parameters . . .	83
6.1	Transit parameters for candidates in first observing run	114
6.2	Transit parameters for candidates in second observing run	115
6.3	Candidate observational parameters for first SAAO run	117
6.4	Candidate observational parameters for second SAAO run	118

Chapter 1

Introduction

Since the first detection of a planet around a main sequence star by Mayor and Queloz (1995) over 3,610 planets have been found in 2704 planetary systems ¹. An astounding twenty percent of these systems contain more than one planet. Over the past two decades many exciting and unexpected discoveries have been made. For example, the existence of so called ‘hot Jupiter’ systems have required that our understanding of planetary formation and migration be re-evaluated. Some of the most significant discoveries lie in the sheer diversity of the planetary systems we have uncovered, including a whole new class of ‘super-Earth’ planets for which we don’t have a Solar system analogue. Though a large parameter space still remains to be explored, these discoveries have helped take an important step towards unmasking the underlying statistical distribution of planetary radii, masses and orbital separations and have brought us closer to understanding the full diversity of planetary systems in our galaxy. Future discoveries will help to unravel the detailed mechanism behind planet formation, helping us to understand not only how our planet has formed, but also how unique it may be. Perhaps in the next few decades we will not only be able to answer questions about what planets are made of, but also whether they may potentially host life.

Within the rapid activity of the exoplanet community, the Next Generation Transit Survey (NGTS) fills a unique niche. NGTS is capable of discovering mini-Neptune and super-Earth sized planets around bright stars, from the ground. NGTS is a fully robotic ground-based survey consisting of twelve 20cm telescopes targeting host stars in the range of 8th to 13th magnitude in V. Bright host stars make detailed follow-up easier. Many of the discoveries made from NGTS will help pave the way for future missions such as TESS² and PLATO³, which

¹<http://exoplanet.eu/catalog/>

²<https://tess.gsfc.nasa.gov/>

³<http://sci.esa.int/plato/>

are set to discover thousands of planetary systems from space. NGTS ambitious design specification will make it the most precise wide field ground-based transit survey to date, achieving an unparalleled precision of 1mmag down to a V_{mag} of 13 with ten minute cadence. This will enable one of the key scientific goals of NGTS, which is to discover enough ‘hot’ and ‘warm’ Neptune’s to have a significant sample in order to measure their bulk density and help constrain the formation history of such systems. In addition NGTS will also find many planets around host stars which are bright enough to be suitable for atmospheric characterisation, helping to expand the number of such known systems, particularly for those in the Neptune size range.

The rest of this chapter is devoted to a very brief overview of the field as it currently stands, with section 1.1 detailing current discovery and follow-up methods, as well as describing a selection of the most important results so far. Section 1.2 will discuss the scientific aims of NGTS and give a brief overview of the current status of the experiment. Finally an overview of the rest of the thesis is given in 1.3.

1.1 Background

What can be learned about an exoplanet is directly related to the method used to discover the planet. Each method in turn has it’s own set of biases and sensitivity, meaning that it is difficult to paint a complete picture of the field without covering them. The most focus will be given to transiting exoplanets as they are the class of planet which NGTS will discover. Also at present the transit method has been the most successful, with approximately 75% of planets having been discovered using these this method. Radial velocity is next, accounting for 20% of all exoplanet discoveries and methods such as direct imaging, micro-lensing, transit timing variations and pulser timing covering the rest.

1.1.1 Transiting Exoplanets

The transit method of exoplanet detection requires that the planet-star system is viewed at an edge-on inclination, such that the planet will occasionally occult the star creating a temporary decrease in brightness. This dip in the flux is typically around 1% for a hot Jupiter and around 0.1% for a super-Earth or Neptune sized planet. An Earth sized planet around a star like our Sun would

produce a dip of 0.01% which is beyond the capability of detection by a ground-based transit survey. The ratio of the flux in-transit to out of transit, f_{in}/f_{out} is approximately equal to the square of the ratio of the radii or $(R_{planet}/R_{star})^2$. Smaller planets produce a larger transit signal around small stars and so NGTS uses red sensitive CCDs (optimised for late-K and early-M type stars). Several transit surveys exist targeting M-dwarfs, such as M-Earth (Irwin et al., 2015) and TRAPPIST (Gillon et al., 2012), however they target single stars. The transit shape is also affected by the inclination of the planet. Planets seen directly edge on, at $i=90^\circ$, will have a deeper transit and shorter ingress and egress times. Transits at lower inclinations will be shallower and have longer ingress and egresses. The stellar limb-darkening will also affect the ingress and egress times, as the star is fainter towards the edge and so the amount of light blocked by the planet will change. Limb-darkening is a wavelength dependent effect and, in general, will affect the shape of the transit less than the inclination.

The probability of a planet transiting decreases rapidly with increased orbital separation and so most transiting planets have small orbital periods, typically in the 1 to 10 day range. Ground based transit surveys are less sensitive to planets with periods longer than about 15 days, due to the gaps in observation during the day and the length of a typical season. Space based missions are able to monitor a field continuously and for longer lengths of time and so may find longer period planets.

The vast majority of planets have been found by the Kepler satellite, which has yielded data of extremely high precision. In particular, Kepler has found many smaller planets which can not be detected from the ground, due to atmospheric effects such as scintillation and seeing. Kepler's key aim of finding an Earth like planet in the habitable zone of a star necessitated looking at one field (105 square degrees) continuously for a long period of time. As a result many of the planets (or planet candidates) are around dim stars and can not be followed up efficiently with current technology. Bright stars are rare. Many Kepler candidates do not have an RV measurement and so we can not determine the bulk density.

Ground-based transit surveys

Transit surveys rely on finding planets whose orbit is viewed edge on from our point of view. The planet will then periodically transit in front of its host star as viewed from our point of view. This can be seen in Figure 1.3. The orbital

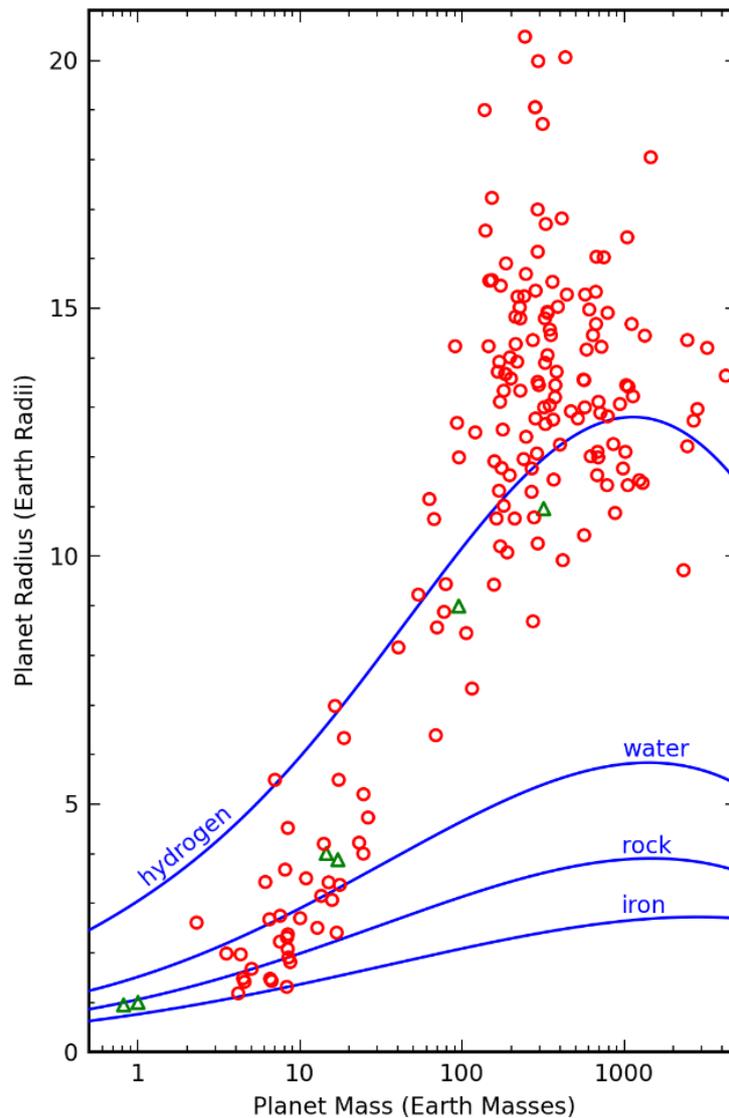


FIGURE 1.1: The figure has been taken from Howard (2013) and shows the “masses and sizes of well-characterised planets”. The blue lines are model mass-radius relationships for a planet composed of only the materials in question. Green triangles are solar system planets. Note the number of “inflated Jupiters” in the upper right hand portion of the image as well as the different possible compositions for lower mass planets.

period can be calculated from the time between transits and therefore the orbital semi-major axis can be derived. Detailed spectroscopic measurements of the star in and out of transit can be used to determine characteristics of the planet’s atmosphere. While similar measurements taken before and after the secondary eclipse (when the planet is occulted) by the star can be used to measure the

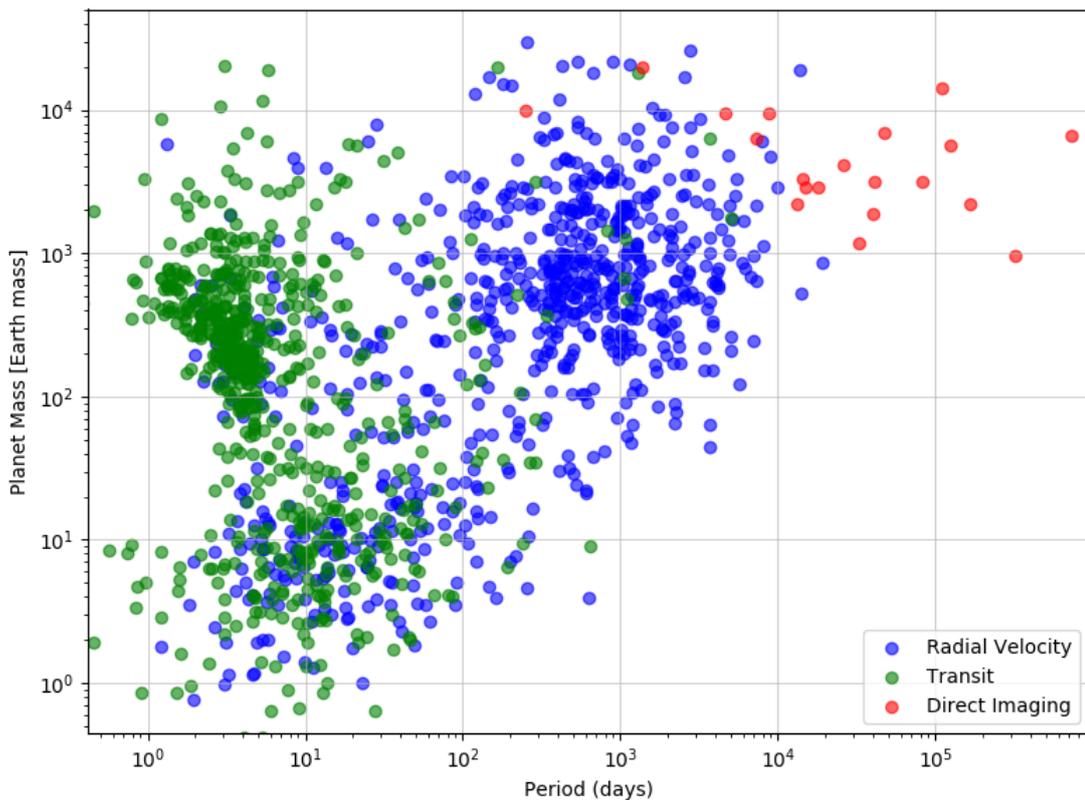


FIGURE 1.2: Mass plotted against period for all planets discovered as of the 20th of June 2018 (with known masses). Three distinct clusters can be made out: so called “hot Jupiters” in the upper left quadrant of the plot, “super Earths” in the bottom left quadrant, and longer period Jupiter size planets from radial velocity surveys. The bottom right quadrant is the frontier of the search for exoplanets and represents the most difficult to find systems.

day time reflection spectrum. By combining radial velocity and transit measurements the bulk properties of a planetary system can be measured, one of the most useful of which is the density. Figure 1.1 indicates the mass versus radius relationship for planets which have been accurately characterised. One of the most interesting features is the myriad possible compositions smaller planets could have. Finding more planets in the ten Earth mass range could lead to a better understanding of the composition of Neptunes and therefore their formation history.

Figure 1.2 shows the mass versus period distribution, for currently known exoplanets with well constrained masses. Figure 1.2 is useful for understanding the biases of the transit and radial velocity methods respectively. The transit method preferentially discovers short period planets, due to the rapidly decreasing geometric probability of transit, with the vast majority of those having periods of

less than 10 days, and the majority of the remaining planets having a period of less than 100 days.

To date the next most productive ground-based survey for transiting exoplanets is SuperWasp (Pollacco et al. (2006)), followed by HAT (Bakos et al. (2002)) in terms of number of planets found, Howard (2013). However the planets found in these surveys are far more suited to RV follow up, compared to Kepler, allowing valuable information about the object mass and density to be determined. This is because they are around bright stars ($V < 12$) whereas many of the Kepler planets are around stars as faint as 16th mag. Many SuperWASP planets have been found in the southern sky and so can be followed up with the VLT and other large observatories. This allows these planets to be studied with tools such as transmission spectroscopy (measuring the spectroscopic signature of the planet as it occults its host star). NGTS has been designed to build on the legacy of SuperWASP by finding even smaller planets which can be studied in the same way. Observations of planetary atmospheres provide a direct way to measure the composition of the outer layers of an exoplanet, but they also help determine what the planet may be made of. The bulk density of a planet is not sufficient to fully constrain planetary composition particularly in the case of smaller planets, which may be made up of a combination of rock, water and ices. Measurement of planetary atmospheres is the only way to assess the chemistry of the system and perhaps link it to the it's initial chemical composition at formation.

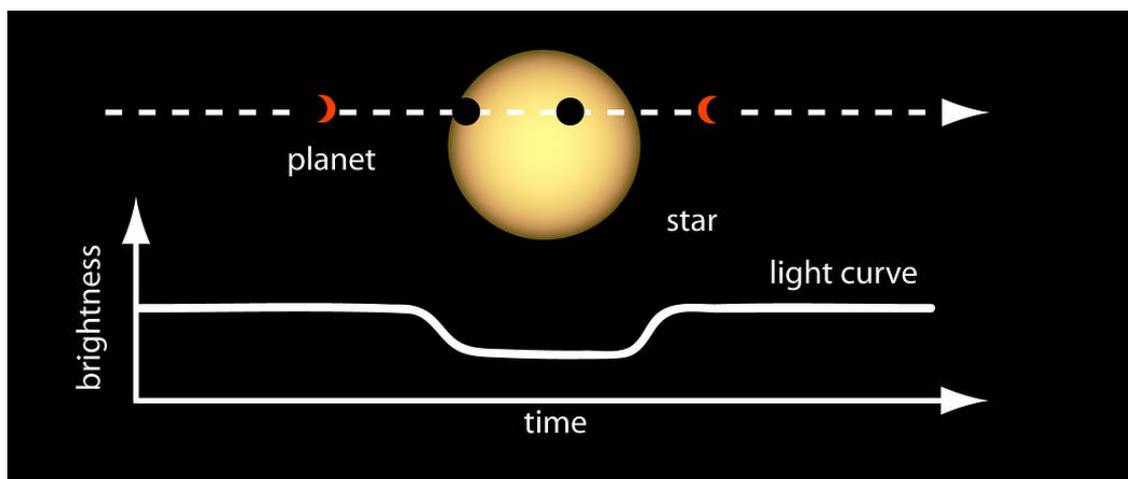


FIGURE 1.3: Artist's impression of an exoplanet planet transit and lightcurve. Photo credit: NASA/Ames.

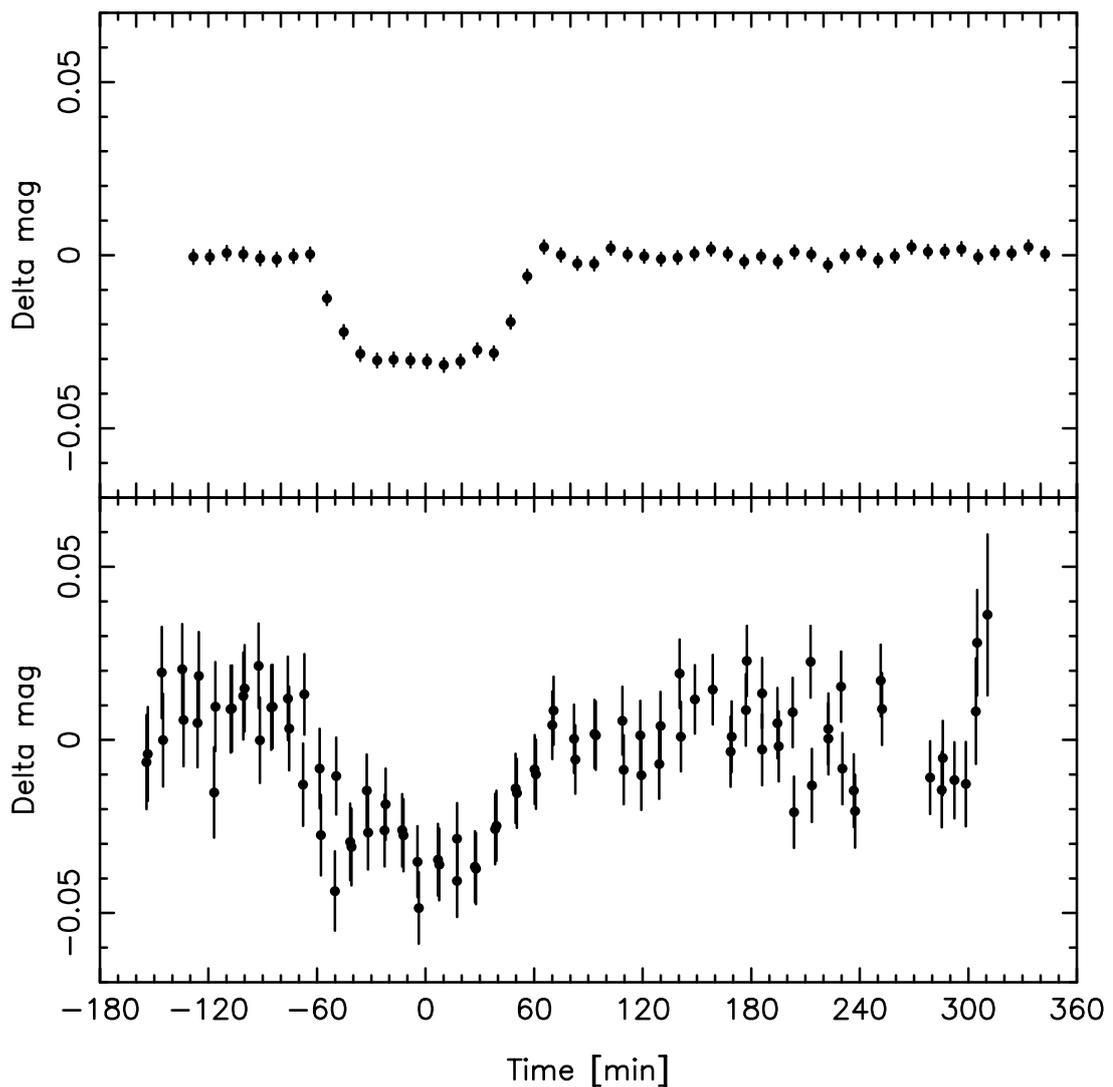


FIGURE 1.4: Comparison between NGTS (top) and SuperWASP (bottom) photometry for a single transit of the exoplanet WASP-4b. Figure taken from Wheatley et al. (2017).

1.1.2 Radial Velocity Measurements

Radial velocity (RV) measurements are measurements of a star's line of sight velocity (Doppler shifts) made by very high resolution, ground-based, echelle spectrographs. Figure 1.5 shows an artists impression of this process. By carefully observing repeating changes in the orbital motion of the star some properties of orbiting planets (or other companions) may be derived. The RV techniques allows a measurement of $M \sin i$, the effective lower limit of the mass, where M is the planet mass and i is the orbital inclination relative to the observer. This will give the mass for NGTS planets as the inclination can be constrained. The period of the orbit and the eccentricity can also be determined by

fitting a Keplerian profile to the RV measurements. The sensitivity of RV measurements is dependent on a number of factors such as, “the observing history, including the number, precision and time of the baseline of the RV measurements” (Howard (2013)) as well as noise from stellar activity and other sources. Radial velocity measurements are required to confirm the existence of a transiting planet, by setting an upper limit on the mass. Combined with the transit modelling, the true mass of the planet can be inferred. Without mass confirmation, the planet can not be validated as the transit signal could be due to an astrophysical false positive such as a background eclipsing binary.

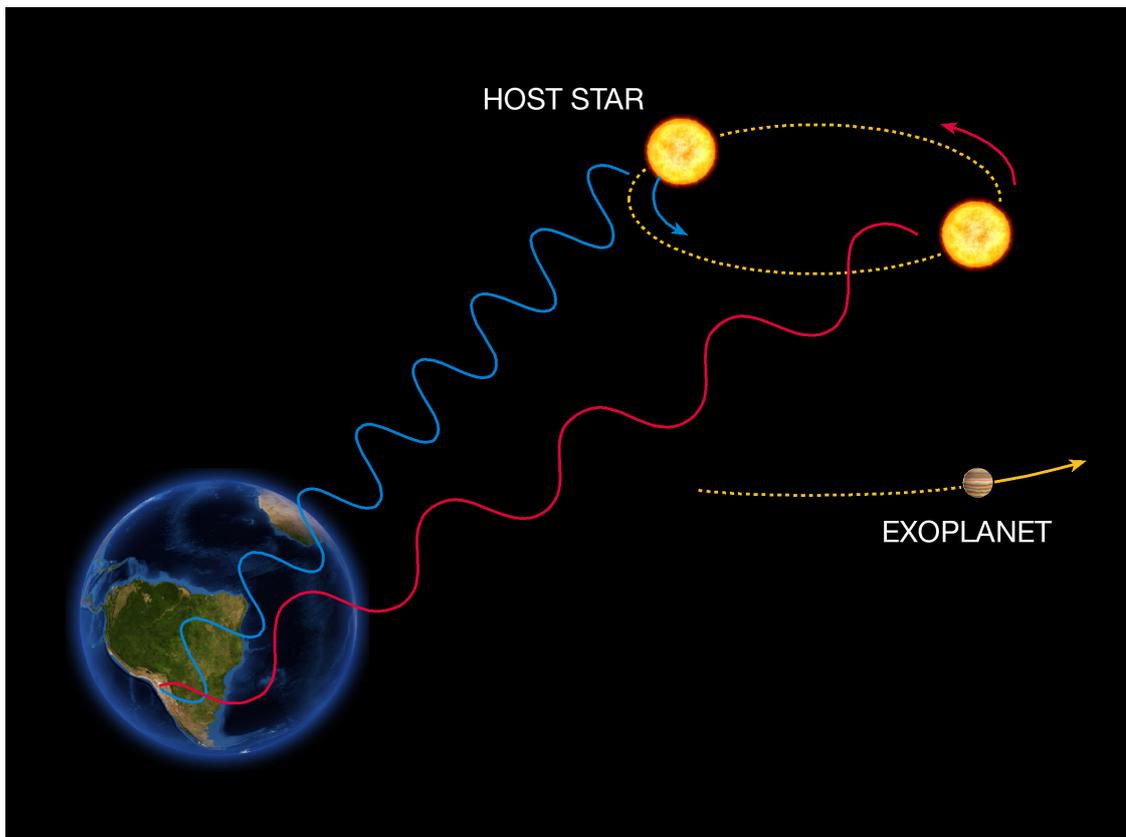


FIGURE 1.5: Artist’s impression of the radial velocity method.
Photo credit: ESO.

The radial velocity method was responsible for the first discovery of an exoplanet around a main sequence star (Mayor and Queloz, 1995). Since then approximately 20 percent of currently known planets have been discovered using dedicated radial velocity surveys. Radial velocity measurements are limited by stellar noise, and the extremely small signals they are trying to detect. Planets are much smaller than their host stars and so a resolution of a few m/s is often need to detect even the biggest planets. The instruments capable of achieving

these kinds of precisions are few and often in high demand, so any observing time granted for the confirmation of NGTS planets has to be carefully used.

Radial velocity surveys are responsible for the discovery of many of the longer period high mass planets currently known. These can be seen in the upper-right quadrant of Figure 1.2. Radial velocity methods are not as constrained by the low probability of transits at high periods, and so can potentially find planets at much larger orbital separations. However, it is difficult to detect smaller planets. For example a 1 Jupiter mass planet at 1AU around the Sun will have a radial velocity signal on the order of 30 m/s, while an Earth sized planet will have a signal of approximately just 0.1 m/s. The latter is impossible to detect with current technology.

1.2 The Next Generation Transit Survey

The NGTS survey consists of twelve fully robotic 20cm telescopes mounted separately, each with an ikon-L camera produced by Andor. The survey is based in Paranal, Chile with some of the best observing conditions in the world year round. As previously discussed NGTS is searching for transiting planets around bright host stars, and is optimised for K and M host stars, due to making use of a red-sensitive CCD. The photometry produced by NGTS is of extremely high quality for a ground-based survey, for example, Figure 1.4 shows a comparison between a single transit of WASP-4b obtained with NGTS and one obtained with SuperWASP. NGTS has been operating for approximately one year and has produced it's first planet discovery (Bayliss et al., 2017) despite unusually bad weather and a rodent infestation (Wheatley et al., 2017). More planet discoveries are on their way.

1.2.1 Science Aims

Background, scintillation and seeing effects from the atmosphere impede the discovery of smaller planets from ground-based transit surveys, due to the additional noise. NGTS would not be able to find an Earth sized planet around a G star, but around an M dwarf it would certainly be possible. While small stars are numerous in the night sky, they also tend to be fainter and so NGTS is red-sensitive to maximise the signal to noise achievable. The main objective for NGTS is to find planets of radius 2 to 5 Earth radii around bright host stars

(Wheatley et al., 2013). By searching only bright stars, we can refine the bulk density of the discovered planets by combining radial velocity and transit measurements. A statistically significant sample of Neptune sized planets found by NGTS could be used to make inferences about the formation history of such planets. Furthermore atmosphere characterisation of Neptune sized planets can help constrain the formation history, for example as shown for HAT-P-26b (Wakeford et al., 2017).

NGTS is also expected to discover several super-Earth sized planets. Super-Earths are a type of planet not seen in our solar system where, as their name suggests, the largest rocky planet is the Earth. However, many other solar systems contain super-Earth sized planets, and understanding their atmospheric composition is a key goal of the exoplanet community. The first evidence for an atmosphere around a super-Earth is that of 55 Cancri e (Tsiaras et al., 2016), a 2.0 earth radius planet around an exceptionally bright host star (I_{mag} of 5.0). Further studies of super-Earth atmospheres have revealed some limited features (Southworth et al., 2017), but discovering more super-Earths which can be studied by transmission spectroscopy would greatly help understanding their atmospheric composition.

In Günther et al. (2017) and Wheatley et al. (2013) the authors perform a detailed simulation of the NGTS planet catch, based on four years of observations with three months of observations per field. The population of planets was then filtered down based on the characteristics of the NGTS prototype, realistic source and background spectra, weather and scintillation for Paranal to give the total confirmable planets that NGTS could find. For radial velocity confirmation using 10 hours of observations with HARPS/HARPS-N this could be 37 Neptunes compared to only 7 from Kepler. Using ESPRESSO on the VLT instead, would mean that for NGTS a total of 231 Neptunes and 39 super-Earths could be confirmed over ten years. Figure 1.6 shows the parameter space which NGTS will search to look for these planets. At a photometric precision of 0.1%, Neptune-sized worlds ~ 6 Earth radii could be found around a Sun like star. Most of the super-Earths would be found around smaller cooler stars. Considering Figure 1.1 and Figure 1.6 it is easy to see how NGTS could contribute to the understanding of super-Earths and Neptune sized planets and their formation history.

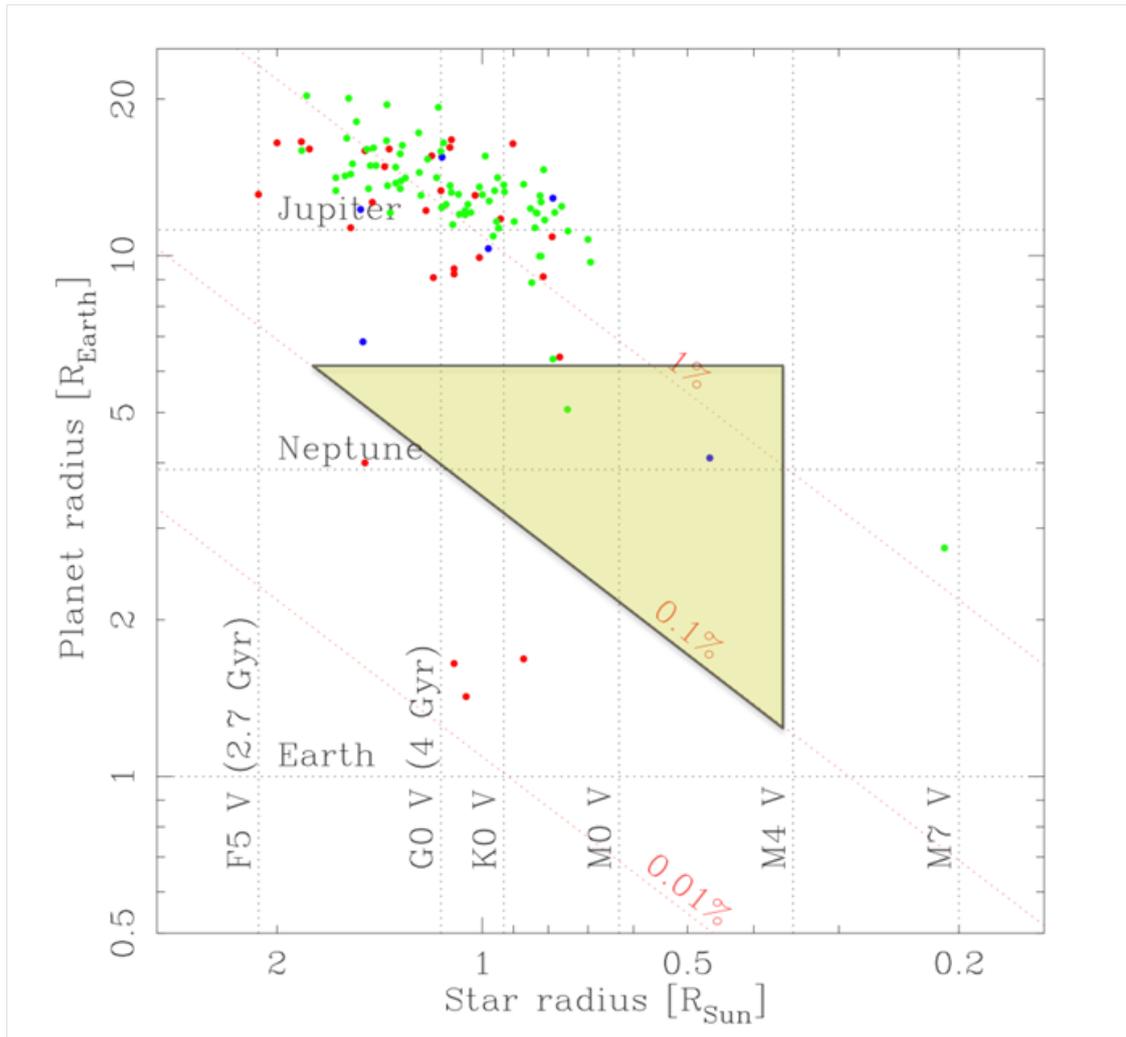


FIGURE 1.6: The figure has been taken from the NGTS website (www.ngtstransits.com). From the website, “known transiting systems are marked in green where they were discovered in ground-based transit surveys, blue if they were originally identified in radial velocity surveys, and red if they were discovered from space.” The figure shows the parameter space which NGTS will search.

1.3 Thesis Overview

The thesis contains four results chapters; two chapters are concerned with optimising NGTS photometry, one chapter is concerned with exploiting the data to search for single transits and variability and finally there is a chapter on performing photometric follow-up of NGTS candidates using the 1-m SAAO telescope. Below there is a brief outline of the four chapters and the key issues which they address.

A brief overview is provided of exoplanet transit surveys and the science goals of the NGTS experiment. The instrumentation covers the instrument design and

specification as relating to the final photometry, as well as the main photometric pipeline and observation strategy.

NGTS uses the standard reduction procedure, sky flat field frames are taken as well as bias and dark frames. The necessity for extreme precision drove the need to develop a high precision reduction pipeline to ensure that no noise was added into the science images as a result of reducing them. The pipeline is developed with special attention paid to the flat field and how best to generate the highest quality master frames possible.

The NGTS survey is expected to reach a precision of 1 mmag down to a V_{mag} of 13. While the system has been designed from the ground up to achieve this precision, the data still has to be properly reduced and processed to produce the best result possible. On the level of parts per thousand, ground-based transit surveys often have the problem of red-noise (or correlated noise) which is can be difficult to fix. Chapter 4 deals directly with diagnostic tests to try and determine which pipeline components are introducing the correlated noise. This is in effect a bottom-up approach for removing excess red-noise as opposed to the top down approach of using detrending algorithms, e.g. Tamuz, Mazeh, and Zucker (2005) and Ofir et al. (2010), which work to ‘blindly’ remove correlated noise.

Within the NGTS data there may be some single transit candidates with relatively long period transits, which are unlikely to be detected by the main BLS (Kovács, Zucker, and Mazeh, 2002) search, due to their relatively long periods. This is only possible due to the high precision of NGTS photometry and the relative short exposure time of 10 seconds. The systems are scientifically interesting, though relatively difficult to find. In this chapter a technique from the family of Bayesian change-point detection (BCP) is applied to a subset of the NGTS dataset in order to search for such transit candidates. The ability of the method to find transit like events is evaluated. BCP has several interesting properties which make it suitable for searching for such transits in noisy data. In addition we also investigate whether BCP techniques are applicable to finding variable stars and other transient events.

High precision photometric follow-up is essential to vet transit candidates and to obtain an independent lightcurve of the transit. The final chapter details the reduction and processing of two seasons worth of high precision photometric follow-up of NGTS transit candidates taken at the South African Astronomical Observatory (SAAO). The full process from candidate selection through to observing is covered including the development of a custom photometric pipeline.

Chapter 2

Instrumentation

NGTS is a ground based transit survey based in Paranal (Chile). The site at Paranal was chosen for its excellent observing conditions which are necessary for achieving the high photometric accuracy needed to detect small planets. The survey goal is ambitious, to provide wide-field photometry at 1mmag in 10 mins for stars brighter than 13th magnitude in the I-band. In order to achieve this the mounts, telescopes, CCDs and optics have been customised to provide the best possible wide field photometry, while drawing on experience with the SuperWASP survey to minimise known sources of correlated noise. This chapter discusses the requirements on the instrument, provides a detailed description of the hardware and discusses the routine operations of the survey. This includes the 'official' NGTS photometry, detrending and transit search pipelines which are used to process all data in this thesis unless otherwise specified. (The reduction pipeline is covered in detail in Chapter 3.) The focus of this chapter is to provide an overview for the reader, while a full comprehensive guide to the instrument can be found in Wheatley et al. (2013) and Wheatley et al. (2017).

Both the prototype telescope and the fully functioning units currently installed on-site in Paranal, have demonstrated 1mmag accuracy in the light-curves. The facility is commissioned and has been operating in survey mode since April of 2016.

2.1 Instrument Design

NGTS is a transit survey and as such has several key constraints which lead the design of the instrument. Due to the relatively low geometric probability of a transit, the survey has to monitor many tens of thousands of stars each night in the correct brightness range to observe a transit. The planets which

are discoverable by NGTS are large and in relatively short period orbits, these systems are rarer than smaller planets in slightly longer orbits (Batalha et al., 2010) and so the more stars in the field of view the better the odds of seeing a transit of detectable depth. Transit depths of Jupiter sized objects around a main sequence stars are on the order of 1% or so, with Neptune and super-Earth sized planets having transits closer to 0.1%. The field of view of the instrument is set by the telescope mirror and the focal length of the optics, so both of these need to be carefully considered.



FIGURE 2.1: Telescopes in the dome. Photo Source: ESO/R. Wesson



FIGURE 2.2: NGTS enclosure in Paranal, Chile. Photo source: ESO/R. West

2.1.1 Telescopes and Mount

The NGTS experiment consists of twelve 20 cm f/2.8 independently mounted commercially available telescopes with an 3 degree wide field of view. With all twelve telescopes working this results in a total field of view of 96 square degrees. Each telescopes can be pointed and guided independently of the others making it possible to observe a wide range of fields at the same time. The telescopes are mounted equatorially in two rows of six telescopes, spaced out so as to prevent them from blocking each other's field of view above an elevation of 30 degrees. Custom built optics have been fitted to keep the point spread function (PSF) to a full width half max (FWHM) of 1.1 pixels across the field of view. The NGTS filter bandpass from 520 to 890nm is specifically designed to optimise sensitivity to late K and early M dwarfs while avoiding the strong water absorption bands above 900nm (Wheatley et al., 2017).

2.1.2 CCDs

Each of the telescopes are mounted with a 2048 by 2048 pixel CCD deep depleted, back illuminated camera. The cameras have improved sensitivity to K and M type stars as they are preferentially red sensitive. The manufacturing technique used ensures that the chips do not suffer from fringing in the red, unlike previous generations of CCDs. The CCD temperature is set to -70C using a 4-stage thermo-electric cooler. Both due to the precision required of the NGTS photometry and as deep depleted CCDs are a new technology, an extensive programme has been undertaken to understand their on-sky performance in the Space Research Center at the University of Leicester. Each of the thirteen CCD cameras, of which one was a spare, have been studied extensively in order to measure the noise properties of the camera and to determine the wavelength dependence of the flat field.

2.1.3 Auto-guiding

Drawing from issues with SuperWASP, Wheatley et al. (2013) have stated that special care has been taken "to minimise known sources of correlated noise related to imprecise pointing, focus and flat fielding." In order to achieve this it is advisable to keep each star on the same pixel of the CCD throughout the

observations of the field. This has been achieved by using a precision auto-guiding algorithm, based on the DONUTS package (McCormac et al., 2013) which achieves a better than 0.1px accuracy across the stars (Wheatley et al., 2017). The auto-guiding works in real-time on the science frames by applying a correction determined using a cross-correlation between the master reference image for the field and the previous science image. Laser guide stars used at the Very Large Telescope (VLT) are detected and masked out before determining the corrections.

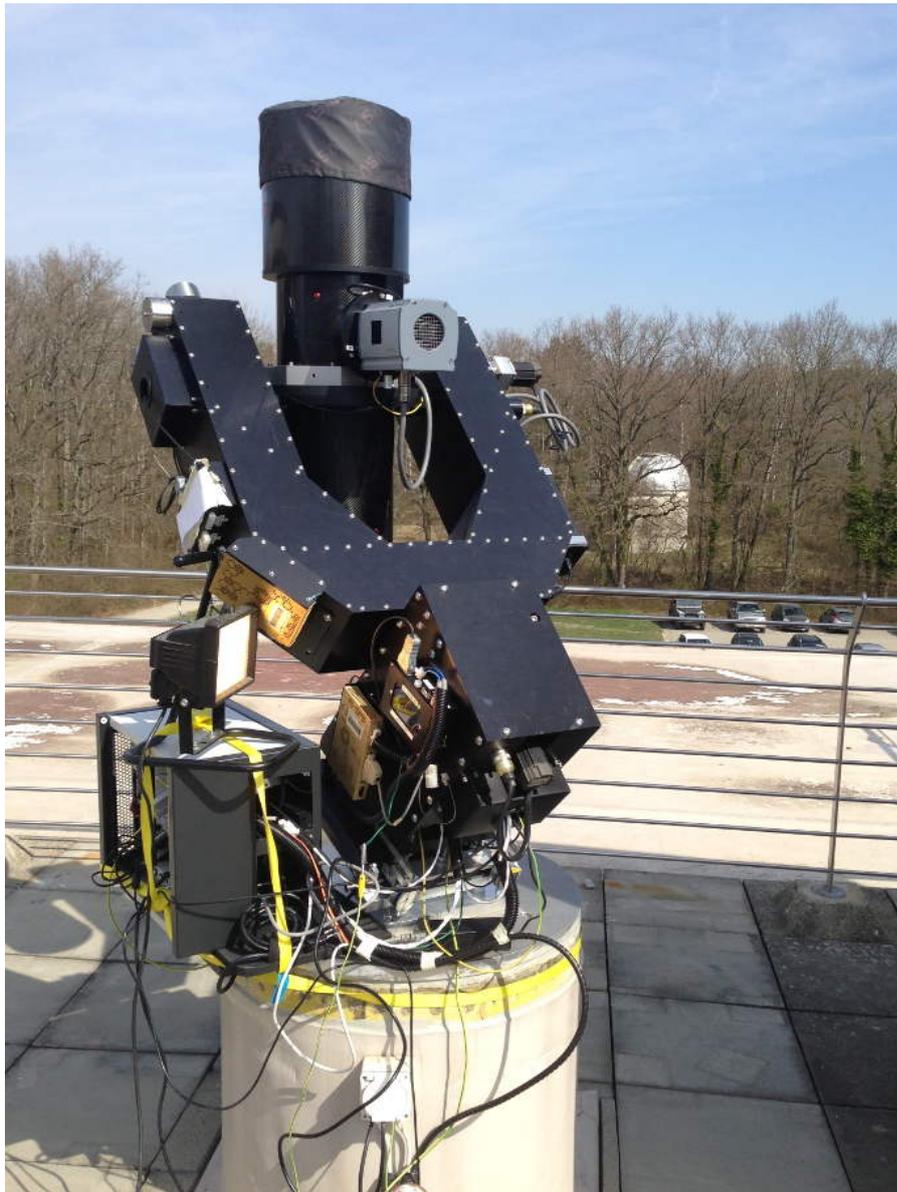


FIGURE 2.3: Close up of NGTS telescope, during commissioning in Geneva. Photo source: M. Burleigh

2.2 Operations

The NGTS telescopes are fully robotic telescopes run from computers stationed in Warwick University. Each afternoon, shortly before dusk, the telescopes are initiated and are left to obtain calibration and science frames over the course of the night. The process is completed automatically, with the dome shutting in the case of bad weather or an earthquake. In this section I will outline the operations of the telescope from field selection through to the collection of the data.

2.2.1 Field Selection

At any one time there are a total of 20 to 24 active fields depending on which telescopes are currently operational. Each telescope will typically observed two fields, one at the beginning of the night and one in the second half. The fields are selected to maximise the potential yield of exoplanets and so can not be too crowded or too sparse. The effect of crowding on NGTS photometry has not been studied in detail, but from past observations fields with 100k stars or so are sufficiently crowded so as to complicate a transit search. In the case of blends it is often difficult to work out on which target the transit is, and the true depth will be diluted meaning that a shallow transit may be missed entirely, and a deep transit may be appear normal. For this reason, fields too close to the galactic plane are avoided (typically > 20 deg on either side).

Fields which are unobservable from Paranal are naturally excluded, as well as those too close to the ecliptic. The remaining fields have then been individually imaged using an NGTS telescope, in a mini-survey. The images can then be used to study the brightness profile of the target stars, and to check whether there are any bright targets in the field which may be saturated. (Of course this is can be done with existing surveys, but it is more convenient to use an NGTS image with the same exposure time as the science frames.) The mini-survey image is then combined with a planet yield estimator Günther et al. (2017), in order to calculate the dwarf to giant star ratios and to work out the potential number of astrophysical false positives.

The fields are also reviewed with respect to existing missions such as the K2 mission and their coverage of the night sky, as well as the future James Webb Space Telescope (JWST). Overlap with K2 is avoided, while finding planets with reach of the future JWST mission will help with the characterisation of those systems. While discovering planets is important, they can only be verified using

radial velocity measurements and so fields must be chosen strategically to fit in with observing schedules at follow-up facilities, producing candidates over a wide range of times.

2.2.2 Observations

Once a set of fields have been selected the telescope will proceed with the observations. Around dusk the dome will be opened to equalise the interior and exterior temperatures. Firstly calibration frames are taken, followed by a series of exposures with which to focus the telescope. Once it is dark and the first field is visible, NGTS will proceed to take the science frames. Typically NGTS will change field once, at some point during the night. Finally at dawn the remaining calibration frames are taken.

Calibration Frames

In order to reduce the data a series of calibration frames are taken following the standard procedure for astronomical data reduction. Bias frames, dark and flat field frames are taken to remove their respective components. A bias frame is a frame taken by reading out the CCD with an exposure time as close to zero as possible. Dark frames are taken with the shutter closed and with long exposure times, typically chosen to match the exposure time of the science frame. Flat field frames are taken by pointing the telescope to a patch of the sky which is expected to be uniformly illuminated after sunset or before sunrise. Typically this is the anti-solar point. However, at 8 square degree, the NGTS field of view is relatively large and the sky is not expected to be perfectly flat. Some residual gradient may be present and so we fit for and remove it downstream. NGTS takes two sets of flat field frames each day, one at dusk and one at dawn.

The dark and bias frames are used to measure how many extra ADU are found in each pixel due to dark current and the bias voltage. They are taken when the dome is closed, to block out all external sources of light which would contaminate the results, and soon after the actual science observations are completed to ensure the ambient temperature is the same as when the science images are taken. CCDs are sensitive to temperature changes. The camera shutter is closed. The dark frames are taken with 10 second exposure times, the same exposure as the science images. The bias frames are taken with the shortest exposure time

possible, just long enough to read out the CCD. The dark and bias frames are taken in the morning after the dome has been shut.

The flat field frames are taken to measure the vignetting and pixel sensitivity of the NGTS cameras. On a typical NGTS CCD this sensitivity ranges by plus minus a percent. While this is only a small amount, as stars drift across the CCD spurious signals can be introduced, increasing the noise floor of the experiment. Flat fields are taken by pointing the telescope to the anti-solar point in the sky at dawn and dusk. The sky illumination here is expected to be most even, however given the relatively large NGTS field of view some residual sky gradient is expected. Frames are then taken with varying exposure times, in the range of 0.1 to a few tens of seconds, with the aim of keeping the counts per frame constant at around 30,000 ADU. While the full well depth of the CCDs is much higher, care must be taken not to increase the counts too much as the CCD may begin to exhibit non-linear behaviour as some of the pixels begin to enter a regime close to the saturation level. Lowering the counts too much on the other hand increases the time wasted, due to high overhead from reading out the frames and slewing the telescope. In between each exposure the telescope is moved slightly so that it is not pointing at the exact same position in the sky to mitigate the effect from stars as the sky gets darker and the exposure times increase. Flats are taken roughly twenty minutes after sunrise and twenty minutes before dawn. At dusk the telescope dome is opened for half an hour before taking flats so that the ambient temperature inside the dome and out can equalise.

Science Frames

Before a field is scheduled for observation, a single nights worth of data is taken to produce a dithered stacked frame of that field. This is used as a reference frame so when the telescopes start observing a field the auto-guiding uses it to ensure that the stars are over the same pixels. The NGTS auto-guider has a precision of 0.1px Wheatley et al. (2017) which is intended to minimise systematic sources of error. The exposure time is ten seconds.

2.3 Photometric Pipeline

The NGTS photometric pipeline processes dark, bias, flat and science frames in order to produce photometry for all targets observed in the NGTS fields. The

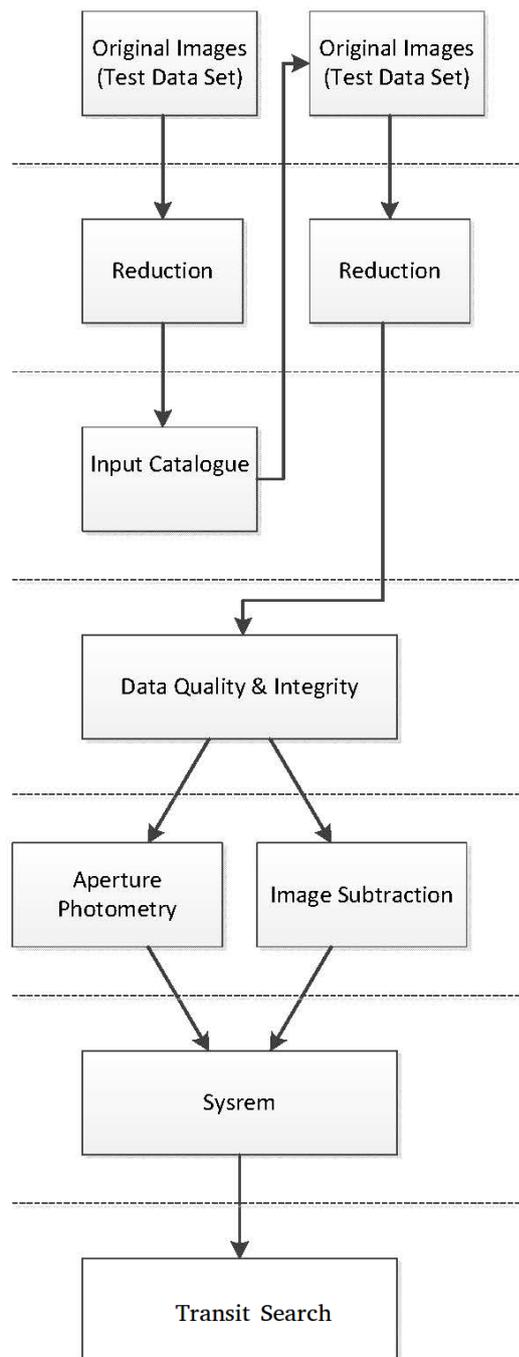


FIGURE 2.4: Flow chart showing the data flow of the NGTS photometric pipeline. Figure is adapted from a plot made by S. Walker.

pipeline is composed of several key components which come together to produce the final photometry for each target, an overview of each is given below. The sections are listed below in order of when which each step is applied to the data. Figure 2.4 shows a data flow chart of the pipeline, given an overview of the entire process. For a complete discussion of the pipeline, refer to Wheatley et al. (2013) and Wheatley et al. (2017).

2.3.1 Input catalogue

At the start of observations for each field a source catalogue is built from the dithered reference frame. For the dithered stacked frame 100 images are taken at low airmass with 10 second exposure times. The images are systematically offset from each other by about 6 pixels. The resulting images are then super-sampled and combined in order to improve the sampling on of the stellar PSF and increase the accuracy with which the position of the stars is determined. (NGTS is undersampled with a PSF of 1.6px.) The source catalogue is then generated using the ‘imcore’ software suite from the CASUTools package¹. The source catalogue is cross-matched with a number of available catalogues and the zero-point of the NGTS images is determined. All sources less than I of 16 are removed. This helps to limit the number of faint objects for which it is unlikely to detect planetary transits (except perhaps in the case of a white dwarf) and for which it would be time consuming to produce lightcurves.

2.3.2 Reduction

Each image is reduced by subtracting the over scan, zero-level bias and the dividing through by a flat field frame. The calibration frames used here are high quality master frames created by combining a seasons worth of data. See Chapter 3 for further details. The science frames are not shutter corrected, as at a 10 second exposure time the shutter travel time does not significantly affect the images. Furthermore the high precision auto-guiding ensures that the stars do not move much more than 1px over the course of a night, which would further dampen any affect the shutter map may have.

2.3.3 Astrometric Solution

In order to perform aperture photometry, stellar sources need to be identified and then the positions of those sources in each image need to be precisely known. To build the source list, a high signal to noise frame is created from 100 images

¹<http://casu.ast.cam.ac.uk/surveys-projects/software-release>

which have been dithered between exposures. The source catalogue is then created by running the ‘imcore’ software from the CASUTools software suite. CASUTools is provided by the Cambridge Astronomy Survey Unit, further details can be found on the CASUTools release page².

Once this is complete, a World Coordinate System (WCS) astrometric solution is computed for each NGTS science image. This is then stamped on the headers of the respective science frames. A solution is needed for each image in order to accurately track the motion of the stars and account for changes in position due to atmospheric refraction and instrumental effects. In addition, the NGTS images show radial distortion, which is solved for by fitting a 7th order zenithal polynomial projection to the images (Wheatley et al., 2017). The polynomial coefficients and centre of distortion is solved for using a custom-built code which employs a Markov Chain Monte Carlo (MCMC) method to find the best fit. The distortion coefficients are re-calculated when any hardware changes are made to the telescope.

2.3.4 Photometry

Once the x- and y-positions for each source have been determined, the image is background subtracted before photometry. The background is calculated by first splitting the image into 64 by 64 pixel boxes. The median of each box is found and then bilinear interpolation is used to create a background map. This map is then subtracted from the image. A 3-px radius aperture is placed over each source and the total flux is measured. Both of these steps are implemented using the ‘imcore’ program from the CASUTools packaged.

The final output of this pipeline is a FITS file, containing information about the source catalogue, image list, flux, background flux, CCD stellar positions and other pipeline by-products. At this stage the photometry is still unsuitable to be used for transit detection as it needs to be detrended first. The lightcurves will contain a significant airmass component, which will dominate the correlated noise, on the order of 10% or so, as well as other systematic issues.

²<http://casu.ast.cam.ac.uk/surveys-projects/software-release>

2.3.5 SYSREM

A custom implementation of the SYSREM algorithm (Tamuz, Mazeh, and Zucker, 2005) is used to detrend the data. In total, three components plus the airmass are removed. Additional systematic noise which is correlated with airmass and the background flux has been discovered in the data. These may arise due to imperfect sky subtraction, flat fielding or aperture placement and are further discussed in Chapter 4

2.3.6 BLS Search

A custom implementation of the box least squares fitting (Kovács, Bakos, and Noyes, 2005) method is performed on each lightcurve to search for transits. This produces a list of candidates which are then vetted to remove as many likely astrophysical false positives and produce a ranked list of targets for further follow-up.

2.4 Summary

NGTS is a ground based transit survey utilizing twelve independently mounted cameras with 9' field of view and deep-depleted red sensitive CCDs to search for transiting exoplanets. A photometric pipeline based on list-drive photometry and CASUTools³ is used, together with custom implementations of SYSREM and BLS to search for transits in the data. NGTS telescopes have been built from the ground up to maximise survey yield and minimise known sources of correlated noise at the hardware level.

³<http://casu.ast.cam.ac.uk/surveys-projects/software-release>

Chapter 3

NGTS Reduction with On-Sky Calibration Frames

3.1 Introduction

The Next Generation Transit Survey is a ground based transit search and so faces an enormous disadvantage in precision when compared to space based surveys, due to the additional scintillation and sky background noise caused by the atmosphere. Nevertheless it is still possible to achieve impressive results from the ground, for example the two most successful ground based transit surveys SuperWASP (Pollacco et al. (2006)) and HATNet (Bakos et al. (2004)), have together discovered more than two hundred exoplanets. As ground based surveys face such a large disadvantage in precision, it is imperative the all aspects of the pipeline are optimised to produce the best possible photometry, including the reduction pipeline. While the effect of the atmosphere can't be overcome on the ground without adaptive optics systems, a well designed and carefully calibrated instrument can still achieve 1 mmag photometry for bright stars (V_{mag} 13 and above), especially over the course of a season of observations.

The NGTS experiment has been designed from the ground up to minimise known sources of systematic errors, known to previous studies of this kind, and to produce photometry which is as close to the technological limit of the telescope and CCDs as possible (Wheatley et al., 2013; Wheatley et al., 2017)). A key part of achieving this precision is the calibration pipeline and the production of high quality calibration frames with which to reduce the data. Poor flat-fielding can result in correlated noise in the light-curves, because the motion of stars across the CCD chip is not random but a complicated function of tracking, stellar colour and airmass. As shown in Pont, Zucker, and Queloz (2006), any correlated noise in the lightcurves will lower the effective sensitivity of the survey and so it is

paramount that it minimised. It provides a strong upper limit on photometric precision as it does not decrease with the addition of repeated measurements.

In this chapter I discuss the end-to-end process for producing high quality calibration frames for each of the NGTS cameras from data taken on-sky over the course of normal science operations. This includes cloud rejection, sky background subtraction and shutter correction, all of which need to be done very carefully to achieve the highest possible precision. I will discuss the noise properties of the bias, dark and flat field master frames and how they affect the data. This is covered in detail in Section 3.5.4 where a numerical simulation is used to gauge the effects of the flat field at a certain precision. NGTS produces on the order of 600-700 calibration frames per day, which require considerable time to process fully. I will discuss the best strategy with which to process the data under the limit of the computational resources available. Finally I will discuss potential avenues for improvement in the reduction process and how they may be implemented in the future.

3.1.1 Calibration process

NGTS science images are reduced according to the standard procedure i.e. a bias, over-scan and dark current component is subtracted from each image, before it is divided through by a high quality flat field which corrects for the optical vignetting and the pixel to pixel variations in sensitivity from the CCD detection. It is useful to consider this process in the form of an equation:

$$\begin{aligned} \text{Reduced_Counts}_{x,y} = \text{Raw_Counts}_{x,y} \times \text{Pixel2Pixel}_{x,y} \times \text{Vignetting}_{x,y} \\ + \text{Bias}_{x,y} + \text{Dark}_{x,y} + \text{Overscan}_{x,y} \end{aligned} \quad (3.1)$$

, where $\text{Raw_Counts}(x,y)$ represents the unreduced pixel values at position (x,y) and $\text{Reduced_Counts}(x,y)$ represents the post-reduction pixel value. So in total there are two multiplicative components and three additive components, each of which need to be subtracted to accurately map the raw counts in the image to a final reduced counts for each target star. This simple model neglects the non-uniform exposure time across the image, a result of the finite time taken for the shutter to open and close. This effect needs to be corrected for when producing the flat field master calibration frame, as the flat fields are taken with a range of exposure times from 0.1 - 30 seconds. The NGTS science image exposure time is selected to be 10 seconds, and so the difference in exposure time across the chip is negligible. (From the centre of the image through to the edge the difference

for a 10 seconds exposure is only approximately 27 milliseconds.) Furthermore due to the precise auto-guiding the NGTS stars move much less than a pixel on average and therefore it is unlikely that there would be any relative differences in brightness due to the lack of shutter correction, unless there is a strong sub-pixel response. So far there has been no evidence of this in the extensive lab testing of the cameras.

The additive component in Equation 3.1 is composed of the bias and dark current. The bias component originates from the extra noise which is introduced while the CCD pixel values are being read out and itself consists of two parts: the over-scan level which is related to the bias voltage applied to read out the CCD, and the zero-level bias which are individual variations in how much extra noise the pixels accumulate. Figure 3.3 shows what the zero-level bias frame looks like. The dark current is due to electrons which accumulate in the pixels over the course of the integration, due primarily to the non-zero temperature of the CCD. In practice, the CCDs are cooled to -50C, so this current is negligible. Figure 3.13 shows a master dark frame built up of 128 actions, or about 25,000 frames, which has a mean value of 0.015 ADU per pixel.

The multiplicative components are the commonly referred to as the flat field, however it is actually composed of two parts, the colour-dependant sensitivity variations from pixel to pixel (in NGTS CCDs by around 1%) and the vignetting produced by the telescope optics. Both these effects need to be calibrated out. The vignetting presents itself as a smooth, rotationally symmetric gradient \sim 10-15% across the CCD chip which decreases as a function of the distance from the image centre. The vignetting is not always centred; the optics may not be perfectly aligned, and so any time the camera is removed from the telescope or the shutter is changed, the vignetting needs to be recalculated. Any changes in the optical alignment might also affect the vignetting. In practice both components are determined at the same time, but it is useful to consider them as separate for the purpose of developing the pipeline and understanding the noise properties of the data.

3.1.2 Calibration Data

Each night of observations produces two sets of flat field frames, one set of dark frames and one set of bias frames as discussed in Chapter 2. These frames are separated into 'actions' which are the discrete building blocks used to structure NGTS observations. The bias and dark frames typically number around 150-200

frames per action. The flat field frames number roughly 150 per actions, with about 70% of these frames having an exposure time of less than 1 second. Flat field frames with short exposure times are dominated by the shutter map, as the shutter opening time of about 27 milliseconds is a significant fraction of the exposure time. On cloudy days flats are usually still taken, if it is expected the telescope will operate later that night, so cloudy flats need to be discarded in post-processing. These actions of bias, dark and flat frames form the basis from which the master calibration frames will be built, with a set of master frames created for each camera.

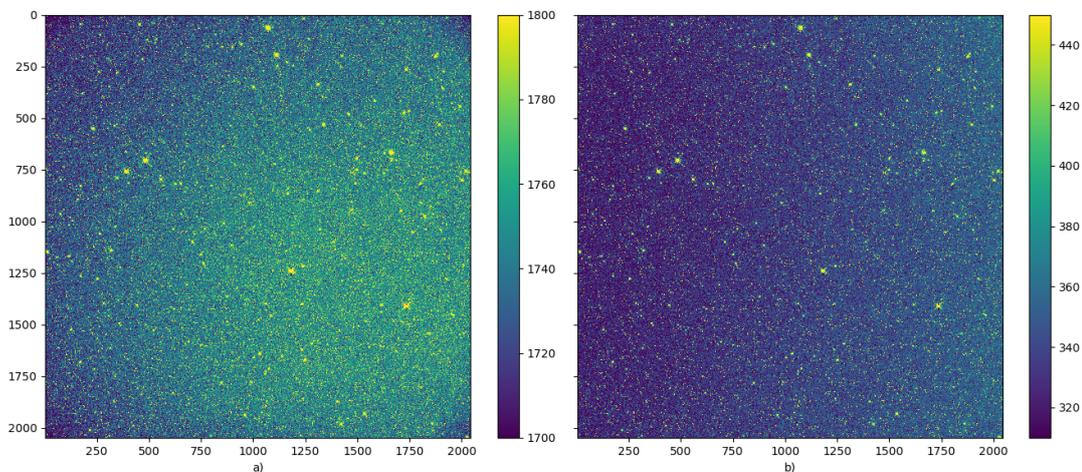


FIGURE 3.1: Comparison between an unreduced and reduced science frame. The image on the right has been bias and over-scan subtracted as well as having been divided through by the flat field. Any residual gradient is due to the sky background in the science image.

3.1.3 Reduction Pipeline Overview

The NGTS reduction pipeline, which was developed by the author, is split into four main components. They are called ‘biaspipe’, ‘darkpipe’, ‘flatpipe’ and ‘metapipe’. The first three pipeline components take as input a list of bias, dark or flat actions respectively. The component’s are designed to pre-process these files to produce per_action calibration masters and other files which are stored for later use. Finally these per_action masters are fed into ‘metapipe’ which combines them to produce per_camera master calibration frames and the error statistics which accompany them for quality control. The vast majority of the processing is done in the first three pipelines, with the combination step in ‘metapipe’ taking a fraction of the time. Further details on the time taken to

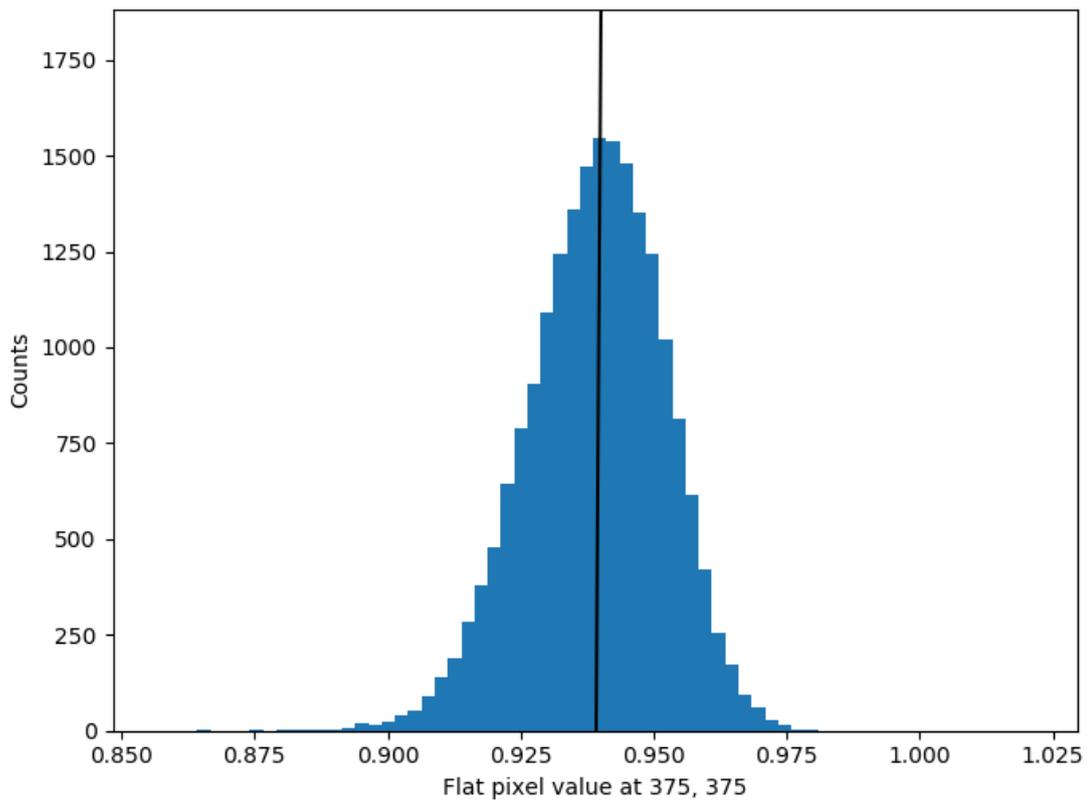


FIGURE 3.2: A histogram of the pixel value 375, 375 for an NGTS flat-field. The values are taken from a seasons worth of data from a single camera. The plotted black line is the mean of the distribution.

compute the reduction pipelines and the implementation are given in Section 3.4. The exact process of combining the calibration frames using the ‘metapipeline’ pipeline is also described in Section 3.4.

Throughout the reduction pipeline all of the calibration frames are built up and combined from the mean of actions below them. Figure 3.2 shows the distribution of values for a single pixel from a flat field over the course of a seasons worth of data. The mean of the distribution is shown in the black line, and proves accurate. Significant outliers due to clouds have automatically been flagged and removed, as will be discussed later in this chapter.

3.2 Dark and Bias Pipe

The process of reducing the science images, begins by creating a master bias frame. The bias pipeline takes in a list of bias ‘actions’, these are sets of bias

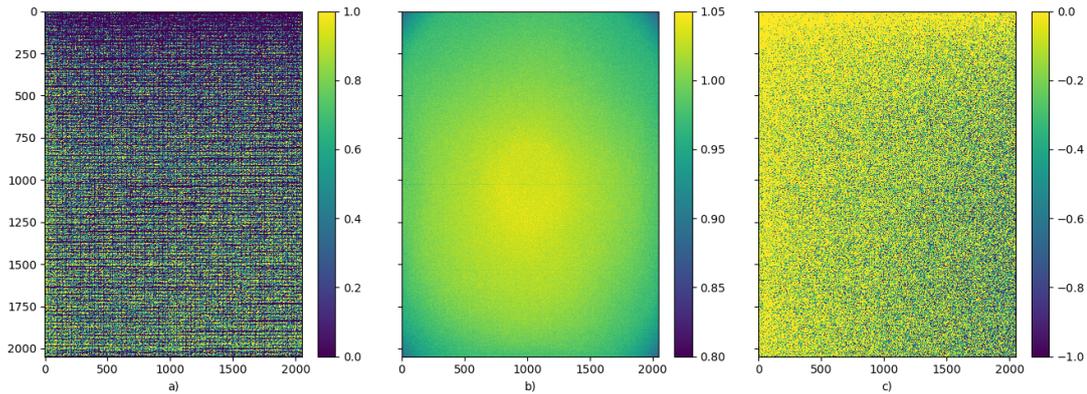


FIGURE 3.3: Shows a master bias, dark and flat field calibration master frame after being processing by the ‘metapipe’ pipeline.

frames taken in one sequence typically after a night of observing, and over-scan corrects each frame from one action. The dimensions of every CCD chip used in an NGTS camera are 2088 by 2048, of which only 2048 by 2048 are used for science. This leaves a strip of silicon twenty pixels wide on either side of the science area which are the pre-scan and the over-scan respectively. Laboratory test conducted by Andrew Grange have shown that of the twenty available columns only those from 2073 to 2088 are used to be used due to charge bleeding from the science pixels. The top four rows of each chip are also ignored due to producing inconsistent readings in laboratory tests (Grange et al. PhD Thesis, in prep).

Once the images from an action have been over-scan corrected they are mean combined and the resulting output is saved. A standard deviation image is also computed and saved. It can be used for diagnostic purposes or to check the data quality. Very occasionally a cosmic ray can hit the camera while a bias frame is being taken. One hundred actions were checked by eye to look for evidence of cosmic rays, and only two frames were found which looked like they had been hit. Producing a master bias from these frames with and without sigma-clipping produced near identical results and so sigma clipping was not included in the bias pipeline, in order to decrease the computational costs of running the code. Sigma-clip implementations available to the consortium typically required producing a copy of the data array to work on, which effectively doubles the memory footprint of the code. A typical action takes up anywhere from 9 to 12 GB and so sigma-clipping is slow as the working memory limitations of the NGTS cluster are reached (16 GB).

The dark frame is created in the same way, but the images are sigma clipped first using a 5 sigma clip to remove cosmic ray hits from the frame. Each dark frame is then over-scan subtracted using the same columns as before, bias corrected

using a master bias frame and then combined together to produce a master dark frame. A standard deviation image is also computed for each master dark frame to help assess the noise properties of that data set.

3.3 Flat Pipe

The process for correcting and applying the flat field is more complicated as the sky flat is a combination of the pixel-to-pixel sensitivity variations, vignetting, sky background gradient and the shutter map function. The sky background has to be removed from each image before applying a separate algorithm which derives the shutter map for the camera. In addition the flat field frames must be carefully sigma-clipped to remove stars from the frames. The flat field frames are then combined and normalised. A sum and standard deviation image is also produced for each action. These are used to as an estimate of the per pixel error in that action as well as to calculate the formal expected error from the number of photons per pixel i.e. the signal to noise. Before any of these steps are undertaken each flat field frame must be bias, over-scan and dark current corrected.

3.3.1 Sky Background

As discussed in Section 3.1.2, flat field calibration frames are taken in two actions, one at dusk before observing and the other at dawn after observing has been completed. It is not expected that the flats should vary between these two actions, however a comparison shows that there is difference in the sky gradient (Figure 3.4). One possibility for this difference may be the temperature difference between the start and the end of the night, however it is unusual that the difference would manifest itself as a gradient across the CCD. Furthermore we would expect much greater night to night variations as the temperature is variable. While taking flats the telescopes are slewed so that they are pointing near the anti-solar point in the sky, however this position is not fixed over the year as the position of the Sun at dawn and dusk changes. A more plausible explanation is that the sky background gradient is slightly different due to a different relative pointing of the telescope at dusk and dawn when taking flats.

Further investigation shows that the gradient is varying as a function of time over a period of a year, lending evidence to the idea that the background variations are due to a difference sky gradient. In terms of the reduction pipeline, this extra component is a problem as we don't expect it to be in the science data. Sky gradients in the science images are due to proximity to the moon or light pollution and not the Sun. The vignetting function is spherically symmetric, so the excess gradient can be removed by fitting and subtracting a linear trend across the image, i.e. $z = a \times y + b \times x + c$. Figure 3.5 shows the ratio of two flats, taken at dawn and dusk, after the linear trend has been subtracted.

3.3.2 Shutter Map

The NGTS cameras have an iris shutter which creates the characteristic petal shape seen in Figure 3.6. The shutter map is a result of the finite travel time of the camera shutter across the image. Each pixel in the camera then has a unique exposure time associated with its location, which when uniformly illuminated as in the flat field, can be measured. In practice the shutter map is only significant for very short exposures, as typically the total shutter time is on the order of 30 milliseconds. This becomes significant in the case of the flat field as each action contains many short exposure frames and a few long exposure ones.

Each flat field action contains frames with exposure times ranging from 0.1 to 30 seconds. The exposure time is chosen in such a way as to keep the mean counts in the frame equal to 30k. So one solution is to only use frames of exposure time greater than a second, however in practice this means losing many frames and not making use of the full statistical power of the flat field frames available. A second solution is to attempt to measure and remove the shutter map from each frame before combining them and making the final master calibration flat. The drawback to this method is that unless the shutter map is measured to an extremely high accuracy, about 1 in 10,000, we risk introducing unnecessary error into the master flat. The solution is to combine many shutter maps taken across a season of data and to use this to calibrate each flat field frame before combining them. This is the approach used in the NGTS calibration pipeline. Each action master frame is "post-corrected" for the shutter map before being mean combined in meta-pipe. The procedure for this is described in Section 3.4.4.

The shutter map is derived using the (Surma, 1993) algorithm for shutter free flat fielding. While a laboratory shutter map is available (Grange et al. PhD

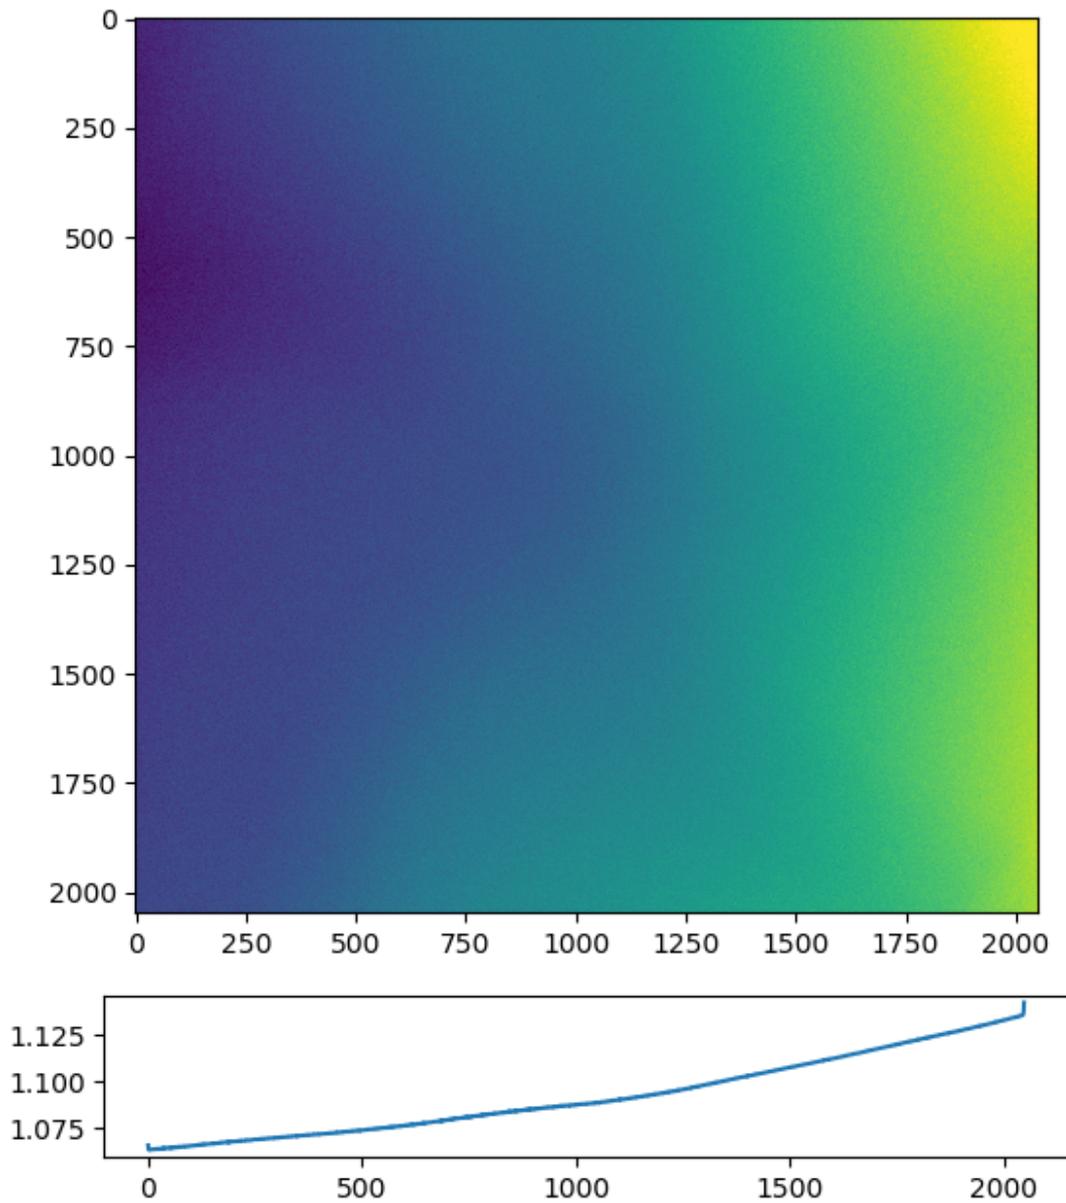


FIGURE 3.4: The plot shows the ratio between to flat fields from two different actions, one taken at dawn and the other at dusk. This ratio is taken before background subtracting the individual flat field.

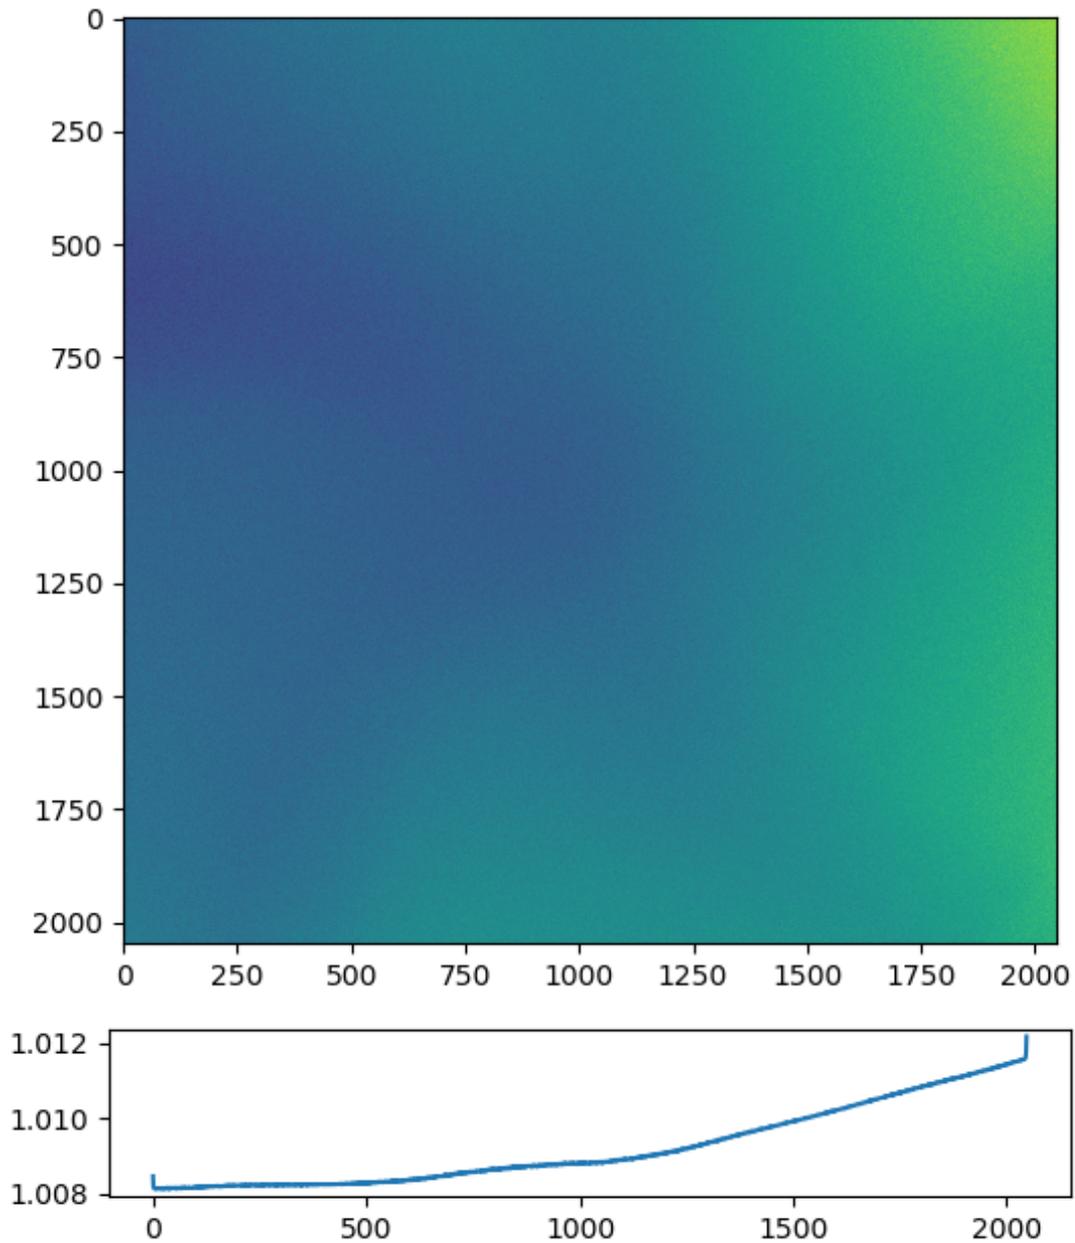


FIGURE 3.5: The plot shows the ratio between two flat fields from two different actions, one take at dawn and the other at dusk after subtracting the sky background ratio. While there is some residual gradient present in the ratio, it is much smaller than before the background subtraction.

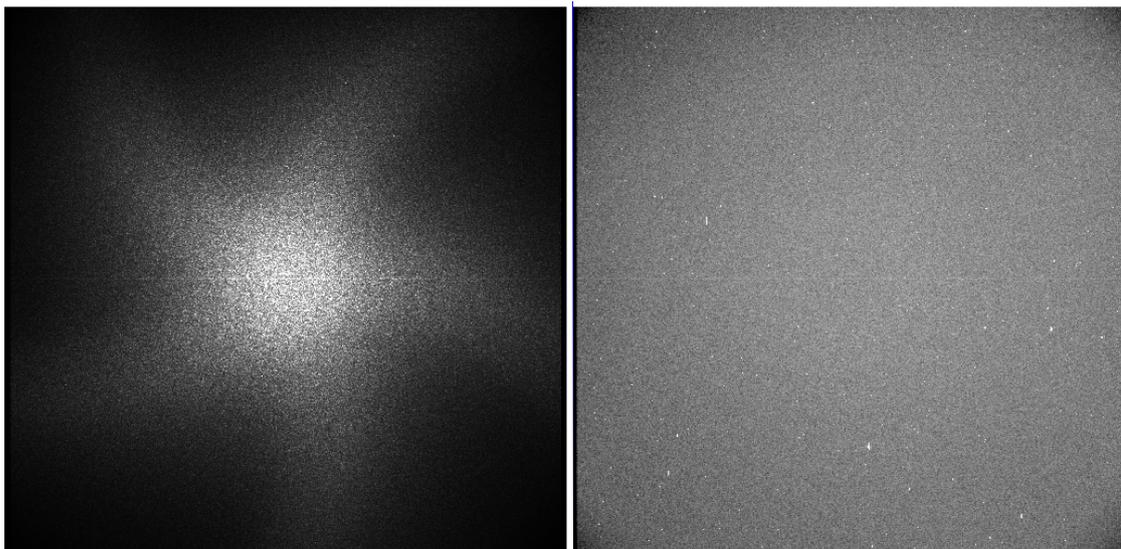


FIGURE 3.6: Two flat field frames with exposure times of 0.12 and 30.3 seconds respectively. The scale of the short exposure image has been modified to emphasize the shutter map.

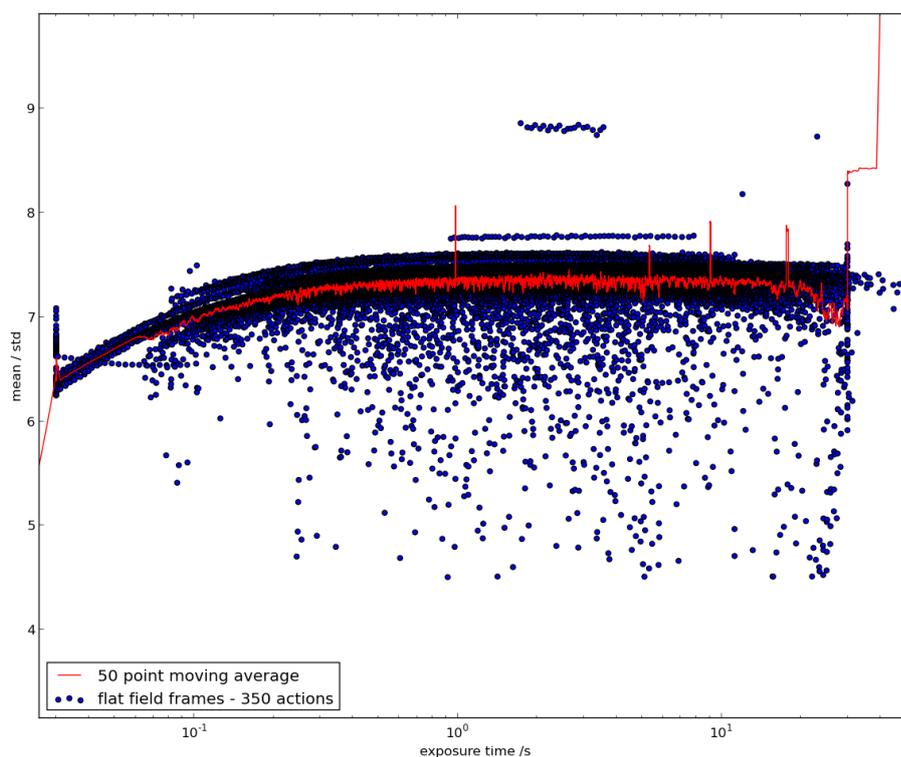


FIGURE 3.7: A plot of the mean over standard deviation shown against exposure time for a few thousand flat field frames from camera 802. Outlying points above the main curve are due to flat field frames which were taken with 15k and 25k counts as opposed to the 30k normally used in operations.

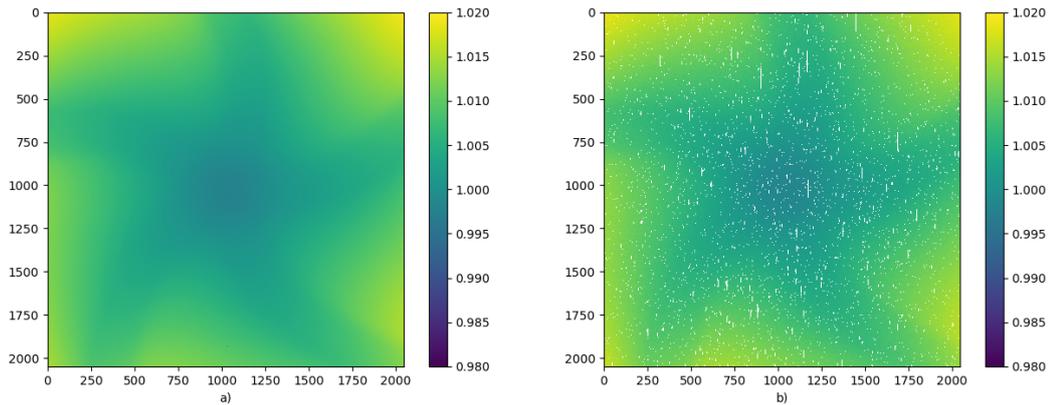


FIGURE 3.8: Figure a) is a shutter map made from a seasons worth of actions for camera 802, while b) is made from a single action. The white streaks are due to stars which have been sigma-clipped out of the flat field frames.

Thesis, in prep), the one used for data reduction has to be derived using the on-sky calibration data, as the optical path is different when the camera is mounted to the telescope. The Surma algorithm derives the shutter maps by treating the problem as a linear regression and fitting each pixel as a function of exposure time. One minor complication with this procedure is that long exposure flat frames often have stars in them which need to be sigma-clipped out of the frame as they would affect the shutter calculation. Replacing each sigma-clipped value with the mean of it's frame effectively remove any artefacts caused by stars. In practise this is unnecessary, as clipped pixels can just be ignored until meta-pipe where there are sufficient shutter maps to combine together which effectively removes these pixels.

Clouds

Sky flat fields have the problem that they are occasionally taken on days which are cloudy or partially cloudy. As NGTS is operated remotely, many frames are taken with clouds in them. A clear flat field is relatively uniform, while those with clouds in them show a greater range of variability in intensity. Around July 2016, a total of 562 flat field actions were available for evaluation for 5 different cameras. Of these a subset of three hundred and seventy actions are chosen at random to be used for testing, containing a total of 39,059 flats. Of these 370 actions, 330 show no evidence of cloud, eight show some evidence for cloud (a few flats) and twenty nine have clouds in most frames. This was determined by checking the frames of each action visually. The vast majority of cloudy flats are

in the twenty nine bad actions, these are the actions we would like to remove. Three actions have between 1 to 3 bad frames, however as we are typically averaging \approx two hundred frames per action, and then hundreds of actions in turn, the few cloudy frames can be safely ignored. Any cloud detection method can therefore focus on identifying clouds on a action basis, as opposed to in individual frames.

Visual inspection of 200 frames per camera per day is time consuming and so it would be useful to have an automatic method for detecting cloudy actions and preventing their corresponding action master from being combined together to create the master calibration frame. Several different statistics were tested to try and determine whether an individual flat frame contained clouds or not, however none were particularly effective. The variety of clouds made it difficult to reliably detect them simply. Some clouds were fairly uniform, while others were highly irregular and showed large variations across the flat. A simple machine learning classifier could have been built to detect clouds, however this would likely have been an unnecessarily complicated solution. Most actions are either cloudy or not, and so we can examine the standard deviation image of each action: actions with cloud tend to have much higher standard deviations than expected compared to non-cloudy actions, so we can use this as a simple classifier to reject unusable actions. Using the full action overcomes the issues of trying to identify clouds in individual frames.

The cloud rejection procedure is implemented by taking the mean of the standard deviation image for each action. An iterative sigma clip algorithm is then applied to this data to remove outlying actions. A sigma of 1.0 is used and removes most cloudy flat action. The remaining actions are then used to create the master flat for each camera. Figure 3.9 shows the average standard deviation per pixel of a selection of action masters for camera 812, as well as the rejection criteria used. In this instance 30 actions were rejected out of a possible 278 actions or just over 10%. This is a fairly typical number of actions to be rejected, though of course the number of cloudy nights has strong seasonal variations.

The typical number of ADU per pixel in the flat field is approximately 28,000, which for camera 812 with a gain of 2.0 converts into 56,000 photons per pixel. The expected error in ADU per pixel is then approximately 175 ADU which is substantially lower than that seen in Figure 3.9. The standard deviations are calculated per action, before the flats have had the shutter map or sky background removed. This additional signal is likely the main extra noise contribution. The bifurcation of the standard deviation signal is due to the increased changes in

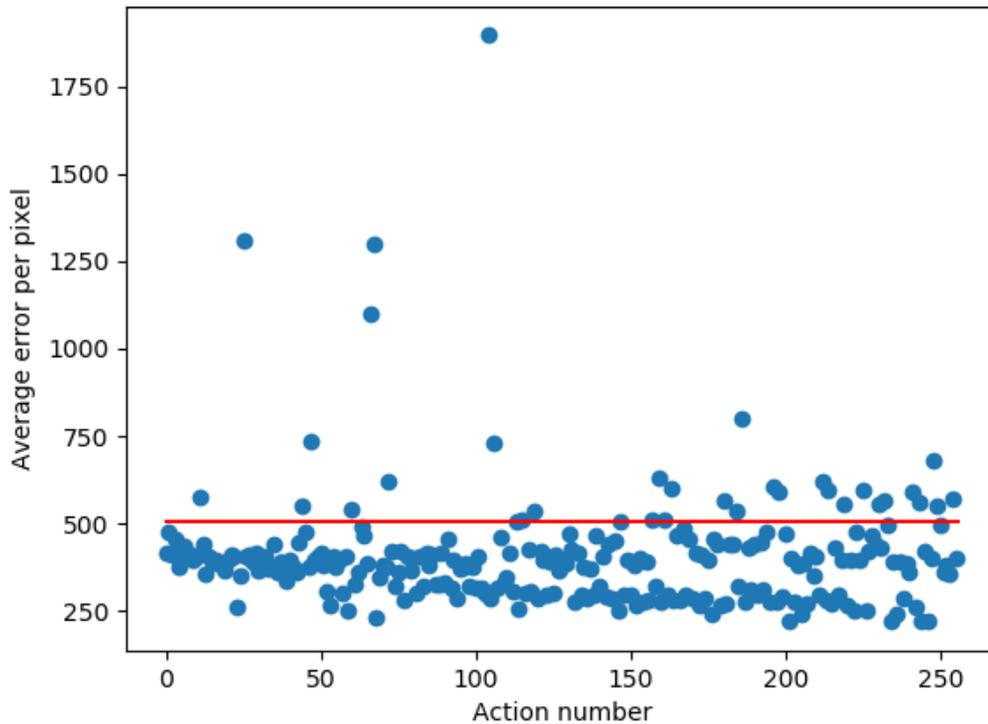


FIGURE 3.9: Rejecting cloudy actions using a simple sigma clip. The plot shows the average error per pixel for 278 actions, of which 30 are singled out as being due to clouds. The red line shows the cut-off chosen using a sigma clip of 1.0 standard deviations.

the sky background between dusk and dawn as the seasons change and further indicates that the sky background is a source of error.

Despite these additional noise sources, cloudy actions can still be clearly identified, and visual inspection of the cloud free actions shows that the criteria works well. Of the 248 remaining actions none had cloud in more than a few single frames. In future the standard deviation clip maybe updated to work on only long exposure frames, hopefully making the cloud removal even more sensitive to noisy flat field actions.

3.4 Decoupling the Pipeline

At the time of writing NGTS has been operational in some capacity for over a year and a half and has produced a huge amount of calibration data. Currently there are 5579 bias actions, 6232 dark actions, 10251 flat actions across each of the

twelve cameras. This is an enormous amount of data considering that each dark action is approximately 250 MB, each bias action 350 MB and each flat action 750 MB per camera. In order to optimally run the pipeline it is necessary to ‘de-couple’ the various pipelines components to some degree i.e. we can not be recalculating the master bias from scratch each time we add a few new bias actions. Similarly we don’t want to recalculate the flat field from scratch when adding a few new flat field actions, however updating the flat as we get more data would be useful to improve the signal to noise as much as possible. This is particularly important for tracking features such as dust.

Furthermore we don’t want to have to re-process all the flat or dark data every time we improve the signal to noise on the master bias. Each pipeline component is designed to work completely independently from each other component so that they may be run as and when needed without having to reprocess large sets of data. Below I discuss how the flat components have been split apart and demonstrate that it doesn’t matter whether the calibration frames are calculated in a de-coupled or traditional coupled way. It is important to note that this procedure only works once the over scan level for each reduction frame has been subtracted. This is done for each frame as discussed in Section 3.2.

3.4.1 Combining bias actions

For the very first test we would like to assess the best way of combining multiple bias actions together to improve the signal to noise. For each action we compute the mean (which is also used for data quality monitoring), however the mean of the action means is not equal to the global mean. To get around this we use a weighted mean, where each action is weighted by the number of images it contains. This can be trivially be shown to be equal to the global mean of the full dataset, considering the values for one pixel say at position 100, 100 on the CCD:

$$\bar{b}(100, 100) = \frac{b_1 + b_2}{w_1} + \frac{b_3 + b_4 + b_5}{w_2} + \dots + \frac{b_{n-2} + b_{n-1} + b_n}{n} \quad (3.2)$$

Each individual term represents the mean of an action. It is trivial to show that by weighting these by the number of images per action, w_i and taking the mean, using $n = \sum w_i$ we get the global mean:

$$\bar{b}(100, 100) = \frac{b_1 + b_2 + b_3 + b_4 + \dots + b_n}{n}$$

While this may work in theory, it is necessary to demonstrate that it is true in practice to verify that the bias frames are being combined properly and that we are not introducing systematic noise through incorrect manipulation of the action means. A test is conducted using 391 bias frames split over five different actions. A master bias is constructed in the usual way first by combining all the bias frames together, and secondly by combining the bias frames using a weighted mean of the action masters. The two frames are subtracted from one another and the mean difference is calculated to be -1.17×10^{-19} which is clearly down to the numerical precision of the calculations. The difference is uniform throughout the images and shows no structure as is expected.

3.4.2 Subtracting the bias

The second part of the decoupling of the pipelines requires that the bias image is subtracted from the mean dark or flat action frames while they are being combined in metapipe. This makes it very easy to change the master bias frame in case an issue is found at a later stage and it also ensures that the dark and flat frames can be rebuilt quickly using a better quality bias master as more on-sky calibration data is collected. The bias frame is an additive component and so is very easy to take out of the dark frame, which I'll show below, and only slightly more complicated for the flat field. If we consider a single pixel in the master dark action we see it is composed of the sum of the bias and dark components,

$$\bar{d}a(100,100) = \frac{d(100,100)_1 + b(100,100)_1 + \dots + d(100,100)_n + b(100,100)_n}{n},$$

where $\bar{d}a$ is the mean of the dark action, d is the dark master and b is the zero-level bias current. Dropping the (100, 100) index we can rewrite the equation as,

$$\bar{d}a = \frac{d_1 + d_2 + \dots + d_n}{n} + \frac{b_1 + b_2 + \dots + b_n}{n} \bar{d}a = \bar{d} + \bar{b}$$

So to correct for the bias in the mean dark action we simply need to subtract the master bias from the master dark at the end of the process. To double check that this does indeed work and that no systematic errors are introduced into the data, perhaps due to a component which has been left out, we run a simple test by creating a dark master frame by subtracting the bias frame first and then create one by subtracting the master bias frame afterwards. The mean difference

between the two master frames was 4.9×10^{-19} which is again consistent with numerical error and shows that this works as well as expected.

3.4.3 Subtracting the Sky Background

The sky background gradient appears in all flat fields, despite precautions taken to point the NGTS cameras at the most uniform part of the sky, the anti-solar point. (The NGTS cameras have a relatively large 2.8×2.8 degree field of view.) As the sky background residual appears as a gradient across the flat (Figure 3.5 it can be removed by fitting out a linear trend from each frame. Some residual may remain in the flat frame however it is difficult to separate from the vignetting. Since a linear trend is being removed from each frame, it may be possible to remove the gradient from the final action mean frame instead of removing the gradient from each individual frame.

To verify that the sky removed in both situations is the same, a comparison is made between an action master which has had a sky estimate remove from each frame and one where the sky has been removed from the final master frame. The two frames are subtracted from each other and then the mean is taken of the difference frame, which is 1×10^{-17} indicating that the two frames are identical to within numerical error. The test has been undergone with the usual shutter and bias correction used in one frame and the bias and shutter correction done after the fact for the second frame.

3.4.4 Decoupling the Flat Field

$$flat = sky * vignetting * pixel2pixel * (1 - shuttermap/exptime) + bias + dark + overscan \quad (3.3)$$

The flat field can be regarded as a combination of a number of different components as shown in Equation 3.3. A number of these terms can only be estimated for the particular frame in question, these are the over scan level, and the sky level and the sky gradient. The sky gradient is removed as discussed in section 3.4.3. The rest of the terms will be common to all flats taken with the same camera, though the vignetting and shutter map can and do change over time. The dark term is taken to be negligible and so is ignored in the following analysis. The overs can is subtracted as discussed in the bias pipe section, while the sky

level is taken to be the median of the frame, and the sky background is removed as discussed in the flat pipe section. So the flat equation would look like,

$$f = V * P * (1 - SH/t) + B,$$

where f is the flat, V is the vignetting, P is the pixel-to-pixel variations, t is the exposure time, SH is the shutter map, B is the bias and S is the sky brightness level or the frame median. Ignoring the bias component for the moment we have:

$$\begin{aligned} \bar{f} &= \frac{V_1 P_1 (1 - SH_1/t_1) + V_2 P_2 (1 - SH_2/t_2) + \dots + V_n P_n (1 - SH_n/t_n)}{n} \\ \bar{f} &= \frac{VP(1 - SH_1/t_1) + \sigma_1 + VP(1 - SH_2/t_2) + \sigma_2 + \dots + VP(1 - SH_n/t_n) + \sigma_n}{n} \\ \bar{f} &\approx VP \left[\frac{(1 - SH_1/t_1) + (1 - SH_2/t_2) + \dots + (1 - SH_n/t_n)}{n} \right] \\ \bar{f} &\approx VP \left[\frac{n - SH * (1/t_1 + 1/t_2 + \dots + 1/t_n)}{n} \right] \\ \bar{f} &= VP \left[1 - \frac{SH}{1/\bar{t}_n} \right] \end{aligned} \tag{3.4}$$

So we can correct the flat action means using a computed shutter map and the mean exposure time for that action. This can be done after subtracting the bias term as shown in the equation above. The flat fields are taken by varying the exposure time to reach some fixed number of counts. So we expect the median sky level for each frame to be distributed in some roughly normal way, though the error will be much larger than for the deviation of the flat field. The camera points to a similar position on the sky throughout the action, though of course there is some small movement as the telescope is slewed to avoid stars. Under the assumption that the sky gradient does not change very much and the sky level is normally distributed the errors should be relatively small.

Nevertheless, it is necessary to investigate how small the final error term is, which can be done numerically by computing the flat normally and computing the flat from the action mean products. Two tests were conducted using an action comprising just 3 frames and a full action comprising 110 frames. The mean difference between the master flats in the former case was approximately 4 ADU per pixel, while for the latter it was 0.291 ADU per pixel. The frame median is

23,370 ADU and so this represents a small error, of 0.012 mmag, which will decrease further as several hundred flat actions are combined to produce the final master flat.

3.5 Reduction Data Quality

3.5.1 Results

The previous sections of this chapter describe the structure and design of the NGTS reduction pipeline, as well as how it is practically run on the NGTS data. Chapter 2 describes how the reduction data is acquired for each camera, on-site in Paranal. This section describes the final master calibration frames produced for each camera, and quantifies their quality. Particular focus is paid to the flat field frames which, due to the fact that they are divided out from each science image, are expected to contribute the most to the NGTS error budget.

Flat Frames

A test set of 136 consecutive actions, taken with camera 806, are selected as a test set to investigate the final quality of the master calibration frame. The actions were processed first by ‘flat-pipe’ and then combined using ‘meta-pipe’ as described in section 3.4.4. After processing with ‘meta-pipe’ eight actions were removed by the cloud rejection filter. The remaining 128 actions were combined to produce a single master flat field calibration frame. This calibration frame can be seen in detail in Figure 3.10.

In order to determine the quality of the final flat, and also whether the flat pipeline has produced the best possible frames it is necessary to quantify the error per pixel for the calibration master. This can be worked out theoretically as well by simply counting the total number of photons and then calculating the expected shot noise. The error can be measured experimentally by comparing the average standard deviation per pixel for the flat action masters as they are combined together. Figure 3.11 b) shows an example of the series of values measured for the pixel (1010, 1010) for camera 806. As more actions are combined together the standard deviation of the time series decreases.

3.11 a) shows the standard deviation for the entire frame as a function of the number of actions binned together. A shot noise model is added to the plot

based on the total number of photons collected over all the actions. We can see there is good agreement between the measured standard deviation and the expected noise levels from the model. The measure noise is slightly higher, which could be due to an excess shutter map component from the shortest exposure frames or perhaps some due to thin clouds which were not picked up by the filter. Overall the agreement indicates that the flats are of high quality and can be used in the final NGTS photometric pipeline.

Bias and Dark Frames

A similar process is used to determine the quality of the bias and dark master frames, which are shown in Figures 3.12 and 3.13 respectively. A clear gradient can be seen across the bias master frame, which is typical for all NGTS cameras, in addition to the fine 'lattice' structure through the frame. The dark frame similarly shows a gradient across the chip, with some additional horizontal structure.

The average standard deviation or error per pixel is measured as a function of the number of actions combined together. The values can be seen in Figures 3.14 and 3.15. The values measured for 128 combined actions are 0.051 ADU per pixel and 0.065 ADU per pixel on the bias and dark calibration frames respectively. The cumulative effect on the formal error on the flux, for a 3px aperture, is then equal to 1.44 and 1.84 ADU respectively. This is a negligible amount and a much smaller contribution than the flat field, background or airmass correction. Further more due to the fact that the bias and dark frame are subtracted from each image, the net effect of the error would be to add a small constant offset to the flux. As the aperture drifts this offset will change slightly introducing an error which maybe time correlated. However, even so, the error would be proportional to the area of the aperture which has shifted, and given that NGTS guiding is accurate to 0.1px, this effect would be negligible. Finally typically on the order of 300-500 bias and dark actions are combined together to produce a calibration master, thus decreasing the error further by a factor of 1.7 to 2.5. The bias and dark frames are therefore unlikely to strongly contribute to the white or correlated noise in the lightcurves.

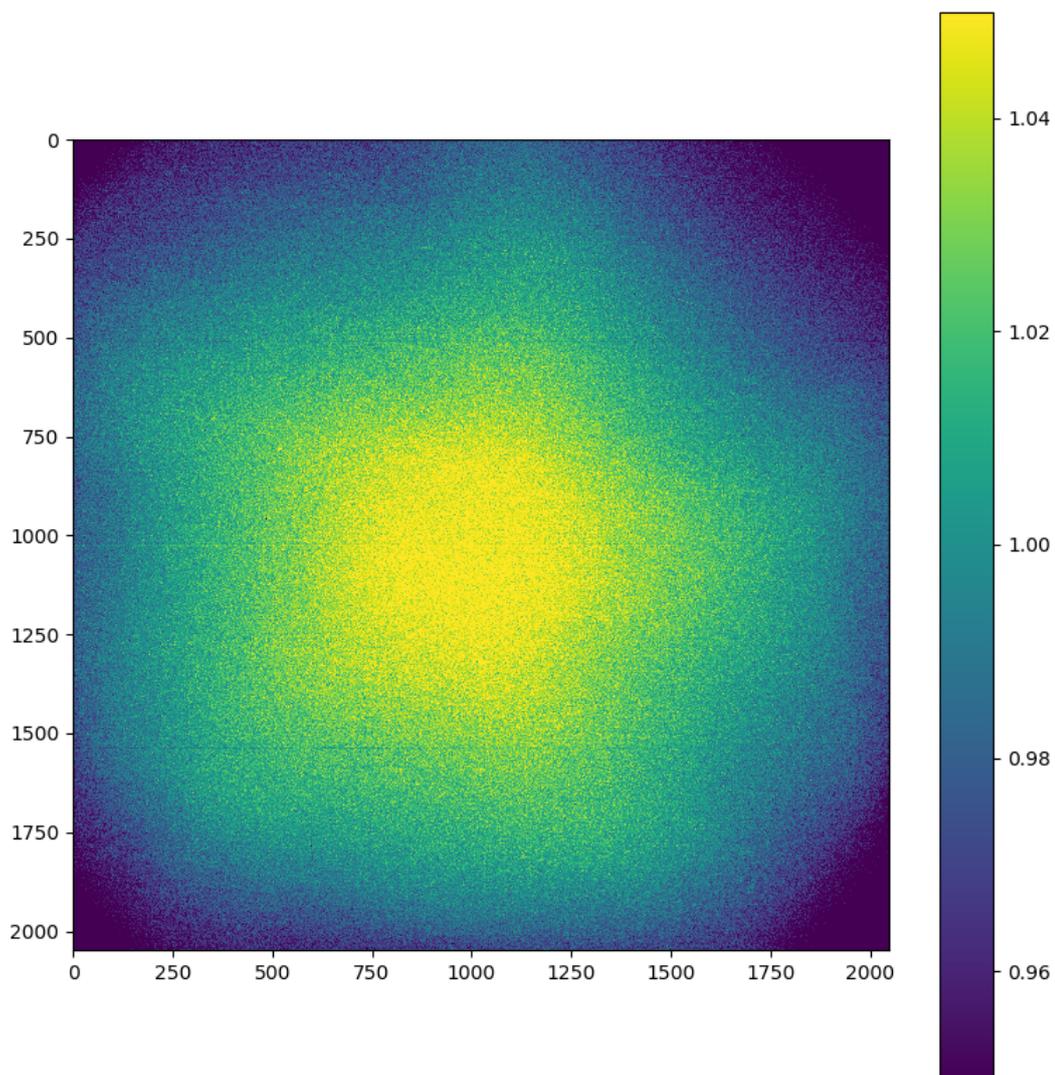


FIGURE 3.10: Example master calibration frame for camera 806. The vignetting across the flat field can be clearly seen here, as well as a slight 'star' shape likely due to the shutter map being imperfectly removed.

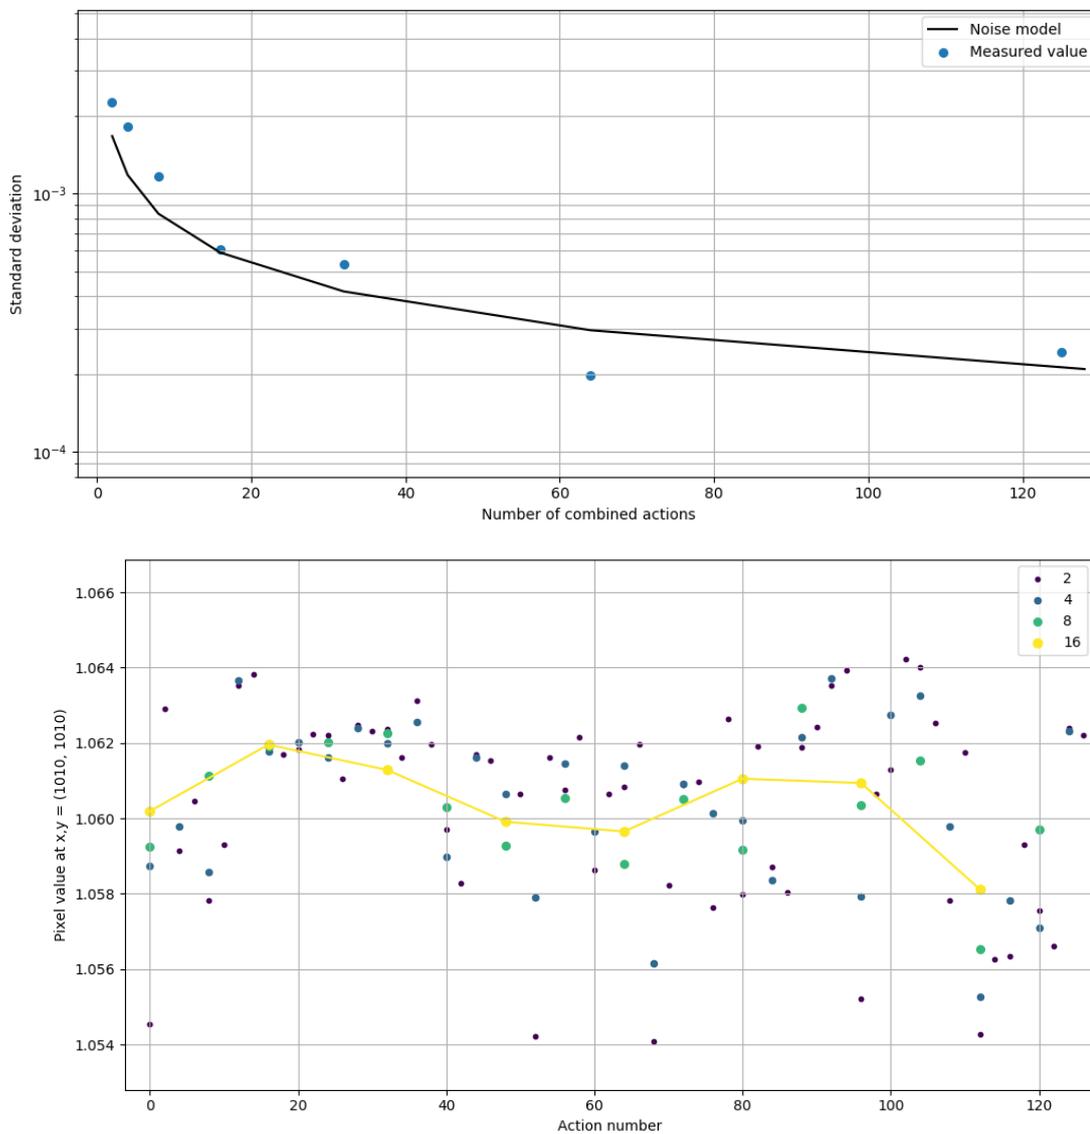


FIGURE 3.11: The top figure shows the average error per pixel as 128 flat actions are combined together to produce a master frame. The black line is a shot noise model based on the number of photons collected across the flat frames. The bottom figure shows the pixel series for (1010, 1010) for each of the flat actions. The different colours show the series converging as more actions are combined together. The error per pixel is the measure standard deviation of this series.

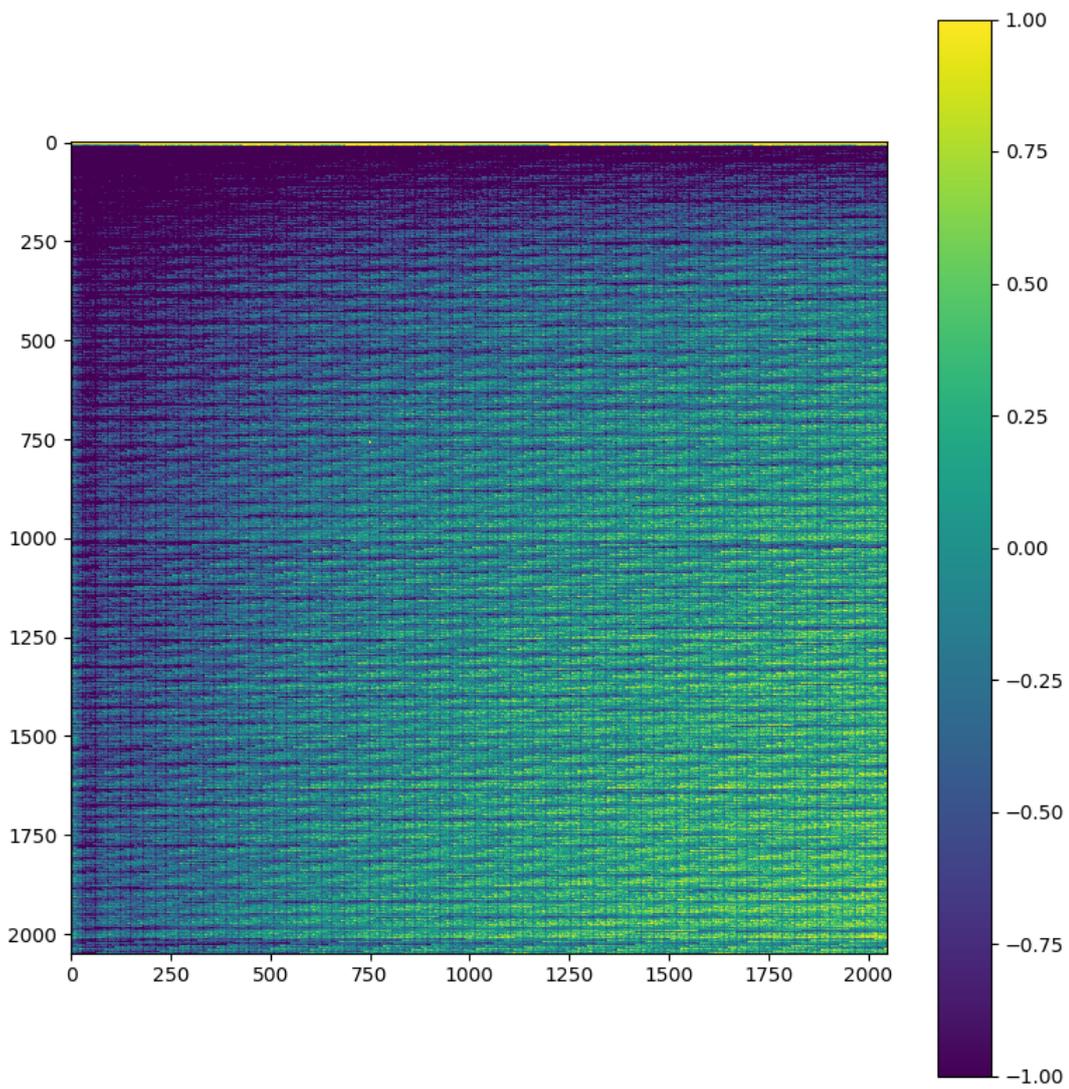


FIGURE 3.12: Example master calibration bias frame for camera 810.

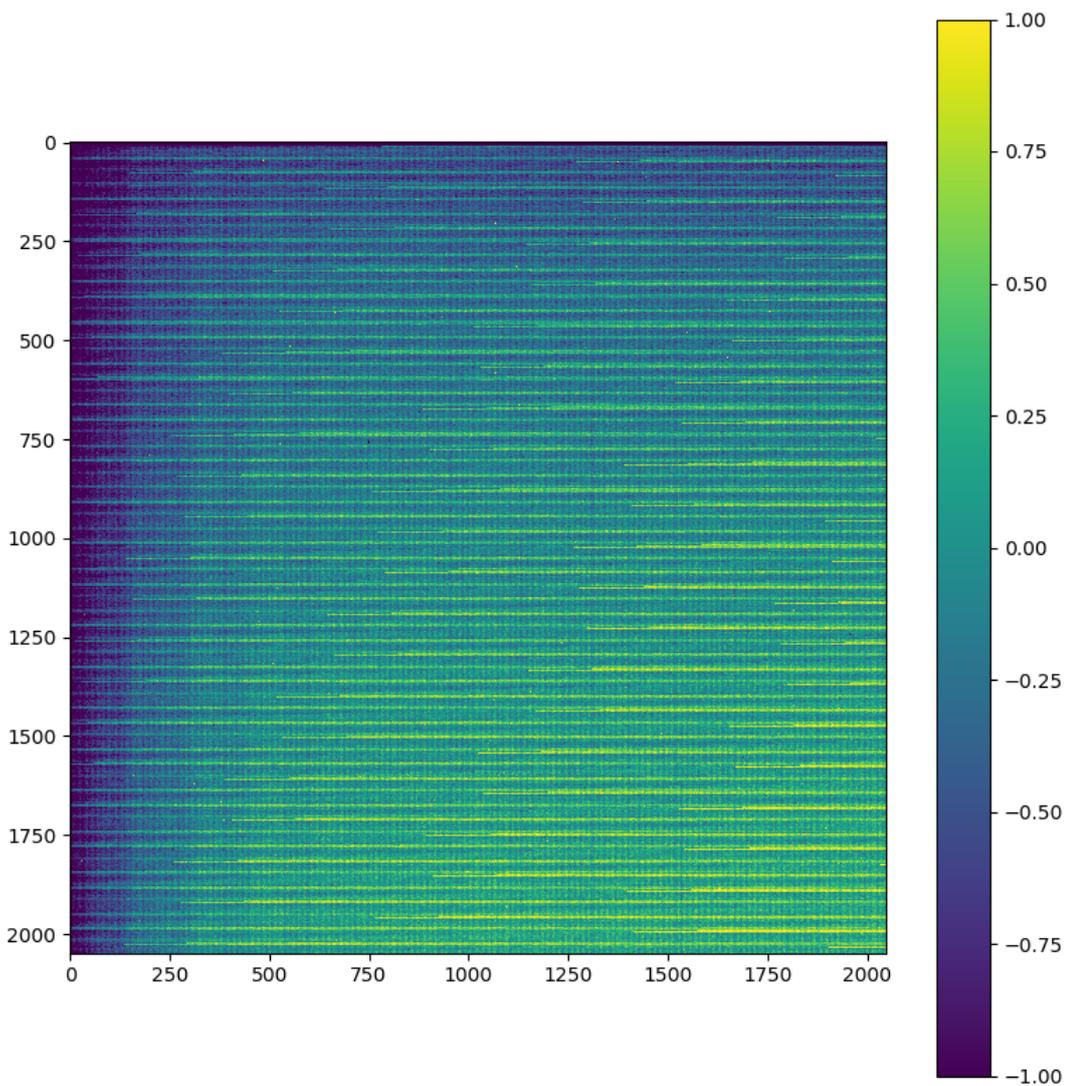


FIGURE 3.13: Example master calibration dark frame for camera 812.

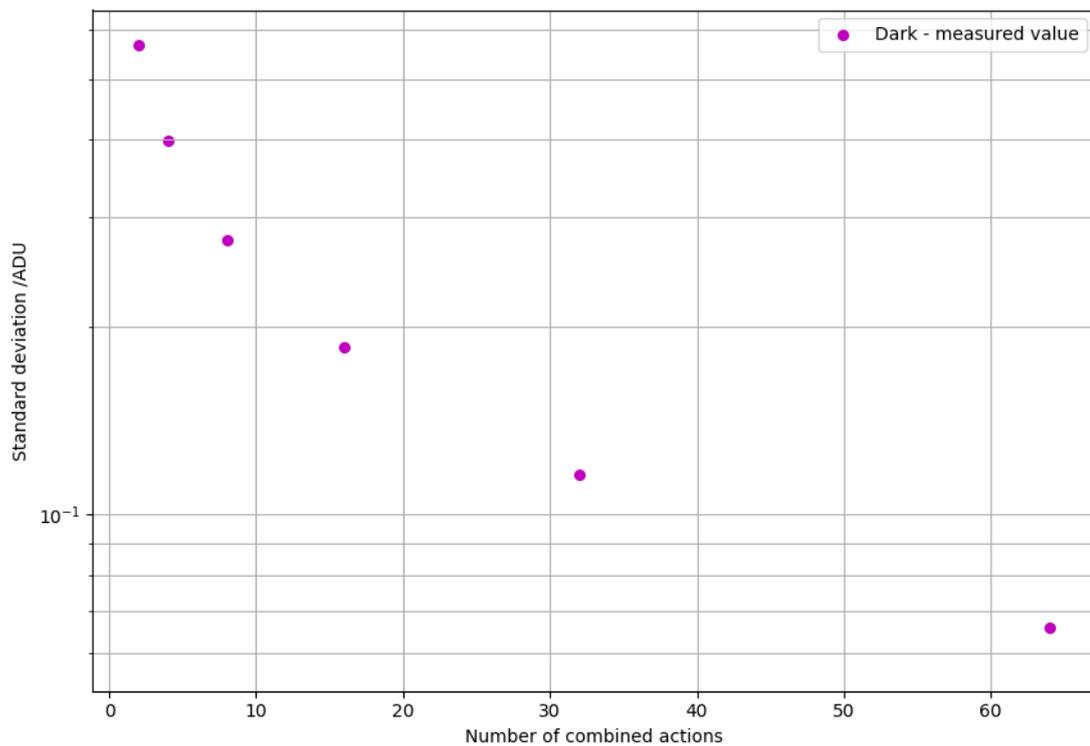


FIGURE 3.14: The average error per pixel of the dark master frame as a function of the number of combined actions. The error is measured as the standard deviation of the values measured for each pixel, in the same way as shown in Figure 3.11.

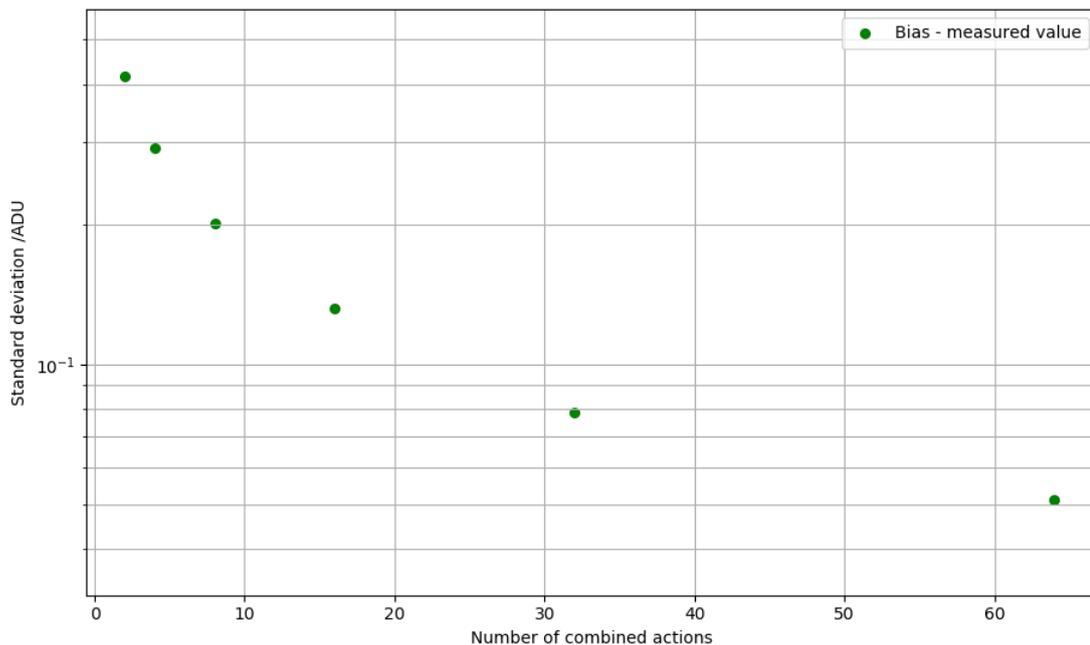


FIGURE 3.15: The average error per pixel of the bias master frame as a function of the number of combined actions. The error is measured as discussed in Figure 3.14 and Figure 3.11.

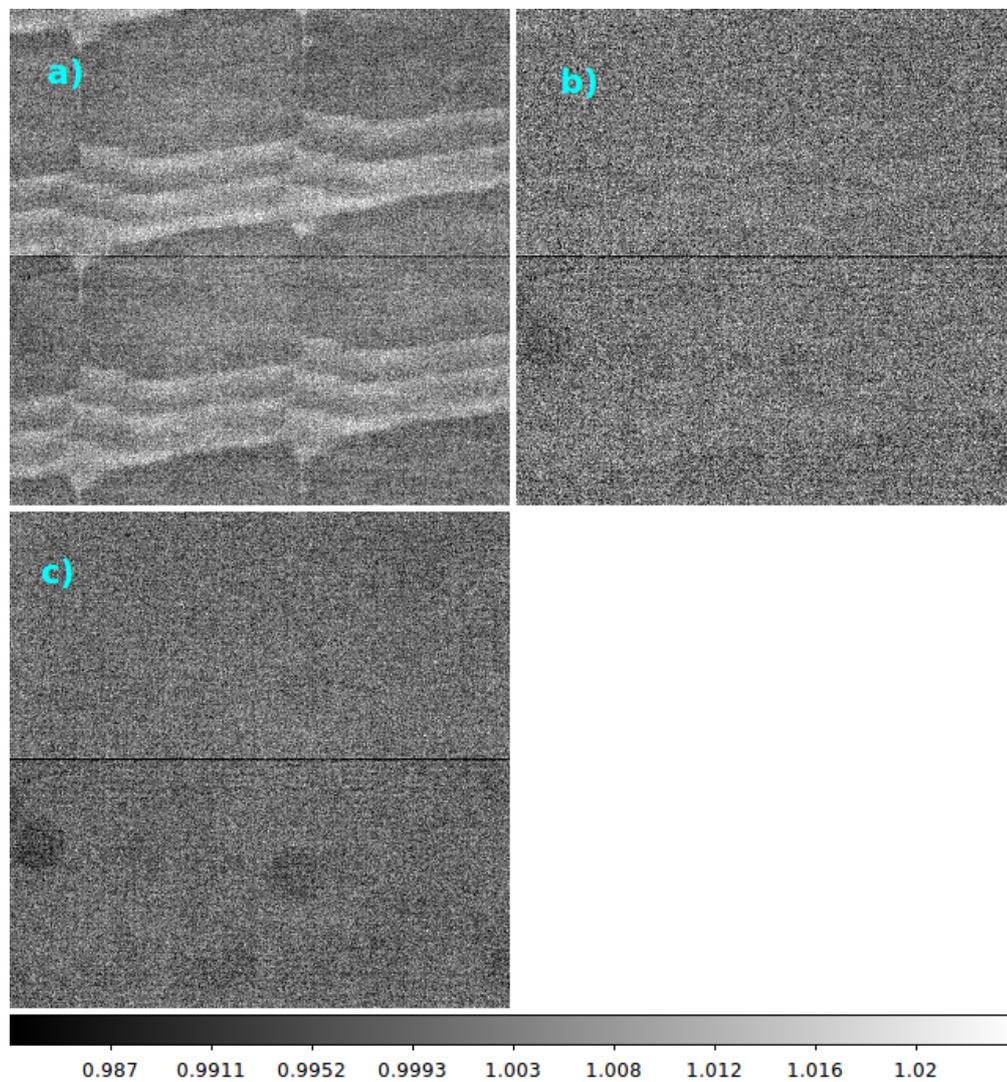


FIGURE 3.16: A 500 by 500 pixel section of three flat fields taken in the Leicester laboratory at different wavelengths: a) 450nm b) 650nm and c) 880nm. The pixel to pixel structure in the flats is consistent between flat fields, however the 450nm flat shows the blue diamond effect, while the 650 and 850nm flats show clear signs of dust.

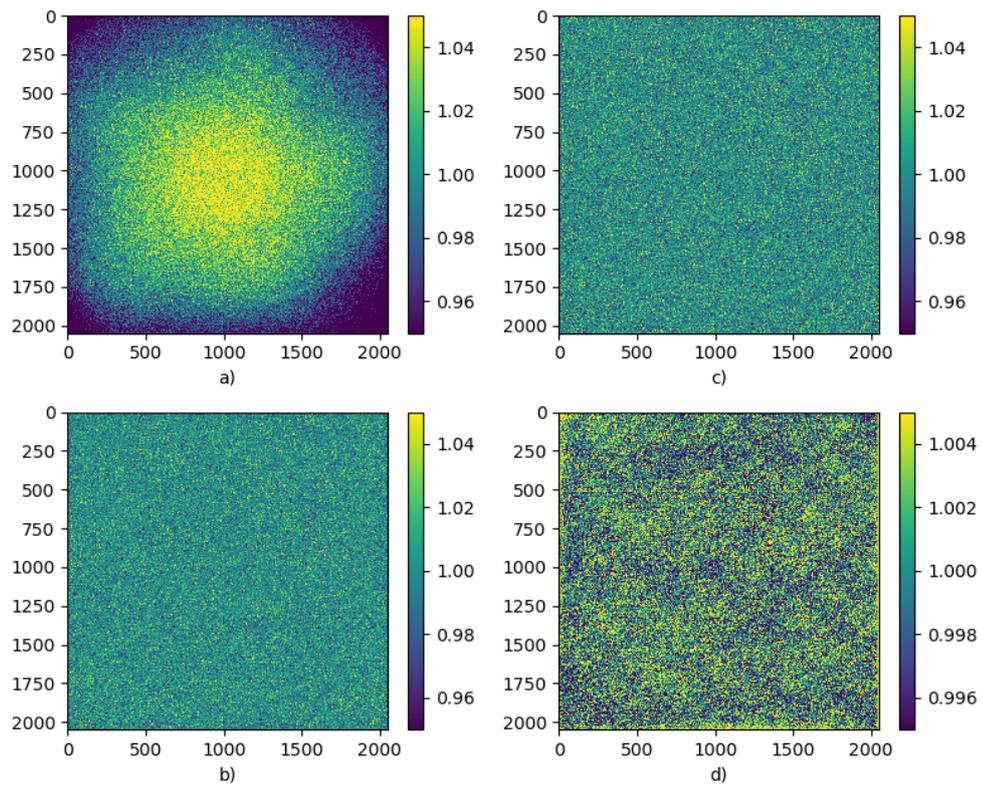


FIGURE 3.17: A comparison between the average laboratory flat and the sky flat for camera 802: a) the sky flat, b) the mean laboratory flat, c) the pixel to pixel variations for the sky flat, d) the residual image between the sky and lab flats. The laboratory flat was built using a raster scan in three different colours and then average together.

3.5.2 Lab Flats

As part of the camera characterisation and testing undertaken by Andrew Grange and Michael Goad at the University of Leicester, a series of flat fields were made for each camera at four different wavelengths: 450, 550, 650 and 880nm respectively. The flat fields were generated using a small uniform light source, combined with a filter, to perform a raster scan across each camera CCD chip. The individual images were then combined to produce a single flat frame. This procedure was performed directly on the camera chip, without any of the telescope optics and so there is no vignetting. Flat fielding in four colours captures some of the wavelength dependant response of the pixel sensitivity and provides a useful handle on how much the flat can vary depending on the colour of the light source. Figure 3.16 shows a comparison between the flat fields taken at different wavelengths. By eye the variations are difficult to see, there are small differences between the pixel values and subtle large scale changes due to the blue diamond effect.

Figure 3.17 shows a comparison between the lab frame, made using a linear combination of the four different colour flats, and the sky frame. To better compare the two, the sky frame is convolved using a Gaussian kernel and then divided by the resulting smoothed frame. This removes the vignetting and leaves only the pixel-to-pixel sensitivity differences. Figure 3.17 d) shows the close agreement between the sky and combined lab flat.

3.5.3 Night Flats

A night flat is a measurement of the flat field made by combining science exposures taken throughout the night in the normal observing regime. This is possible because the sky background on an image is not insubstantial and over the course of several tens of nights it can add up to many millions of counts. One problem which NGTS has compared to other surveys is the precise auto-guiding which keeps stars on their respective pixels with an average precision of 0.1px over the course of an observational season (Wheatley et al., 2017). Unlike other surveys this means that stars are not smeared over the chip in the night flat, but instead are highly defined point sources. This can be overcome by combining multiple night flats from different fields.

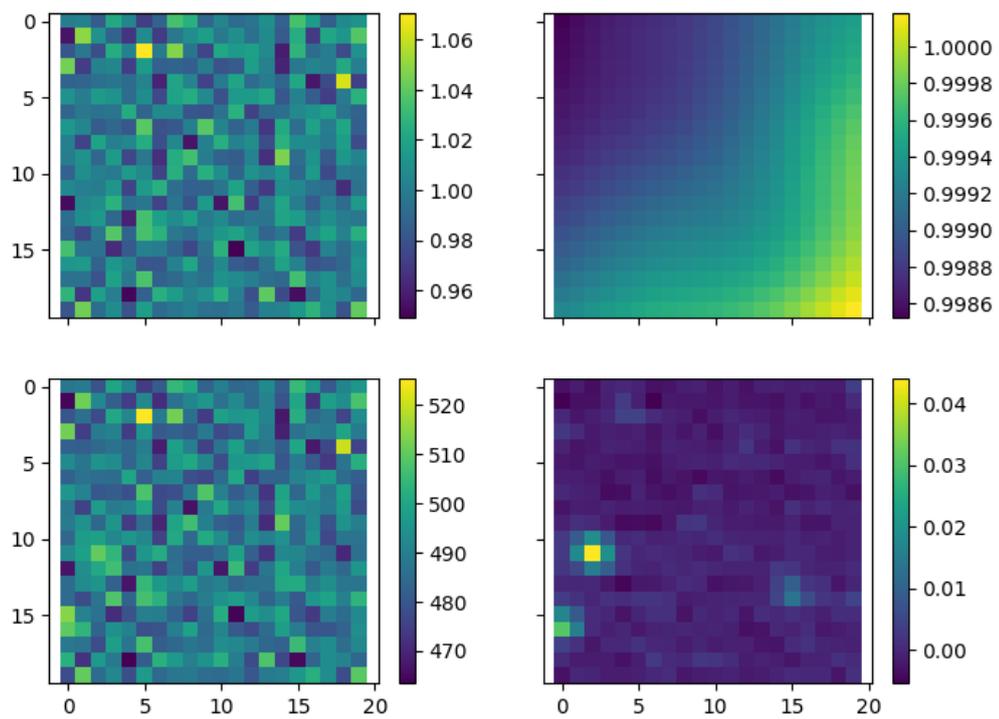


FIGURE 3.18: Comparison of a 20x20 pixel section of the sky flat and night flat from camera 801. Taken from near the center of the CCD. Top left: the pixel to pixel variations from the sky flat, top right: the vignetting from the sky flat, bottom left: the night flat, bottom right: the difference between the night flat and the sky flat.

A night flat was created using 72,000 science images from the field NG1444+0537. These images were taken with camera 801. Each science image was first over-scan corrected before being mean combined with the other images to form a single mean frame for the field. The over-scan was removed in the usual way as detailed in Chapter 3. Figure 3.18 shows the night flat in the bottom left corner and the sky flat, specifically the pixel to pixel variations, in the top left corner. As can be seen both of these show a very similar structure from the individual pixel intensities. The mean level of 500 counts in the stack shows the average background level and gives an indication of the kind of precision the night flat has reached. For 72,000 frames with an average level of 500 counts per frame we average 36,000,000 ADU. Using the measured gain for camera 801 of 2.05 (ref) this gives a precision of 0.0001 % or 0.1 mmag. In practice this value may be an underestimate, due to background gradients which were not removed during this stacking or due to the presence of nearby stars which will systematically raise the level of the pixels. This 20 by 20 pixel box is taken at from the center of the chip at position (900, 980) and so least likely to be affected by these problems.

The bottom right panel of Figure 3.18 shows the difference between the night flat and the sky flat. Three stars are noticeable with a difference of a few percent. Excluding the pixels with stars, the mean difference between the two flats for this section of the camera is approximately 0.0001 i.e they agree to within 0.1 mmag as predicted from the number of counts in the night flat (the sky flat is similarly composed of 100 million counts or so). The vignetting, which was calculated by applying a high pass filter to the sky calibration frame, can be seen in the top right corner. The variation in the vignetting within the 20 by 20 pixel box is on the order of 0.0015% and so is too small to be visibly noticeable in this position of the chip, particularly with the stars present in the night flat.

A full comparison of the night flat to the sky flat can be seen in Figure 3.19. A gradient can be seen across the image which could either be due to a residual sky background in the on-sky flat or it could be as a result of background gradients in the science images. Without background correcting the science images it is difficult to distinguish between the two, nevertheless we see a good agreement between the night flat and sky flat as discussed earlier in this section. Figure 3.20 shows the absolute difference between the night and sky flat, and better highlights the background gradient between the two. The irregular shape of the gradient points to the background in the science images being the main culprit.

Further work would be necessary to create a high quality night flat, either for use in the main pipeline or for comparison with the sky flat. This would require

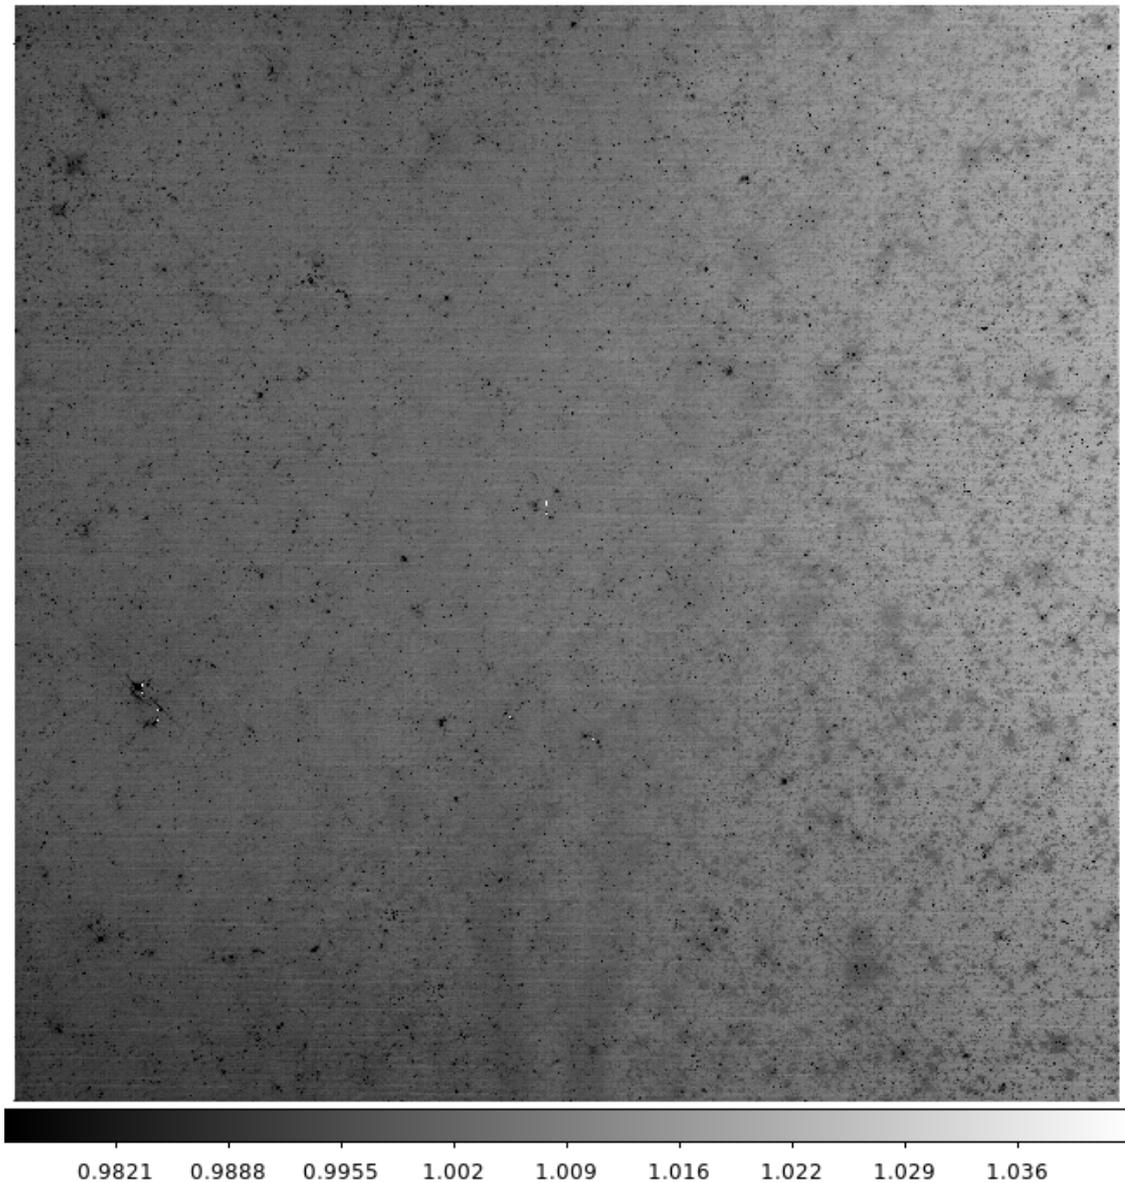


FIGURE 3.19: Ratio of the sky flat to the night flat for camera 801.

careful removal of background gradients in each science image, and perhaps removal of the stars by diving through with the point-spread function. This can be obtained from the night flat itself, and since the NGTS pointing is so stable, it should be possible to fit and remove the stars - with the residuals on the fit being the flat field. Additionally combining many fields together for the same camera should help to remove any systematic errors when building a high quality night flat.

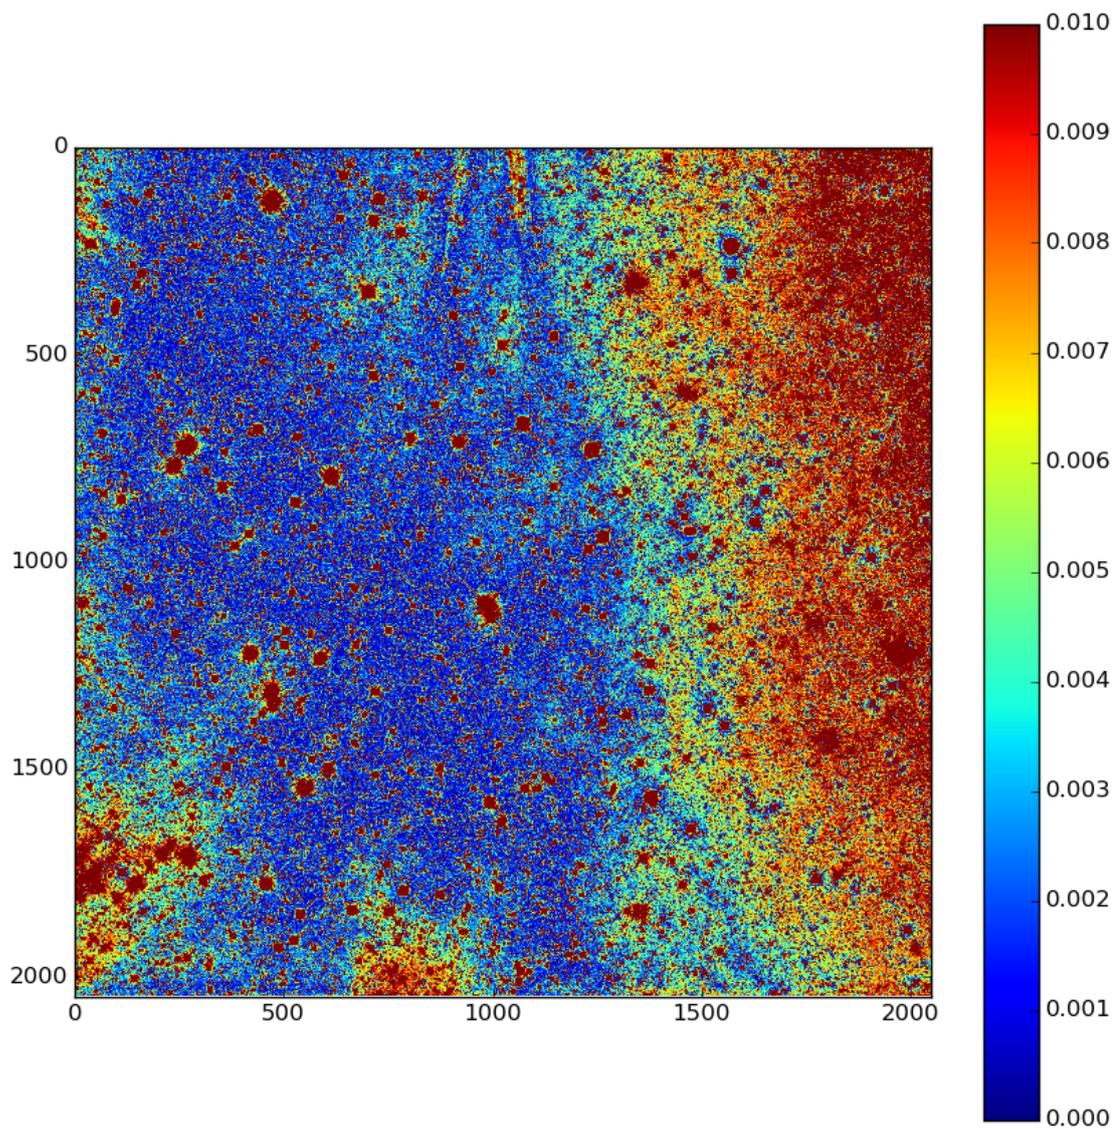


FIGURE 3.20: Absolute difference between night and sky flat for camera 801. This highlights the general background variations between the two.

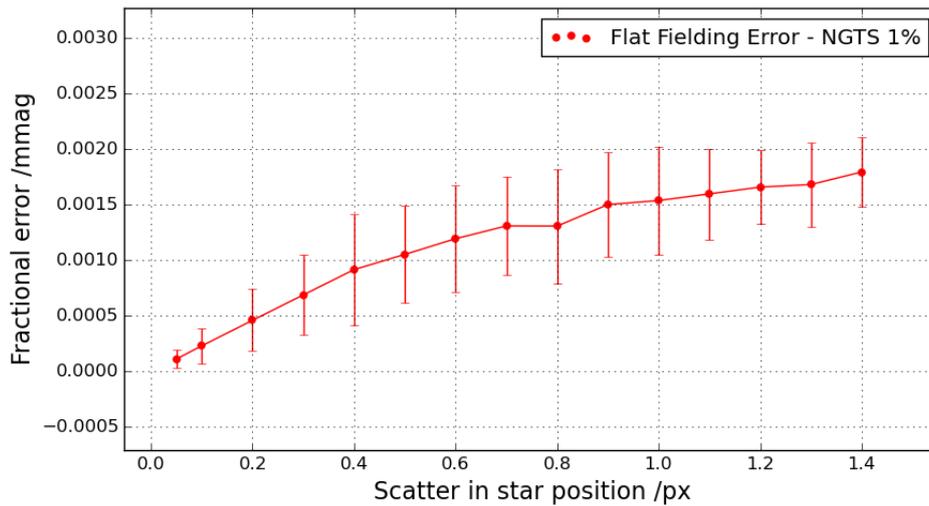


FIGURE 3.21: Results of the simulation of synthetic star and flat field. The plot shows how much noise we expect to be introduced in a lightcurve based on how much a star moves and the NGTS flat field.

3.5.4 Simulating the Flat Field

The previous sections have shown that we can build a high precision master flat field frame from data taken on-sky. Comparisons between the sky flat, lab flat and night flat all consistently show that the same values are being recovered for the flat field. Here we consider how much of an error we expect to be introduced into the data given the level of precision achieved in the flat fielding. The simplest approach to calculating the flat field error is to use a numerical simulation of a star, with and without flat correction and to consider the level of noise in the stellar lightcurve in each case.

A star with a Gaussian point-spread function (PSF) is set up with full width of half max (FWHM) of 1.6 pixels, which is comparable to the average NGTS star on a typical night with usual seeing. A random flat field is created by drawing values from a uniform distribution in the range of 0.99 to 1.01, which is comparable to the pixel to pixel scatter in the NGTS flat field. Each star is then randomly moved from its central position each time for one hundred frames, and in each case a 3 pixel aperture is used to perform photometry on the star. The lightcurve of the star is then compared before and after flat fielding to measure the decrease in the noise from correctly applying the flat, where the noise is taken to be the standard deviation of the lightcurve. In addition we expect stars that move further to have higher flat field noise so the synthetic stars are moved by different amounts ranging from 0.05 pixels all the way through to 1.4 pixels

Figure 3.21 shows the outcome of the simulation. We can see that for most stars we expect an additional noise of between 0.5 and 1.5 mmag given the low level of movement the NGTS stars show due to the excellent auto-guiding. In the case of the no flat fielding we would expect the noise floor to be increased by roughly a mmag or so for the bright stars where flat fielding is a major issue. The noise is expected to be ‘red’ or correlated due to the complex motion of the stars on sky, however given even partially correct flat fielding using the sky-flats we would expect the actual noise introduced by the flat fielding to be much lower than the values shown here. Since we are achieving signal to noise ratios on the order of 1000 the flat field should no longer be making a significant contribution to the NGTS noise budget. The bias and dark frames could also be evaluated in this way however, as shown before dark correction only makes the noise worse and the magnitude of the effect for the bias frames is very small even if they weren’t extremely well measured. As such we expect only the flat field to be a potential source of concern, but as has been shown here, with the kind of signal to noise we achieve the flat field should not be a problem.

3.6 Conclusion

A reduction pipeline is developed which can produce high quality master calibration frames for the NGTS photometry pipeline. The pipeline is set up to minimise the amount of computing time used to process the reduction frames. The final calibration frames are then shown to be of high quality, and in the case of the flat field are compared to flats taken using two different techniques. A simulation is then performed of the flat field to quantify how much the final lightcurves should be affected by the error in the flat, which is shown to be negligible.

Chapter 4

Correlated noise in NGTS

4.1 Introduction

Ground based transit surveys are limited in precision compared to their space based counterparts. This is primarily because of the atmosphere, due to high sky background levels and scintillation which increase the noise in the lightcurve. In principle it should be possible to overcome this limitation (within practical limits) by taking repeated measurements of stellar flux over a long period of time and combining them together or 'bin' the lightcurve. This is the approach taken by most ground based transit surveys, however in this situation the key limiting factor is so called 'red noise' or time correlated noise. Unlike white noise which bins down in time, red noise does not bin down and so it provides a strong lower limit on the depth of the transits which can be detected. A detailed discussion of the effect of red noise on transit detection can be found in Pont, Zucker, and Queloz (2006). For the NGTS experiment to maximise it's planet catch we require the correlated noise to be below 1mmag for stars brighter than 13th magnitude. As such it is critical that the correlated noise properties of the NGTS experiment are understood and that efforts are made to reduce the correlated noise.

Existing Ground Based Transit Surveys

Previous ground based transit surveys have come up against a similar problem. Smith et al. (2006) discuss the impact of correlated noise on the SuperWASP survey. They find that SuperWASP lightcurves contain a significant 'red' component, and find indications of systematic trends on the 2.5 hour time-scale at a level of 3mmag. Smith et al show that the popular detrending algorithm SYSREM (Tamuz, Mazeh, and Zucker, 2005) can help to reduce the correlated

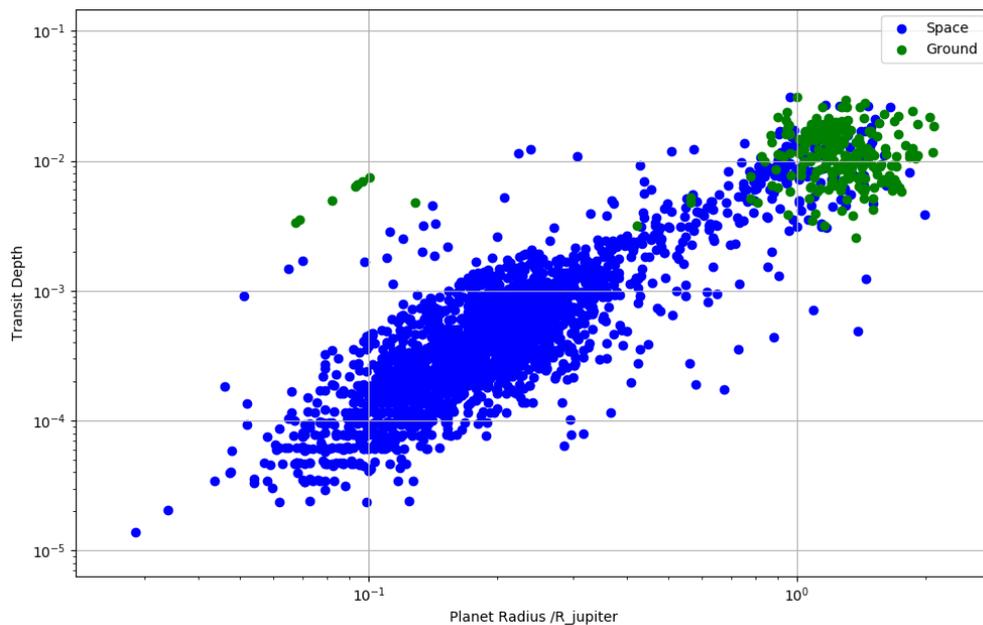


FIGURE 4.1: Transit depth vs planetary radius for transiting exoplanets detected from the ground (blue) and space (green). Most sub-Jupiter sized planets have been discovered from space. Data taken from ‘www.exoplanet.eu’. Error bars were omitted from the plot for clarity.

noise. A few such methods exist, for example the SARS algorithm (Ofir et al., 2010) and the Trend Fitting Algorithm (Kovács, Bakos, and Noyes, 2005) and have been used extensively. More advanced techniques, based on the latest advances in machine learning, are also being developed. For example Zucker and Giryes (2017) apply a convolutional neural network to search for shallow transits in noisy lightcurves. Such networks are part of the ‘deep learning’ repertoire for object detection and classification. These methods are still being developed and in the future will hopefully unlock small transiting planets which otherwise would have been impossible to find.

Figure 4.1 shows a plot of transit depth vs planet radius for all transiting planets discovered so far. Most planets discovered from the ground tend to be approximately a Jupiter mass or so. A few exceptions can be seen in the top left quadrant of the plot, most notably the TRAPPIST planets (Gillon et al., 2017) found around a late M-dwarf host star. The shallowest transit discover so far belongs to the KELT experiment, with KELT-11 (Pepper et al., 2017), an inflated sub-Saturn exoplanet with transit depth of just 2.3mmag. KELT-11 was found around an exceptionally bright subgiant with a $V_{mag}=8$, which no doubt helped to identify the system. Figure 4.2 shows transit depth against the brightness of

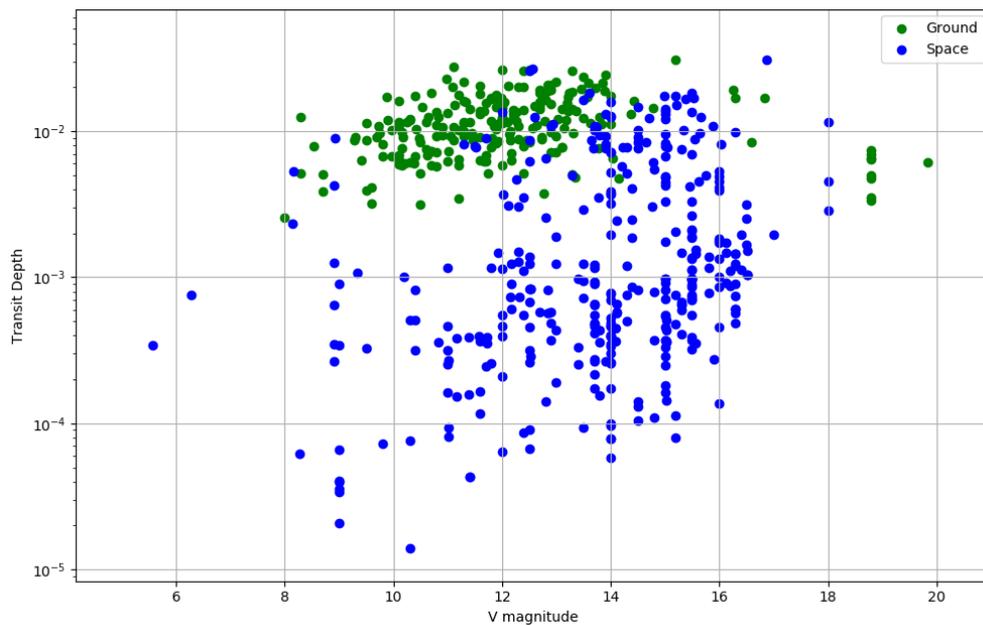


FIGURE 4.2: Transit depth vs V_{mag} for transiting exoplanets detected from the ground (green) and space (blue). Data taken from www.exoplanet.eu and only shows planets with available V magnitudes. Most exoplanets discovered from space-based missions did not have a measure value available. Error bars are omitted for clarity.

the exoplanet host stars. KELT-11 remains one of the brightest exoplanets discovered so far, at least from the ground. Discovering a 1-mmag transit from the ground is a formidable challenge, even around a bright star, and represents an exciting frontier in exoplanet science.

Correlated Noise

There are many potential reasons why correlated noise can be found in a lightcurve; for example the movement of the telescope on sky is very complicated and could be a source, or the camera chip could be non-linear in a regime in which the stars are observed. Sky background levels also show significant correlations across different time scales. Other CCD effects, such as bad flat fielding can introduce correlated noise as well. The NGTS pipeline is designed to correct for many of these noise sources and attempt to remove any correlated noise from the lightcurves. However, in the case that this is done imperfectly, we would like to pin down which is the main source of correlated noise in order to minimise it.

The NGTS cameras have undergone extensive lab testing and characterisation. While this does not exclude all instrumental effects, it minimises the chance that

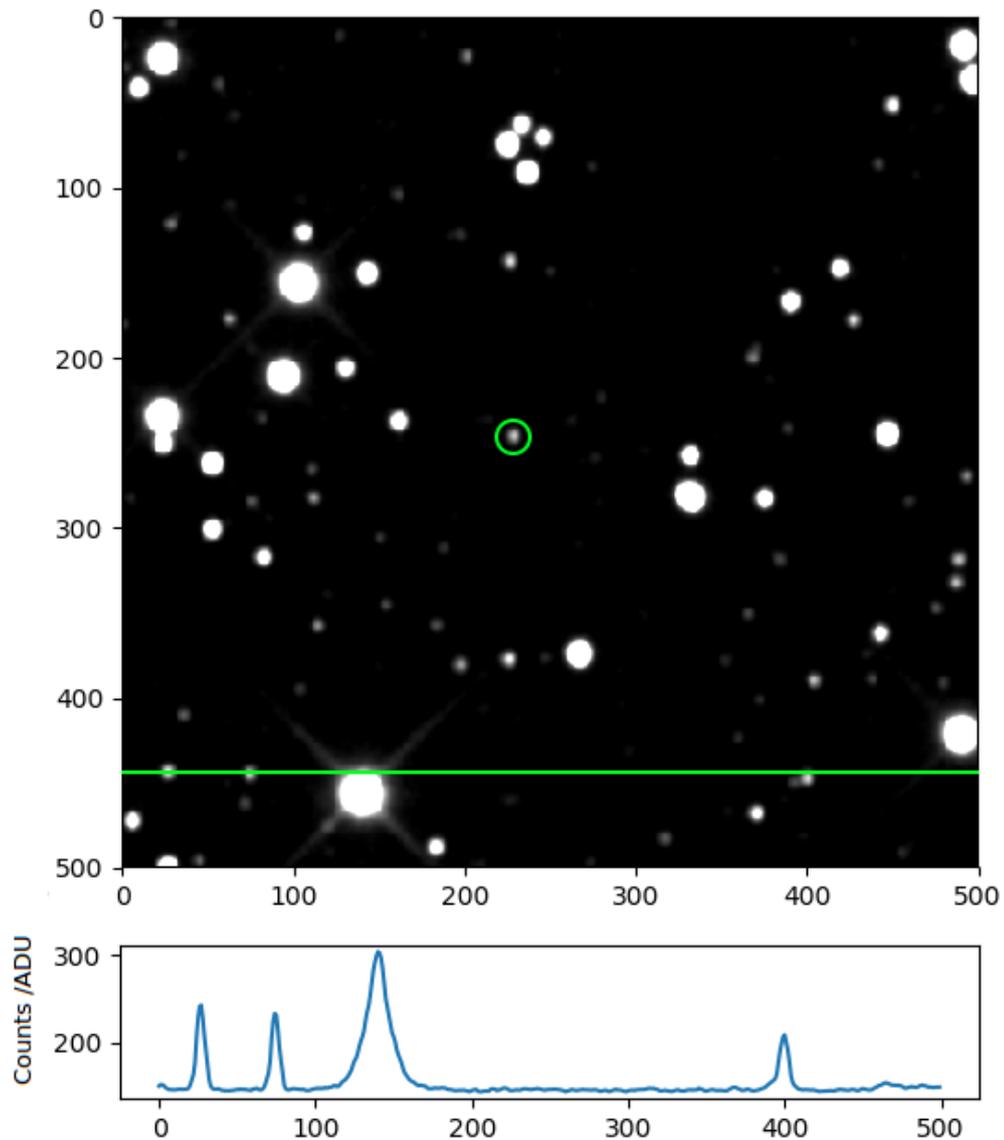


FIGURE 4.3: An example image taken with NGTS camera 803. The image here is a composite image created from 100 individual frames which have been systematically dithered and stacked together. Since NGTS images are normally under-sampled this helps to show the PSF, and is used for source detection. The lower panel shows a cut through the image, indicated by the green line. The green aperture indicates the size of the NGTS aperture, this may appear small compared to the stars, but the image scale has been chosen to emphasize the PSF. A 3px aperture radius contains 95% of the flux from unsaturated sources.

the CCD is introducing correlated noise, and instead points to other factors such as the sky background, flat-fielding or imperfect astrometry as being a likely noise source. The pipeline may be imperfectly correcting for these noise sources, at a level which is significant for transit detection. If it is not then it would be useful to excluded the relevant pipeline components as potential issues.

When photometry is performed on a star, the measurements are a function of two parameters, the CCD x, y co-ordinates of the aperture (which should be centred over the target star) and the sky background estimate. Both the background level and the x, y movement of the star shows cyclical correlated motion, and so if there is a systematic issue with how they are obtained we may end up injecting correlated noise into the lightcurve. Figure 4.3 shows an aperture on a NGTS target. By eye it is not difficult to see where the star is centred and roughly what the background flux may be. However if these values need to be measured to 1% than this becomes much harder. It is not clear how accurately these values need to be known. In addition to the CCD position and sky background, the strong airmass trend, which for most stars is dominant, has to be removed as well. The airmass is heavily correlated across a 1-day time scale and does not bin down well. Even small errors in removing the airmass can leave residual correlated noise in the lightcurves. It is also worth considering the flat field, which is thought to be a source of correlated noise in the SuperWASP experiment (Pollacco et al., 2006).

Detrending Methods

Commonly used methods for detrending lightcurves and reducing correlated noise are all related in that they search and remove correlations between lightcurves. The reasoning stands that most of the correlated noise is due to external factors which should be the same for all stars in the image, and so other stars can be used as a basis for detrending. These methods do not consider the cause of the offending correlated noise before attempting to remove it. The correlations are decreased 'blindly' in a sense. As such I believe it is useful to label SYSREM, SARS, TFA and other such methods as a 'top-down' approach to removing correlated noise. While these techniques are very successful at removing correlated noise, a careful consideration of the origin of correlated noise in the NGTS data may be needed to breach the 1 mmag level.

In this chapter I propose a 'bottom-up' approach to augment the use of SYSREM in reducing the correlated noise in the NGTS data. A series of simple tests will

be use to attempt to identify the origin of correlated noise in each component of the photometric process (reduction, astrometry, background removal and etc...) and quantify it's effect. The photometry can then be optimised to minimise the correlated noise, reducing the work which SYSREM has to do. This may tackle systematic effects which are not necessarily correlated across all the stars, for example such as imprecise flat fielding, and which may prevent us detecting transits below 1mmag.

4.2 Measuring Correlated Noise

The practical effect of having correlated systematics in NGTS lightcurves is that they set an effective lower limit on the precision which can be achieved. As stated before, correlated red noise impedes transit detections using the Box Least-Squares method (Pont, Zucker, and Queloz, 2006). Correlated noise is in part dependent on the aperture size use, the background flux subtracted and the precision with which the apertures are placed.

As a working definition we take correlated noise to mean any systematic noise source where the standard deviation does not decrease as quickly as white or Gaussian noise when measurements are combined. In other words, we expect the correlated noise not to 'bin down'. This is not true of all forms of correlated noise, but it captures the type of noise which impedes transit detection (Pont, Zucker, and Queloz, 2006). The 'slow binning' property can be turned on it's head, in order to allow us to measure the noise level and try and identify from which part of the pipeline it comes from. However, first it is useful to quantify how much correlated noise is in the NGTS data and to see whether there is any dependence on the camera, stellar colour or magnitude of the star.

4.2.1 Correlated noise and binning

Figure 4.4 shows the standard deviation of a single NGTS lightcurve as a function of the number of points per bin (n). The standard deviation is evaluated at each step after the lightcurve is binned. As the errors are largely due to a Poisson process we expect the lightcurve noise to bin down as a function of $1/\sqrt{n}$. This is clearly not the case however and we can see that in the limit of large n the lightcurve will tend to some constant value. We take this value to be the standard deviation of the correlated noise present in the lightcurve. To calculate

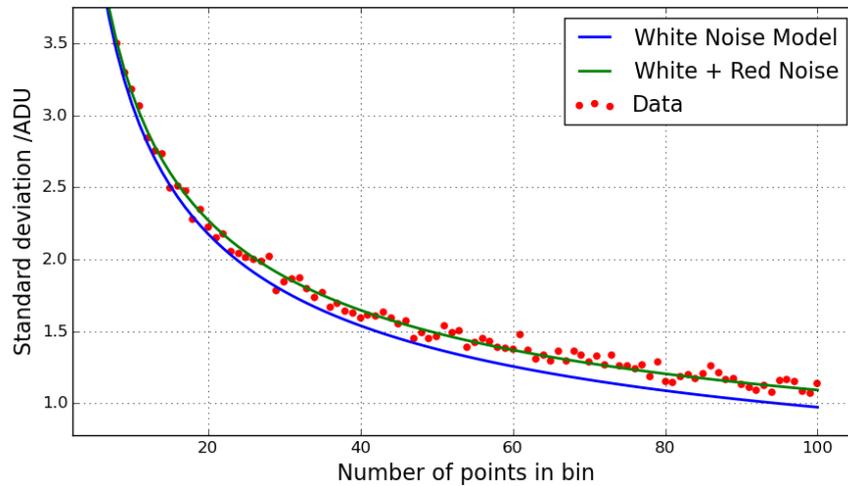


FIGURE 4.4: Standard deviation as a function of bin size for an NGTS lightcurve. The blue line is a white noise Poisson model, while the green line is white noise plus a fixed correlated noise constant.

the value of this constant a model is fit: $std(n) = \sigma_{white} / \sqrt{(n)} + \sigma_{red}$. This gives a value of the white noise and correlated noise for each lightcurve. This process is repeated for each lightcurve taken with NGTS in order to get a measure of the red and white noise for each star. The same method is used in Pont, Zucker, and Queloz (2006). Before binning the lightcurve a 5.0 sigma-clip is used to remove any outlying points.

In fitting the model an upper-limit must be chosen for the bin size which is large enough to get a decent estimate of the correlated noise, while small enough to be computationally feasible. Additionally the standard deviation does not need to be measured at every discrete bin size. Perhaps bin sizes of 5, 10, 15 points are sufficient? In order to understand how this estimate of the correlated noise varies based on the upper-limit chosen for the bin size a series of tests are run as show in Table 4.1. The values in the table are for every target in the field NG0304-2518, including targets down to approximately 16-17th magnitude, which are very noisy, explaining the relatively large values.

Based on the results shown in Table 4.1, using a maximum bin size of 1000 provides the best results, based on the percentage error. The percentage error reflects the fit to the noise model and the level of variation in the standard deviation while binning the data. The percentage error does not represent a "bootstrap" error, or the range of correlated noise values in the field. These are dictated by the distribution of stellar magnitudes in the observed field. Using a

Data points per bin	White Noise / mmag	Percentage error	Correlated Noise / mmag	Percentage error
10	61.83	0.27	1.80	4.7
100	58.20	0.41	3.06	1.68
1000	56.70	0.20	4.02	0.36
100 (multiple 10)	59.13	0.61	3.87	1.05
1000 (multiple 5)	58.30	0.52	3.92	0.74
1000 (10% of data)	60.30	0.37	4.18	0.37

TABLE 4.1: Measured correlated noise as a function of maximum bin size and sampling, where the sampling is given in multiples of the initial bin size where specified. The maximum bin size is the upper limit on the number of points grouped together before calculating the standard deviation. The values of white and correlated noise given are roughly consistent with each other, though some variation can be seen. The percentage errors were calculated from the covariance matrix of the model fit and represent the quality of the fit. These values are important to determine which scheme is the best to adopt.

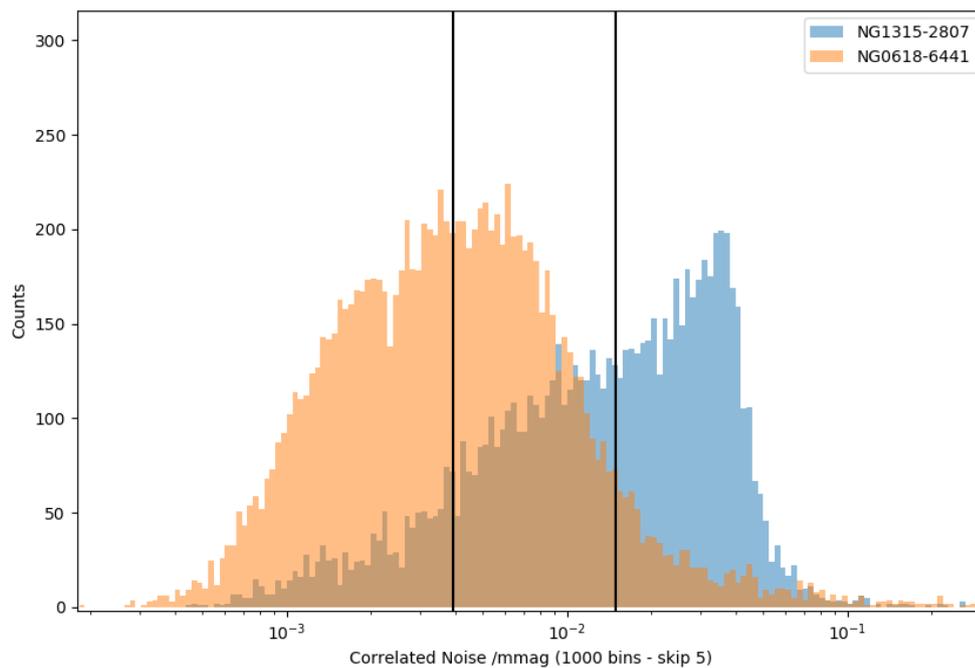


FIGURE 4.5: The figure shows the estimated correlated noise across all the lightcurves in two NGTS fields. The distribution of correlated noise values is quite broad, as is expected, because the fields contain stars ranging from 8-16th magnitude in I_{NGTS} . The black lines represent the median values of those fields. While the median does not accurately capture the width and shape of the distribution it is a good indicator if one field contains more red noise on average than another corresponding field.

larger bin size, would result in having increasingly fewer data points on which to estimate the standard deviation with and would take much longer. The percentage error on the fit using a maximum bin size of 1000 is sufficiently small. Using bin sizes which are multiples of five, significantly speeds up the correlated noise measurement while only sacrificing a little bit of accuracy. This is the scheme which will be used to estimate correlated noise for the rest of this chapter. For measurements involving a whole field or group of stars the median of the distribution will be presented. As can be seen in Figure 4.5 the distributions of correlated noise can look quite different in two fields. Measuring the breadth of the distribution will capture information about the magnitude range of the stars in the field. The median provides a good way of comparing the rough noise level between fields, indicating whether one set of stars is much noisier with respect to the other.

The correlated and white noise components were calculated for each target for each field in the 'TEST18' NGTS dataset. This consists of 386,959 stars across 43

Magnitude (I_{NGTS})	Median Correlated Noise / mmag	Range Correlated Noise / mmag	White Noise / mmag
9-11	0.89	253.6	6.3
11-13	1.11	226.0	20.2
13+	5.93	175.3	89.3

TABLE 4.2: Measured correlated noise as a function of stellar magnitude, calculate in the NGTS bandpass.

Camera ID	Median Correlated Noise /mmag	Median White Noise /mmag
801	2.37	50.3
802	3.48	57.1
803	2.54	56.6
804	4.61	70.8
805	4.28	55.9
806	4.14	64.9
808	4.28	57.4
809	3.42	58.2
810	3.85	61.4
811	8.72	66.5
812	3.91	55.6
813	5.51	72.9

TABLE 4.3: Measured correlated and white noise for each NGTS camera. These results represent a rough measurement of the correlated noise per camera and are not a quality rank to fractions of a percent. The correlated noise values can be influenced by clouds which may explain the unusually high levels seen in camera 811.

separate fields. Each field has a varying number of images from approximately 20k through to a complete field at roughly 200k measurements. Table 4.2 shows the correlated and white noise measurements as a function of stellar magnitude. An upper limit cut-off of 9th magnitude is used as this is close to the point where stars enter a regime close to saturation. The median correlated noise in this magnitude range is 0.89 mmag, indicating that the noise is sufficiently low on most stars to search for transits in the mmag range. The noise increases slightly to 1.11 in the 11-13th magnitude range, giving an overall correlated noise of approximately 1mmag down to 13th magnitude.

4.3 Specific Contributions to the Correlated Noise Budget

The second goal of this work is to attempt to identify which component of the pipeline or which combination of components is contributing the most to the correlated noise. The different parts of the pipeline which could contribute to the correlated noise are the aperture placement, background correction, flat field, and airmass correction. The last component is related to the detrending algorithm used to remove common trends in the lightcurves. For NGTS data this is a custom implementation of the SYSREM algorithm (Tamuz, Mazeh, and Zucker,

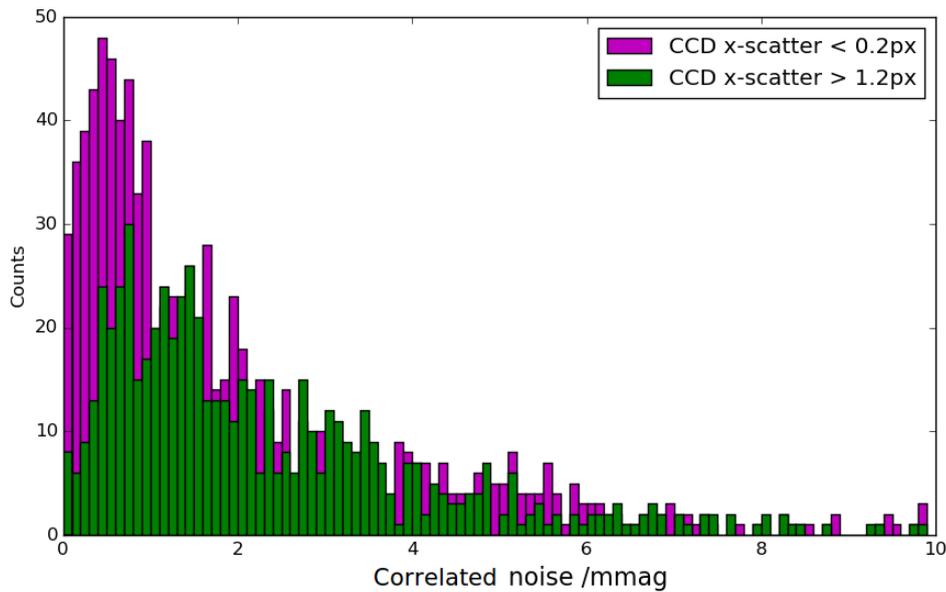


FIGURE 4.6: The plot shows the measured correlated noise for two populations of stars, one where the scatter in the CCD x-coordinate is less than 0.2 and the other where it is greater than 1.2px. These represent opposite populations of stars, those that move very little and those which move very much across the CCD.

2005). For each component the data will be examined to attempt to determine if it is contributing to the correlated noise, by checking how the correlated noise varies as a function of pipeline parameters. It may be the case that the pipeline is working perfectly, this is in itself a useful result and would point to an instrumentation issue giving rise to the remaining correlated noise in the data.

4.3.1 Flat-Fielding

Before the photometry is produced, each science image is divided by a flat-field frame in order to remove the large scale vignetting and small scale pixel to pixel sensitivity changes imprinted into the frames. As the stars drift over the pixels, these sensitivity changes could introduce spurious signals into the data, and so they need to be removed carefully. Chapter 3 details how this is done in the NGTS pipeline and also contains a simulation of the expected flat field noise for the NGTS pipeline. It is expected that flat field error should not be a problem for this dataset, nevertheless it is important to confirm this result.

Stars which move large distances across the CCD chip are more likely to show flat field error, as they drift over more pixels with different sensitivities. Any

error in the flat across the chip, will be imprinted in these lightcurves more than the rest. Figure 4.6 shows the fractional correlated noise for stars which move more than 1.2px on the chip and those which move less than 0.2px for the field NG3056-2518. These movements are as a result of airmass changes in the field, as well as distortion in the NGTS field of view (Wheatley et al., 2017). The NGTS pipeline aperture photometry tracks and compensates for both of these effects within the limits of the precision of the hardware and pipeline. Figure 4.7 shows the respective positions of these two sets of stars on the CCD chip.

As can be seen in Figure 4.6 the correlated noise is on average higher for stars which move further on the chip. The difference in the mean of the two is approximately 1.0 mmag though it is heavily influenced by stars in the tail of the distribution. This result would indicate that stars which move more do indeed have a higher correlated noise on average and points to the flat field as one possible reason for this. A second possibility is that the aperture placement is worse on stars which move more, and if this can be excluded than the creation of the flat field needs to be revisited.

4.3.2 Aperture Placement

When performing aperture photometry on a star, the precise position of the star on the CCD chip needs to be known so that an aperture can be centred on it. Aperture placement is crucial to high precision photometry as even small changes in the position can lead to large placement errors, particularly as NGTS is under-sampled with a FWHM of 1.6px. An analytical equation of the aperture placement error was derived by Irwin et al. (2007) for the data processing and lightcurve production of the Monitor project. This equation can be used to calculate the error for a specific FWHM, aperture radius and offset error.

Numerical Simulation

As a first step a numerical simulation was constructed to simulate the aperture placement error for the NGTS data. For each simulation 1,000 images of size 10 by 10 pixels were generated with a fake Gaussian PSF function with FWHM of 1.6px in the center of each. The total flux of the Gaussian was set to be 1 count i.e. the Gaussian was normalised. The PSF was offset randomly by values drawn from a normal distribution with a mean of 0.0 px and a standard deviation of 0.5px. This introduces a small error into the images as varying amounts of the

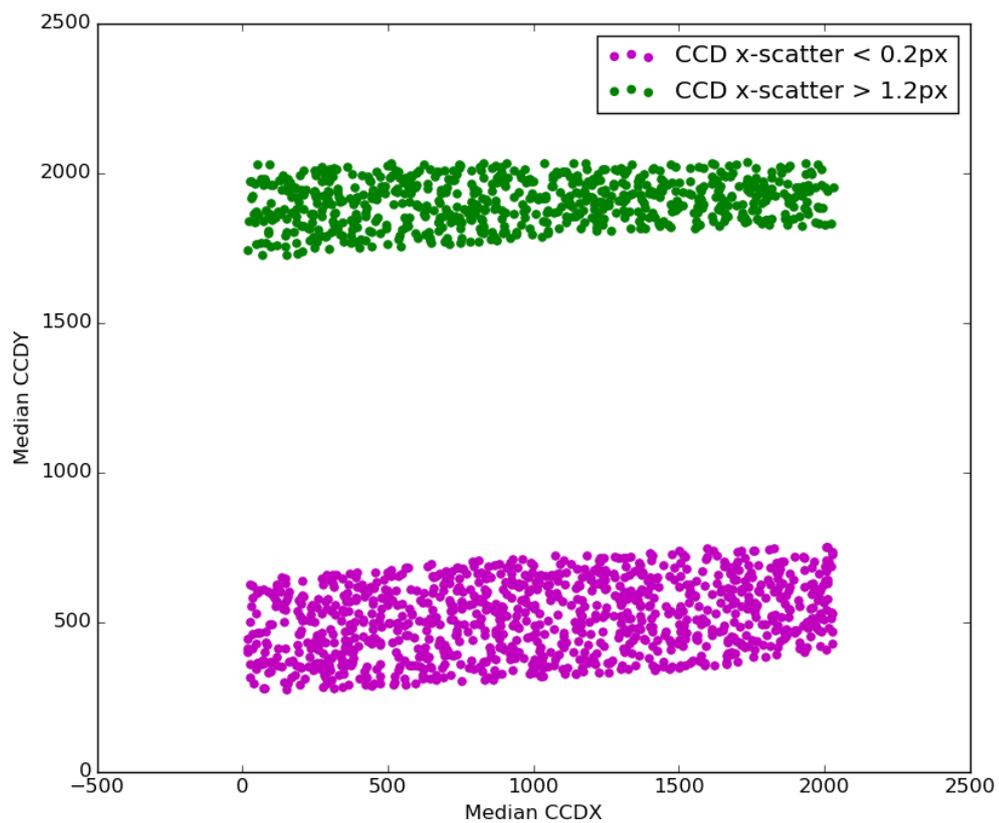


FIGURE 4.7: A plot of the physical mean position of a low-movement (less than 0.2px standard deviation) and a high-movement (greater than 1.2px) group of stars on the CCD. This is only for one field, the stars in each field will have different trajectories depending on its position on sky.

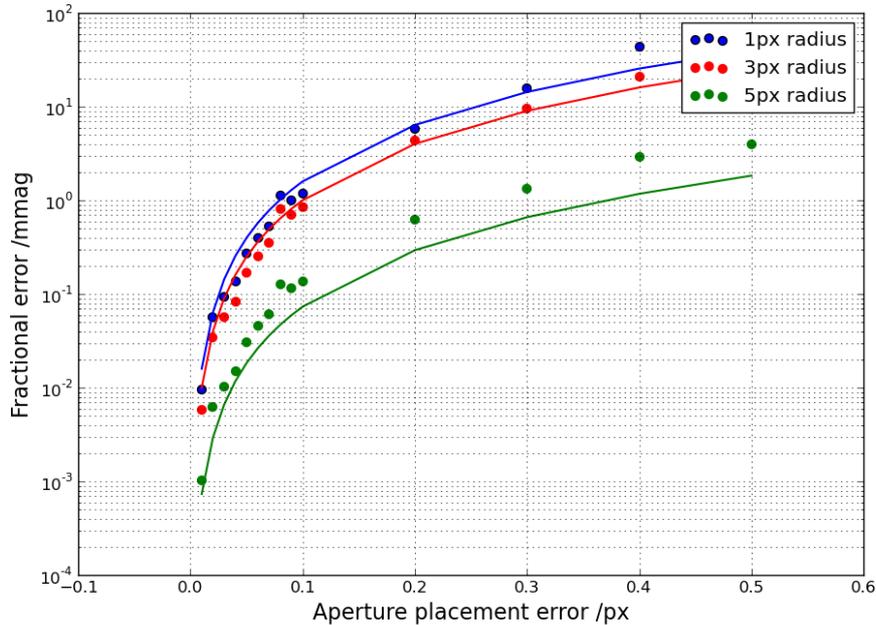


FIGURE 4.8: Aperture placement error as a function of the size of the radius. Both a numerical simulation and an analytical model are shown here.

Gaussian are in the image. No background flux, flat field or other errors were added into the simulation. No shot noise or other changes in flux were added.

Photometry was then performed on these images using an aperture of fixed size, at either 1px, 3px or 5px. The aperture was placed with a random offset drawing from a normal distribution with mean of 0.0 and sigma of 0.01, 0.02, ..., 0.1, 0.2, 0.3, 0.4 and 0.5 pixels respectively. For each aperture error and each radius size a new set of 1,000 images were randomly generated. The error in the lightcurve was then measured and recorded. Figure 4.8 shows the results of the simulation overlaid with the analytical expression derived by Irwin et al. (2007). The results are consistent and show that for the NGTS PSF an aperture error of just 0.1px would potentially result in 1 mmag of noise. This noise could either be correlated or white depending on the error, but given the complicated motion of the stars in the sky with respect to one another, the auto-guiding used in NGTS and the camera distortion it may be expected that some of the noise from aperture placement is in fact correlated.

Correlated Noise vs. Aperture Size

One feature of Figure 4.8 is the steep drop off of the aperture placement error as a function of the aperture radius chosen. This makes sense as the larger the error the more of the stellar PSF is incorporated and then in turn the smaller the effect of any offset. If there is correlated noise due to the aperture placement, then it is expected that this correlated noise would decrease as the aperture radius is increased. Measuring the correlated noise as a function of aperture radius is a way of checking the accuracy of the aperture placement. Aperture placement would preferentially affect the brighter stars, as the faint stars are dominated by background noise.

In order to check for this possibility, aperture photometry was performed on the NGTS field of NG1416-3056 using three different aperture sizes: 3px, 4.24px ($3\sqrt{2}$) and 6px. Only the aperture radius was changed and the rest of the pipeline remains the same. Source detection, astrometry, calibration, photometry are completed as described in Chapter 2. The correlated noise was calculated for each target as discussed in Section 4.2. The total correlated noise for the different apertures was 4.19, 5.68 and 7.71 mmag respectively for the full set of stars. Using only stars with a flux of greater than 2000 ADU results in a correlated noise value of 3.11, 6.65 and 6.63 mmag respectively. The full results can be seen in Figure 4.9.

The correlated noise increases as a function of the aperture radius instead of decreasing, which is the opposite of what would be expected for an aperture placement error. Based on this it is unlikely that the aperture placement is a source of significant correlated noise for the stars, or in other words it is probably not the dominant effect. The increase in correlated noise due for the bright end may be due to blending with other stars or point to an issue with the background subtraction.

4.3.3 Airmass

The airmass correction and subsequent detrending of the photometry is one of the most important steps in the pipeline, as the trends in the lightcurves due to the varying airmass across the night are substantial. They dominate almost any kind of photometric signal. For example see Figure 4.10, which shows the normalised lightcurves from the field NG0304-1115 across a night. Furthermore the airmass trends are periodic, and so may contribute to the correlated noise

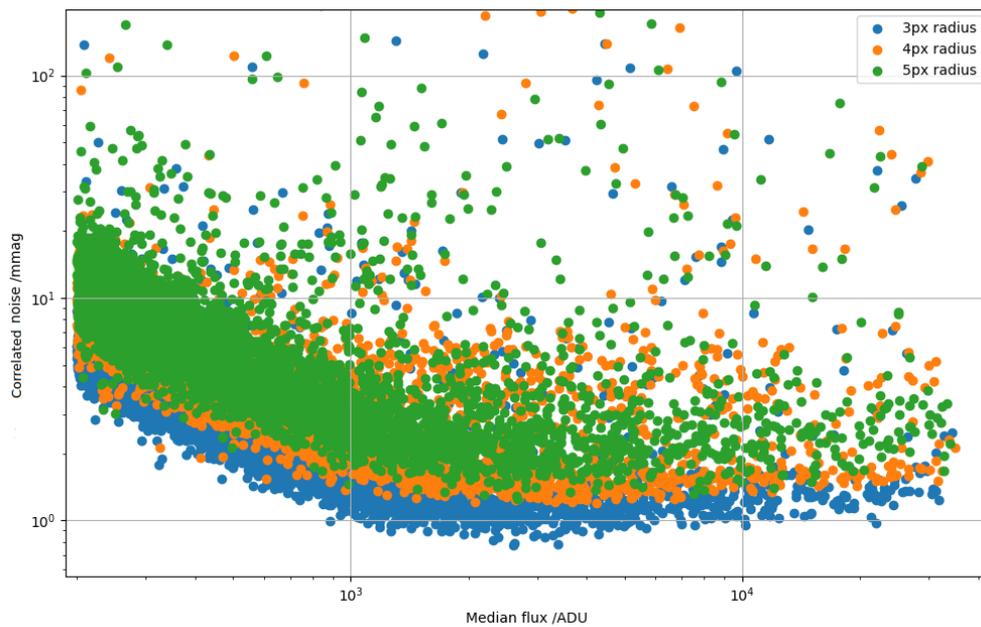


FIGURE 4.9: Correlated noise as a function of aperture size for the field NG1416-3056. The correlated noise was calculated for the same set of objects, but with each lightcurve using a 3, 4 and 5 pixel aperture radius respectively. The correlated noise values are then compared to the mean flux level for each target. The 3 pixel radius lightcurves have the lowest correlated noise.

budget of the lightcurves, particularly if the airmass subtraction is not perfect. The colour of the observed target may also play a part in how it responds to the airmass; differential refraction in the atmosphere could cause a star to change its position with respect to the rest of the cohort (in that particular field). The NGTS field of view (2.8 by 2.8 degrees) is large enough that there can be a sizeable difference in airmass from one side of the field to the other, which may further complicate the detrending process.

Looking for clues in the data for second order airmass effects is important, particular if they can be shown to correlate with the correlated noise. It may however be difficult to disentangle the airmass related effects from the aperture placement and the background correction, because any issues with these components may show the underlying airmass trends. Aperture placement due to differential refraction, and the background correction due to the increased background levels at higher airmass. Nevertheless this section will attempt to qualitatively assess whether the airmass correction is working as expected or not.

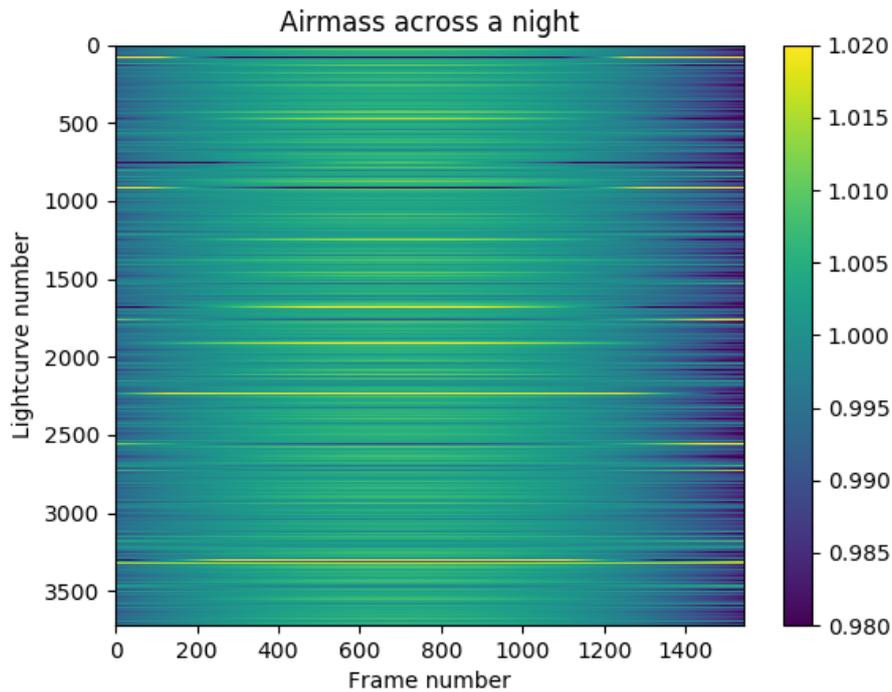


FIGURE 4.10: An example of the airmass trend across the lightcurves in an NGTS field for a single night. The lightcurves was all median normalised.

Airmass cut

The lightcurves are most affected at high airmass, where the path that the light travels through the atmosphere is larger. As such it may be that any correlated noise related to the airmass is more prevalent at high airmass. This could happen if for example if the airmass correction was systematically smaller than it should be. In such a case, the beginning and end of night of a lightcurve (in a field being observed while it rises and then sets) may show a dip.

To check if there is a correlation between airmass and correlated noise, an airmass cut is applied to the data. The field NG0304-1115 is used and the correlated noise is calculated for each full lightcurve, as well as the lightcurve composed of only the low airmass data. The field consists of 3,057 stars with 135,000 frames at 10 second cadence. SYSREM is not applied separately to the low airmass data in order to compare the two on an even footing and assess the quality of the detrending.

Figure 4.11 shows the results of the airmass cut, and while no formal test has been used, a small difference can be see between the two datasets. The low airmass data has a slightly lower correlated noise level. The correlated noise is

computed as in Section 4.2. The mean correlated noise level for the full field (including many faint targets down to NGTS I-band of 16) with complete airmass is 5.81 mmag, while the mean level for the low airmass lightcurves was 4.37 mmag. Removing stars with a flux below 1,000 ADU/second and repeating the comparison yields a mean correlated noise level of 1.39 mmag for the full airmass set and 1.26 for the low airmass set. Both of these levels are higher than the NGTS average shown in Section 4.2, perhaps as this is one of the few fields which were partially observed at this stage. SYSREM has been observed to perform better detrending on longer time series, though this has not been explicitly tested.

This result would indicate that any airmass residuals in the data are affecting the faint targets more than the bright targets. The airmass trend is a multiplicative effect, that is to say it is of the form $a(t) * F(t)$ where a is the airmass at that particular point in time and F is the stellar flux. This effect will be better resolved and removed from bright targets as opposed to faint targets where the error is dominated by the sky background. While the airmass corrections seems to work well, the SYSREM algorithm will also perform an intra-night normalisation of the lightcurves, attempting to compensate for systematic changes in the mean level of the lightcurves between nights. The next section will investigate whether there is an issue with this normalisation.

Intra-night Differences

The SYSREM algorithm attempts to compensate for systematic changes between nights which can occur in NGTS lightcurves. These intra-night changes could occur due to the overscan level, uncorrected for background light from the moon or other stray light, non-linearities in the chip and etc... Residual systematic differences between the mean levels of the lightcurves could contribute to the correlated noise budget, though some of the variability is likely to be due to intrinsic changes in stars themselves.

In order to check the effect of the intra-night normalisation on the correlated noise, each lightcurve is 'levelled out' on a nightly basis so as to remove any systematic changes. The median of each lightcurve for each night is calculated and then every night is offset so that the new nightly median matches the median for the full lightcurve. In this way any changes, systematic or due to real variability, are removed. The correlated noise of the old unaltered lightcurves

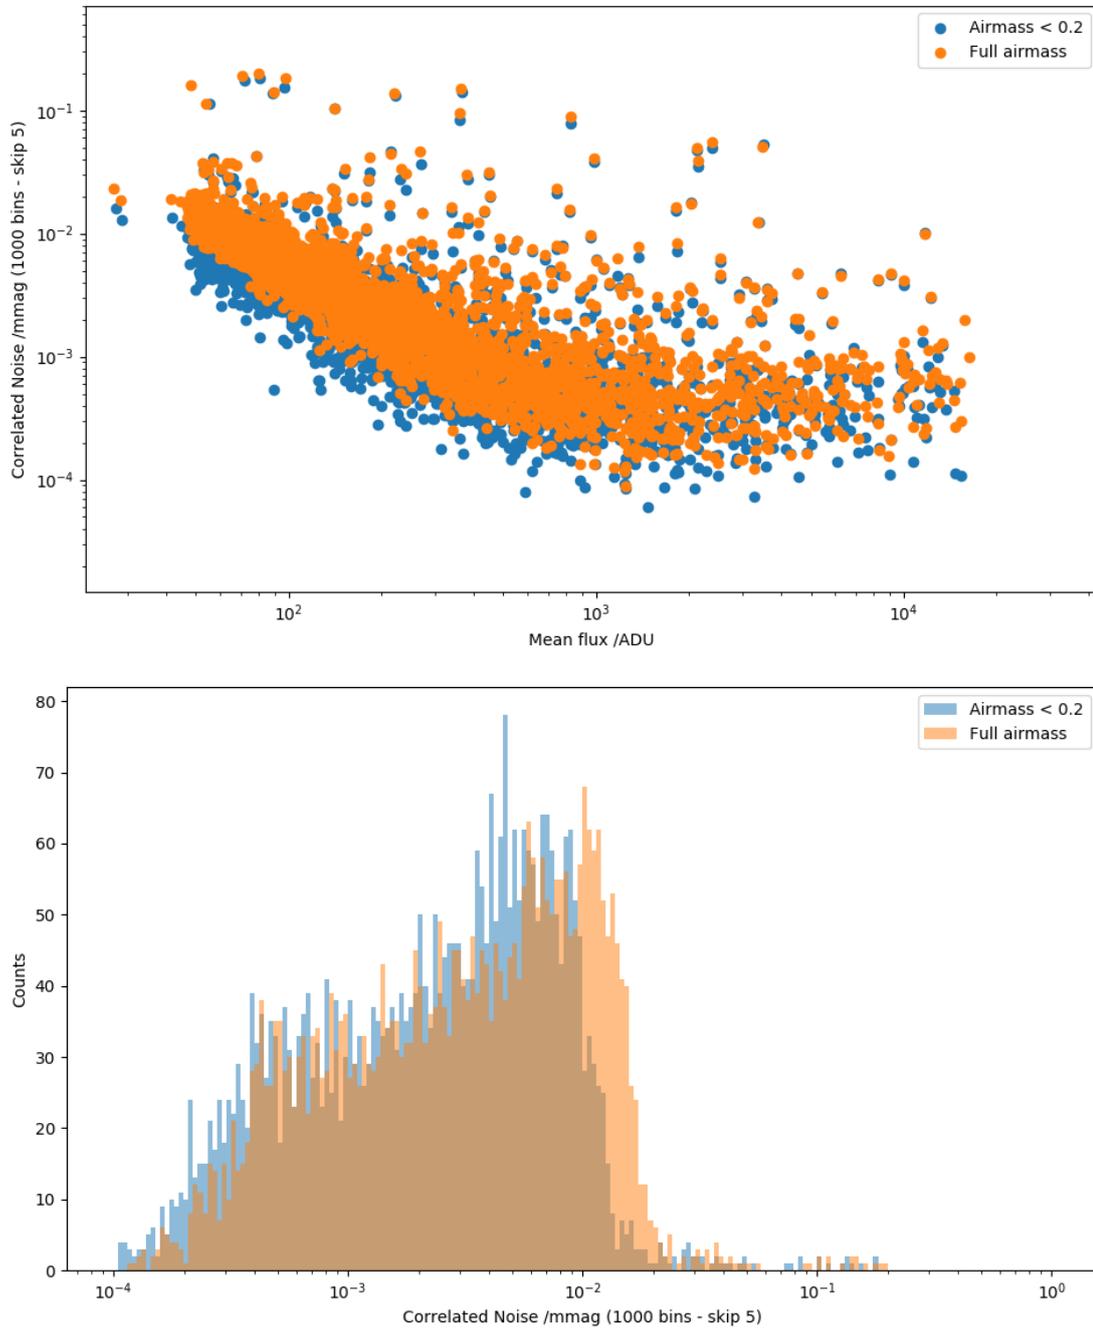


FIGURE 4.11: A comparison of the correlated noise level for field NG0304-1115 with and without using an airmass cut-off of 0.2. A small difference can be seen between the two sets of data.

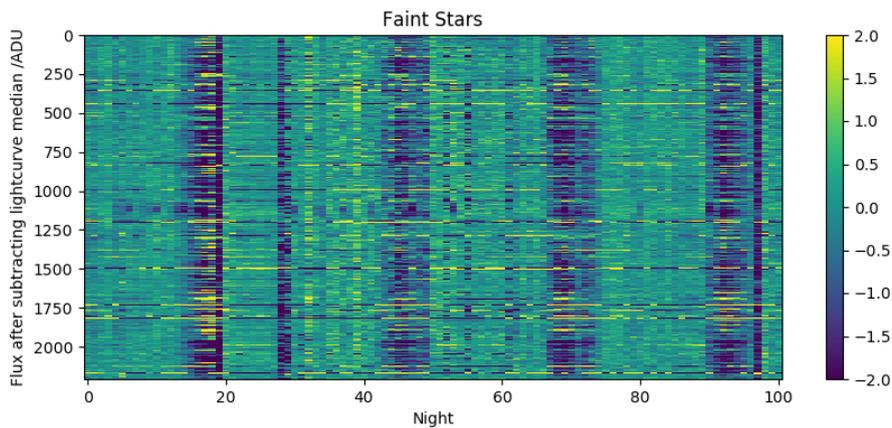


FIGURE 4.12: Values of the nightly median for faint targets (with a mean of less than 500 ADU). The scale is artificially set to show the small scale structure common to all the lightcurves. The mean value of each lightcurve is subtracted.

can be compared to the new change ones, in order to place an upper limit on the effect of intra-night changes.

The reason that each lightcurve is treated separately is that such systematic differences may also not affect the whole cohort of stars, but could affect different sets of stars differently. Using a nightly median for the whole set of stars could in fact make the correlated noise worse. Figure 4.12 shows what the nightly median is for all the faint stars (less than 500 ADU) in the field NG0304-1115 which is used for the airmass testing. The image scale is set to be ± 2 ADU to highlight the low level structure. Some of this may be due to clouds, but there are clear periodic bands which affect many of the stars. Removing some of this structure could reduce the correlated noise. Figure 4.13 shows the same plot but for the bright targets (> 500 ADU), however the targets are now normalised by their individual medians, to highlight systematic differences. Similar bands can be seen in the figure, though the structure is not as pronounced. The deep grooves in both the faint and bright targets are due to cloudy nights.

A subset of 1,000 stars are selected randomly from the field. After each lightcurve is normalised by the nightly medians, the correlated noise is calculated for both the altered and old lightcurves. The resulting plot can be seen in Figure 4.14 showing the correlated noise in ADU against the target brightness. The correlated noise is substantially lower at 0.84 mmag compared to the old correlated noise of 4.5 mmag for the set of 1000. The largest decrease is for the faint targets, while the brighter ones shows a smaller change in correlated noise. Some

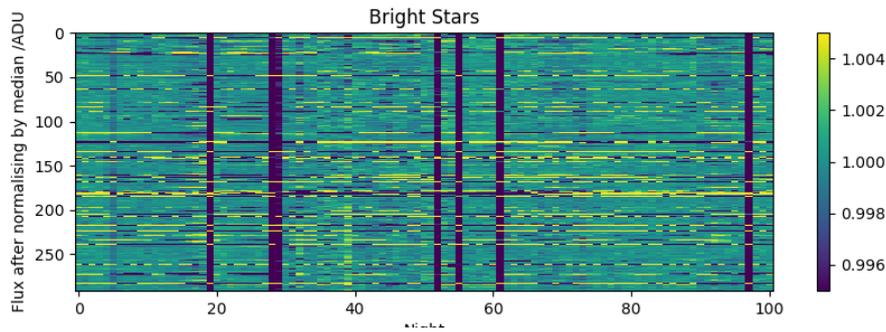


FIGURE 4.13: Values of the nightly median for bright targets (with a mean of greater than 500 ADU). The scale is artificially set to show the small scale structure common to all the lightcurves. Each lightcurve is divided by its global mean to normalise it.

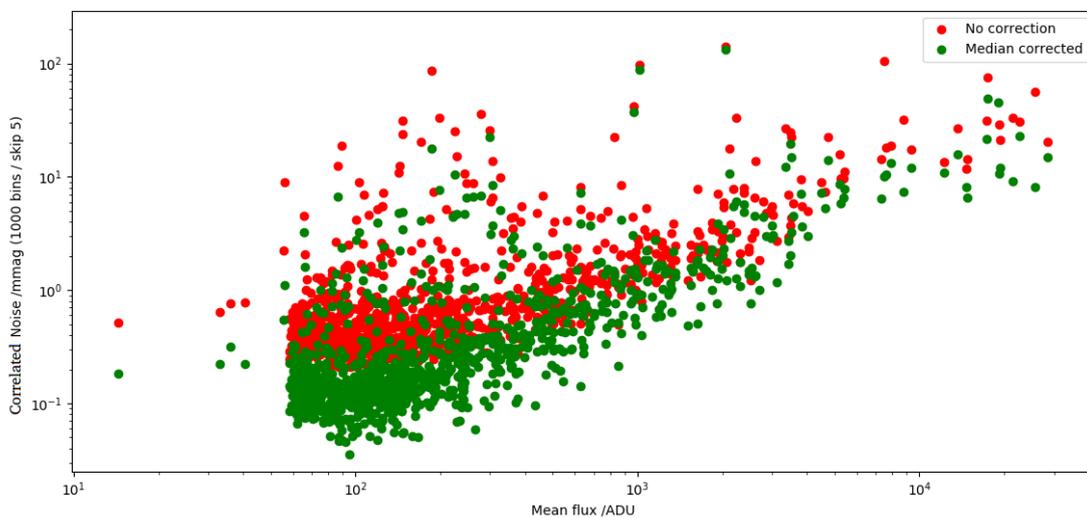


FIGURE 4.14: Comparison of the correlated noise levels before and after applying a nightly median normalisation.

of this change in correlated noise is likely due to intrinsic variability in the targets, particularly on the bright end. However, this shows that there may well be systematic differences due to the intra-night normalisation which if fixed could lower the correlated noise substantially. In practise it would be difficult to apply this algorithm to clean the data without making a careful assessment of the effect of the normalisation on the transit search.

The difference in the effect of the nightly normalisation may point to an additive systematic issue, a small constant offset which would preferentially affect faint stars over bright stars. This could happen for a number of reasons, but may also point to an issue with the background subtraction.

4.3.4 Background Correction

Background correction is an important part of the photometric pipeline. Diffuse light, which is scattered by the atmosphere, has to be removed before performing photometry on the stars, otherwise the lightcurves will show fluctuations due to changes in the sky background levels. These typically vary on the time scale of the lunar month and will also depend on the angular separation between the field and the moon. High level clouds and seeing will also change the mean background levels, though usually not by as much as the moon. The typical background levels observed in an NGTS field can range from 100 ADU on a dark moon free night up to as much as 2,000 ADU at a small separation from a full moon. This is a critical noise contribution, particularly for faint stars which are background noise dominated below 12th magnitude or so.

The background in the main NGTS pipeline is estimated and removed using the CASUTools software package¹, which uses an iterative median sigma-clip to estimate background values in a 64 by 64 pixel box (Wheatley et al., 2017). The values are then filtered and a full background image is created using bilinear interpolation. Full details can be found in Irwin (1997) as described in the documentation for the source extraction package². The CASUTools software package does not currently allow for the filtering parameters to be set, and so for testing purposes the SEP python package³ will be used. This package uses a similar background estimation algorithm, and so should produce results which are similar to the main NGTS pipeline. Further details can be found in Bertin and Arnouts (1996). The reduction, source detection and astrometric components of the pipeline will remain the same. The images will be calibrated using the same calibration frames and apertures will be placed in the same position and use the same aperture radius (3px). Only one component is different to the main NGTS pipeline, and so any changes in the quality of the photometry can be directly ascribed to the background subtraction.

The key purpose of these tests is not to perfectly replicate the photometry produced by the NGTS pipeline, but to understand the interplay between white noise and correlated noise which may be influenced by the filter width or the background box size, to check whether there is evidence of systematic issues and to find the optimal parameters for photometry. Lessons learned using the

¹<http://casu.ast.cam.ac.uk/surveys-projects/software-release>

²<http://casu.ast.cam.ac.uk/surveys-projects/software-release/imcore>

³<https://github.com/kbarbary/sep>

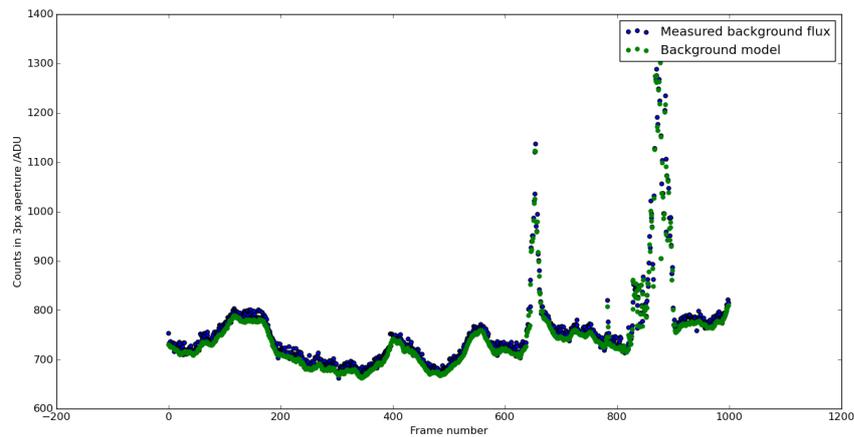


FIGURE 4.15: Comparison of the background flux between the sky and a model of the background. The sky value is determined by placing an aperture in a star-less part of an image, while the background model is built using the SEP python package.

more flexible SEP package can be readily applied to the CASUTools code, and it may be replaced in future if it is shown to be sub-optimal.

Aperture on sky

For the first check of the background the values derived from an aperture placed on the background image will be compared to those of an aperture placed in the same position on a reduced science image, where the 3px aperture is placed in between the stars. A location is chosen such that the aperture is as far away from nearby stars as possible. Using the SEP python package, a background map is created and an aperture is placed in the same position. The two fluxes are then compared in Figure 4.15. There is a systematic difference between the flux in background estimate and the actual aperture of 8.7 ADU. While this is only one example, it may be possible that this is part of a wider issue with the background correction.

Correlated noise vs Aperture radius

In Section 4.3.3 a test was conducted to measure the correlated noise as a function of the aperture radius. Apertures of 3, 4.24 and 6px radius were used to perform photometry on the NG1416-3056 field. The correlated noise was then calculated for each of the three apertures using the same method detailed in Section 4.2. Figure 4.9 shows the full results of the test, and as can be seen the

correlated noise increase dramatically with the aperture radius. For stars with a flux of less than 1000 ADU the mean value of the correlated noise was 4.67, 5.71 and 8.25 mmag respectively. The difference between 3px correlated noise and the 6px correlated noise is just under a factor of two.

One possible explanation for this is that the background correction is not working as expected. As the aperture radius increases the number of pixels increase as $\pi * r^2$. The increase in correlated noise does not match this prediction, but this could be because for example, perhaps only pixels near the star are systematically offset enough to cause a problem. Perhaps SYSREM partially compensates for the extra correlated noise. Nevertheless this, together with the results of the nightly median normalisation, is a strong indication that perhaps the background correction needs to be optimised in the NGTS pipeline to produce better quality lightcurves. The next section will deal with testing and optimisation of the background correction under a range of parameters to find solution with the lowest possible correlated noise.

4.4 Optimisation

As shown in the last section, one possible component which may be contributing to the correlated noise in the NGTS data is the background subtraction. In order to reduce the correlated noise, and increase the survey yield, it would be useful to correct the background subtraction or to eliminate it as a possible cause and look for improvements elsewhere.

4.4.1 Background correction

The background subtraction for NGTS is performed globally, and using the CASUTools packaged. A similar method is implemented in the SEP python package, which allows for easier customisation of the parameters of the algorithm. Visual inspection of the output of both packages over many hundreds of frames shows that they produce very similar results when building a background map. Using SEP we may investigate the correlated and white noise of the properties using the techniques discussed in Section 4.2 and then apply the result to the NGTS pipeline. By optimising the parameters in this way the correlated noise may be reduced.

Filter Width	16px Box	32px Box	64px Box	128px Box	256px Box
2px	8.08	8.16	3.92	3.92	8.40
3px	8.09	7.73	3.92	3.92	8.40
4px	7.48	3.93	3.91	3.94	3.96
5px	8.16	7.64	3.91	3.94	3.96

TABLE 4.4: Mean correlated noise in mmags for all targets in the field as a function of different background subtraction parameters, calculated using 20,000 images.

Filter Width	16px Box	32px Box	64px Box	128px Box	256px Box
2px	31.36	29.37	33.13	33.34	30.64
3px	31.36	30.44	33.13	33.34	30.64
4px	31.52	33.13	33.20	33.62	34.26
5px	30.58	30.21	33.20	33.62	34.26

TABLE 4.5: Mean white noise in mmags for all targets in the field as a function of different background subtraction parameters, calculated using 20,000 images.

For the optimisation photometry is performed using 20,000 NGTS images from the field NG0624-2518. Each time the parameters which define the background correction are changed, varying the box size and filter width. The resulting data set is SYSREM corrected and then the correlated noise is measured for each. The box size parameters are chosen from the set of 16, 32, 64, 128 and 256 pixels and the filter width is chosen to be 2, 3, 4 or 5px. Photometry is performed on the set using the full combination of parameters. Using the subset of 20,000 images yields a correlated noise of 4.31 versus 4.19 for the full set of images. These are comparable and it is expected that a steep decrease in correlated noise in the subset should transfer to the full set of data. Photometry is performed using a 3px aperture radius and using the same x- and y-positions for all cases.

4.4.2 Optimisation results

The results of the optimisation can be seen in Table 4.4 and Table 4.5. The white noise varies only slightly as a function of the different parameters with a 32px box and a 5px filter size producing the lowest noise. A few other parameters come in close, but no discernible pattern can be seen. The correlated noise is more sensitive to the parameters, varying by a factor or two. However, the variation seems to be almost binary, between a ‘good’ correlated noise value of about 3.9 and a ‘bad’ correlated noise value of 7.9. This may indicate that there may be underlying effects which mean that no one combination of box and filter size

may be suitable for all data. The currently used parameters for NGTS correspond to a box size of 64 px and a filter width of 3px, which is a fairly optimal choice. It seems likely that there isn't a strong gain to be made by change the parameters. Switching to an alternative method, such as local background subtraction may yield better improvements.

4.5 Conclusions

The correlated noise is measured and evaluated for the NGTS experiment to be approximately 1mmag across the full dataset available at the time of the study. Future pipeline changes planned for the later half of 2018 will likely change this value. The data is then investigate to search for correlations between the correlated noise and pipeline components such as the flat field and background correction. From these tests there is some evidence that the background may be contributing to the correlated noise. Finally, steps are taken to investigate the effect of different parameter choices in the background subtraction in an effort to reduce the correlated noise. These tests prove to be inconclusive, and so may point to improvements needing to be made further downstream in the pipeline detrending.

Chapter 5

Bayesian Change Point Detection and Single Transits

5.1 Introduction

On an average night, NGTS observes an 8.3 square degree field of the night sky for anywhere between one to eight hours over a period nine months. With a 10 second exposure time, this amounts to an average of about 2,000 frames per night per field per camera. Each field can contain anywhere from 3,000 to 20,000 stars with brightness down to ~ 16 th magnitude in the I-band, the chosen cut-off for the NGTS photometric pipeline. (Many fainter targets are observed down to approximately 18th magnitude, but are not included in the main transit search.) This is an enormous amount of data and provides an unprecedented opportunity to study the variable sky and test new approaches to finding transits, transients and other interesting astrophysical events. In this chapter I will discuss the use of Bayesian change point detection (BCP) in the context of finding single transit candidates in NGTS lightcurves as well as the detection of transients and variables.

Change point detection (CP) or change detection is a family of methods for determining the most likely point where the underlying statistical distribution of a time series has changed (PAGE, 1957). Change point detection can encompass both single and multiple change points, both on-line and off-line methods and can be used to search for changes in mean, variance or a combination of both. Figure 5.1 shows examples of a CP method search for different changes in a time-series. CP is a general purpose technique which has multiple applications from spam filtering to quality control in manufacturing and website tracking, also forming a key part of the methods commonly used in anomaly detection (Prescott Adams and MacKay, 2007).

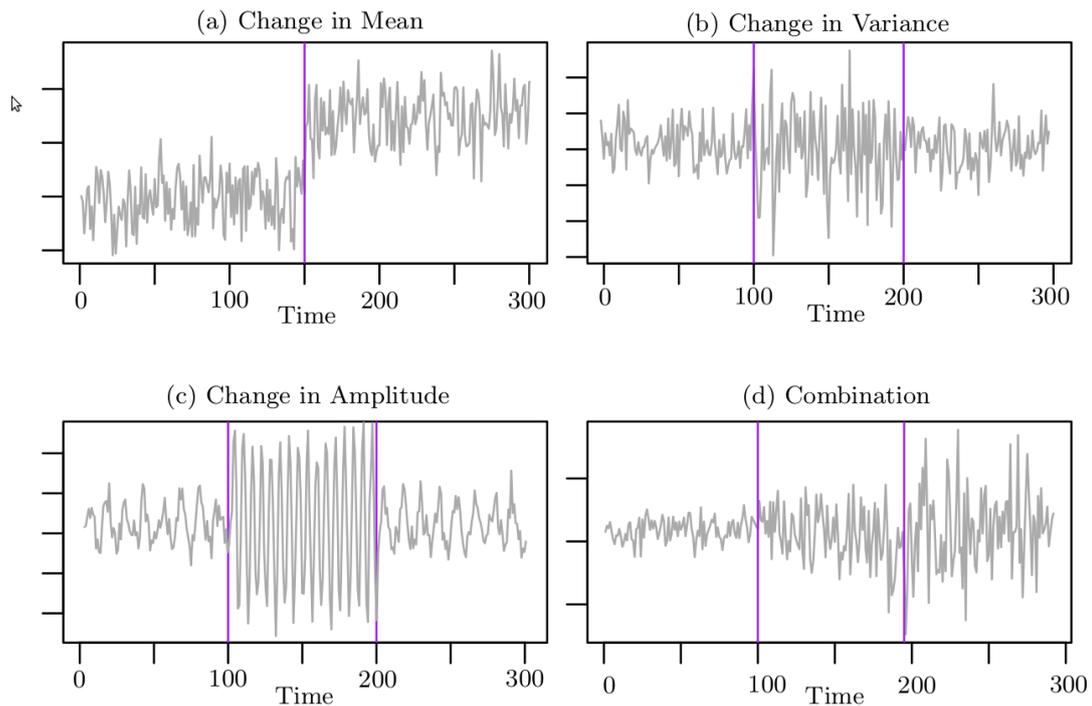


FIGURE 5.1: Examples of change point techniques looking for different changes in the underlying statistical distribution of the time series. Image source: <http://www.lancs.ac.uk/~chapmanj/changepoint.pdf> - Jamie-Leigh Chapman

Single Transits

Within the context of NGTS, the primary goal of this study will be to apply Bayesian change point detection techniques to search for a change in the mean of the stellar flux, which would happen when a planet is transiting across the star. There are of course many other reasons why stellar flux can change, for example pulsations, star spots, as well as the plethora of astrophysical false positives which create transit like signals such as blended objects, binaries and background eclipsing binaries. The former set of false positives should be identifiable, for example by excluding active stars with many variations or by excluding transits which are ‘V-shaped’ and much deeper than 1% of the brightness of the star. The second set of false positives are common to any ground based survey and will be dealt with by careful follow-up observations of any interesting targets which may be uncovered.

Space based missions such as Kepler and K2 are ideal for the detection of single transit events due to their long uninterrupted observing windows. Some work has already been done to identify single transit events in these datasets using machine learning and Bayesian techniques, for example Foreman-Mackey,

Hogg, and Schölkopf (2015) and Osborn et al. (2016). While space is the ideal platform to find single transits, ground based surveys like M-Earth (Nutzman and Charbonneau, 2008) have also search for such events (Dittmann et al., 2017). Ground based surveys have the advantage of targeting bright host stars compared to previous space based missions and so may increase the sample of warm Neptunes and Saturns which can be readily characterised (Osborn et al., 2016). Few such systems are known. Single transit events, may also provide constraints on the formation mechanisms of circumbinary planets (Brown and Armstrong, 2015). Finding single transit events that may be present in the NGTS dataset is an important compliment to the main planet hunting mission.

Variables and Transients

A secondary goal of this project is to discover variable stars and transients in the NGTS data, where transients refer to events occurring in the existing lightcurves, or a-periodic variability, and not to the detection of new targets. (As discussed in Chapter 2 NGTS photometry is list driven and in order to detect new targets a technique like Difference Image Analysis (Alard and Lupton, 1998; Alard, 2000; Bramich, 2008) would need to be used.)

Three main types of techniques exist for detecting variability and they are the use of variability indices, periodic search and template fitting. Box-least squares fitting (Kovács, Zucker, and Mazeh, 2002) is an example of a template search, and other variables of known shape can be searched for in this way. Variability indices rely on the computation of a number to quantify the amount of change or scatter in a lightcurve, thereby making it easy to find targets which are more variable than the majority. Two examples of this technique are Welch and Stetson (1993) for finding variables and Udalski (2003) for finding micro-lensing events. More complex methods such as the machine learning have also been applied, for example Richards et al. (2011), Brink et al. (2013), and Armstrong et al. (2016) among many others.

BCP will be used in this context, and it is hoped that it will prove to be a powerful technique for not only identifying variables, but also finding unusual events which may be missed without a directed search. For example the potential detection of an exoplanet with rings, J1047b, by the SuperWASP experiment (Kenworthy and Mamajek, 2015). While BCP is more complex than an index search, it is as flexible, and when applied to many lightcurves simultaneously provides the potential of ignoring time correlated changes by filtering common change

points in the cohort. BCP is also not dependant on the event being periodic as are many techniques.

5.2 Change Point Detection

Many techniques exist for performing change point detection, it is a general problem which can be solved in many different ways depending on the nature of the dataset under investigation. For processing NGTS data we focus on probabilistic, Bayesian methods, which have already been implemented in freely available code, as packages for the R programming language¹². As this work is a proof-of-concept for single transit detection using change point detection, it is not necessary to find the ‘globally optimum’ change point model or technique for the data, but one which produces results of a sufficient quality.

Initial Testing

As an initial test one hundred NGTS lightcurves with a single injected transit were processed using the Barry and Hartigan (1993) (BH) technique as well as the Binary Segmentation (BINSEG) (Scott and Knott, 1974) and the PELT algorithms (Killick, Fearnhead, and Eckley, 2012). The transit depth was chosen to be three times the standard deviation of the lightcurve, with a width of 2.5 hours. The transit was injected in a random night of data, chosen from a selection of nights consisting of over 6 hours of data. The lightcurves were drawn at random from the full NGTS data set of targets in the magnitude range of 8 to 13th in the NGTS I-band. The ingress point was chosen at random to be with the 1st and 2nd, ensuring that the beginning and end of night affects due to high air-mass were minimised. (Though of course, this depends on the exact field chosen and the time of year).

Visual inspection of the results showed that all three methods did a reasonable job of picking up the transit ingress and egress. The BINSEG method was the slowest, while both the PELT and BH methods finished faster taking only linear time or $O(n)$ in big O notation, to complete the calculation. The BH technique produces a posterior probability for every point under the assumption that the errors in the lightcurve are normally distributed, while the PELT method returns

¹<https://cran.r-project.org/web/packages/bcp/index.html>

²<https://cran.r-project.org/web/packages/changepoint/index.html>

an ordered list of the optimal n change points along with a penalty for each. This is the penalty value, which is chosen, would make the change point significant.

Both the penalty value and the posterior probability can be calibrated by using transit injection and recovery tests to determine the appropriate value for different signal to noise injections. However, the posterior probability is much easier to interpret, particularly as no additional parameters need to be chosen in order to run the algorithm. The PELT method requires the choice of an appropriate cost function and a minimum segment length, both of which will affect the penalty values. Optimising for single transits may affect the ability of the PELT technique to pick up transient events or other variables. The BH method will return a value for each data point, and the underlying assumption, that the errors are normally distributed is a reasonable approximation for NGTS lightcurves. While a may suitable model may exist, for a proof-of-concept the BH provides a convenient and easy to interpret change point method, which will be used to analyse the data in the rest of the chapter.

Barry and Hartigan (1993)

In this section a brief overview is given of the BH method, using terminology and wording as discussed in the original paper. The BH method is designed to model "a sequence of observations undergoes sudden changes at unknown times" (Barry and Hartigan, 1993). The process is modelled by supposing that for every time series there is an underlying sequence of parameters partitioned into continuous blocks of equal parameter values. Between consecutive blocks there are said to be change points. A product partition model "assume(s) that the observations in different components of a random partition of the data are independent," (Barry and Hartigan, 1992). The product partition is used as a convenient way to weight which partitions with the data, thereby allowing an inference to be made about the change points in the dataset.

The resulting problem can be solved exactly in $O(n^3)$ or to a good approximation by using MCMC in $O(n)$. This allows the BH method to applied onto long time sequences, though binning the NGTS data to 10 minutes greatly speeds up the computation (factor of 60). Further details can be found in Barry and Hartigan (1992) and Barry and Hartigan (1993).

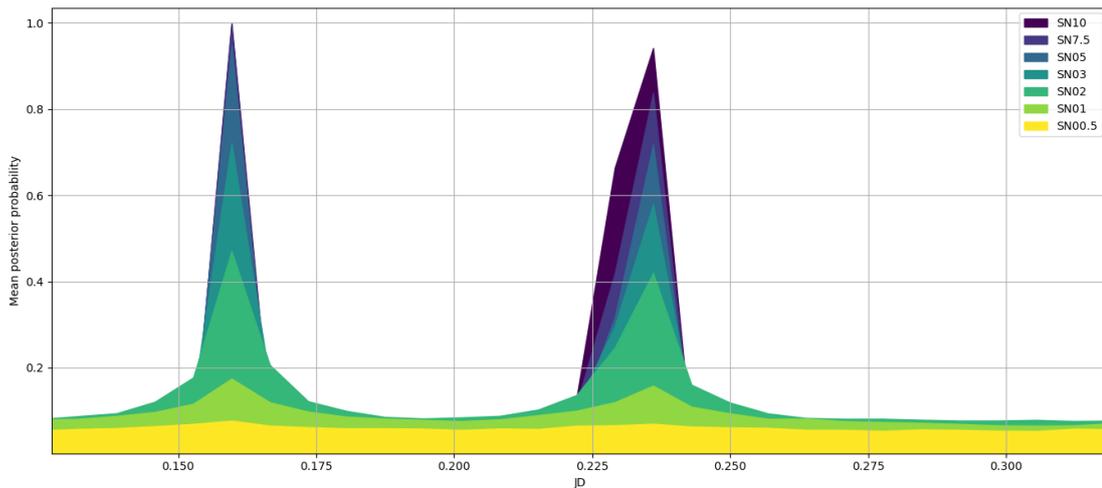


FIGURE 5.2: Single transit injection into randomly generated white noise. The mean values of the posterior probability are evaluated on a lightcurve which has been binned to a 10 minuted cadence. The mean is calculated over 1,000 repetitions per signal to noise level.

5.3 Simulation

Before applying the BCP algorithm to the NGTS dataset it is necessary to run several simulations to determine whether single transits would be detectable using BCP and also under what conditions that would be the case. Transit injection and recovery tests also help to determine what a single transit signal might look like, making it easier to search for one in the real dataset. For example, what would the posterior probability be of a transit event of 5 mmag in a lightcurve with 1 mmag noise? Considering that we would like to detect long period mini-Neptunes and super-Earths, how realistic a case is this? Is the probability affected by the transit duration? How do the systematics present in NGTS affect the result? To what extent can they be mitigated by running BCP over the whole cohort of stars in each field and making full use of this information.

5.3.1 White Noise Test

To address these questions a series of transit injection and recovery tests were performed using the PyTransit module (Parviainen, 2015) and the BCP change-point module³ written in R. For the first test, a single transit was injected into a lightcurve generated using Gaussian white noise. Each transit was set to be

³<https://cran.r-project.org/web/packages/bcp/bcp.pdf>

three hours long and the depth was chosen to be either 0.5, 1.0, 2.0, 3.0, 5.0, 7.5 or 10.0 times the standard deviation chosen for the noise. The length of the light curve was set to be eight hours with a ten minute cadence, and so consists of 48 data points. The transit takes up over a third of the night. The lightcurve was then processed using the BCP code and the posterior probabilities at each point were saved. This process was repeated 1000 times at each value of signal to noise to get a more accurate understanding of how well the single transit is picked up.

Figure 5.2 shows the results of the transit injection and recovery simulations in white noise lightcurves. The results indicate that, for a white noise lightcurve, a transit of signal to noise 10.0 down to 2.0 should be detectable with a posterior probability of >0.4 . For comparison the typical posterior probability for a random data point is about 0.1 to 0.2 in a random NGTS lightcurve within the magnitude range of 8 to 13. For transits of lower signal to noise, a detection becomes more difficult. A signal to noise of 1.0 produces a probability of just under 0.2, which is only a bit higher than the baseline probability for the rest of the points. In practice, it would likely be very difficult to identify such a signal in the data, unless perhaps it was periodic. In that case it would likely be picked up by a box-least squares search of the data and so would already have been found by the main pipeline. A true one off low signal to noise event, while very scientifically interesting, would be incredibly difficult to identify and follow-up as there would likely be many false positives in the data.

An interesting feature in 5.2 is that the probability at the ingress is a little higher than the corresponding probability at egress. The algorithm finds it easier to identify the initial change point than the subsequent one where the flux returns to its normal baseline. If that time series is reversed the result remains the same, however the effect disappears when performed for real data (see Figure 5.3).

5.3.2 Transit Injection into Real Data

While a white noise model is useful to confirm that it is possible to recover a single transit in a night's worth of data, it is not an accurate model of the NGTS data. Real lightcurves always contain a degree of correlated noise which is known to impede the detection of transits. As such the transit injection and recovery test needs to be repeated with real NGTS lightcurves to confirm that

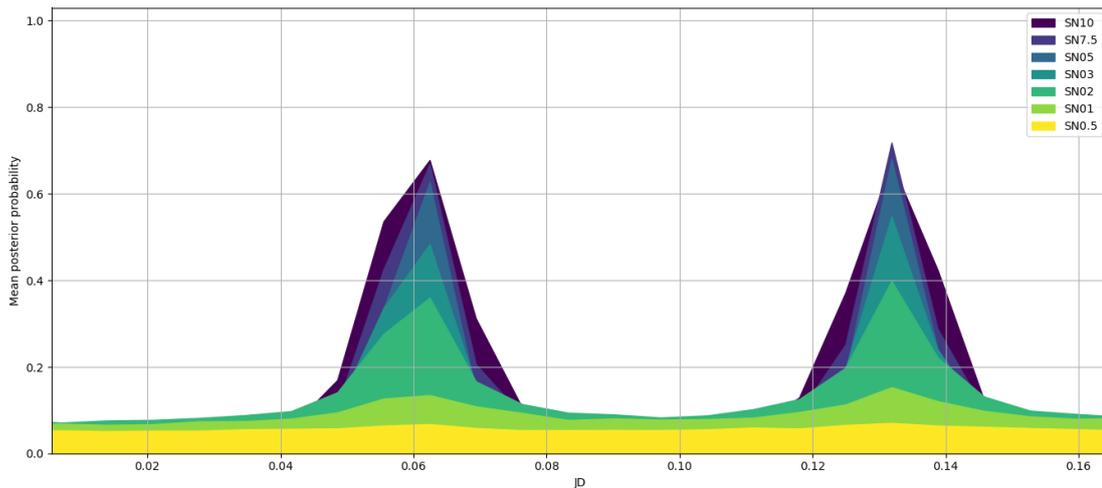


FIGURE 5.3: Single transit injection into NGTS data. The lightcurve is binned to 10 minute cadence and the mean posterior probability is calculated over 1,000 test lightcurves per injected transit.

transits can be recovered and also to roughly quantify what kinds of probabilities are expected. Figure 5.3 shows the results of a transit injection and recovery test completed using real NGTS lightcurves. Each lightcurve was selected at random from those available in the TEST18 dataset, comprising of 43 fields and 386,959 stars. Once the lightcurve was selected, a single night, longer than five hours, was chosen. In total over 3,000 single night lightcurves were selected from different targets, all of which were brighter than 1,000 counts which roughly corresponds to an I-band magnitude of about 13. Most of the transit detection activity focuses on these stars due to a higher signal to noise.

The transit duration remains the same as that of the white noise test, though the lightcurve may be shortened by up to two hours. The lightcurve could be made artificially longer, by combining two nights for example, but this approach is not ideal as there are known beginning and end of night systematics in the data. Placing a transit close to the either end would likely decrease the probability of the detection, and give a poor indication of the ability to recover the transit. When searching for single transits we are looking for a complete event, as it would be difficult to reliably identify an ingress or egress, so we are not interested in replicating this scenario.

A bootstrapping approach was considered, where a lightcurve is composed from many smaller randomly shuffled sections, but this may not preserve the structure of the correlated noise which could artificially inflate the strength of the detection. The transit ingress is also moved to the start of the night as opposed

to the middle of the night as in the white noise test. Since an off-line BCP algorithm is chosen, it doesn't matter where the transit signal is relative to the start of the lightcurve. The whole lightcurve is processed at once.

Figure 5.3 clearly shows a reduced strength of the detection of the transit signal. This is likely due to the presence of correlated noise in the NGTS lightcurves and perhaps due to a shortening of the length of some of the lightcurves. It is interesting to note that beyond a signal to noise of 3.0 there did not appear to be a significant difference in the strength of the posterior probability. Both the ingress and egress level seems to be similar unlike the white noise example, but they are noticeably wider in the correlated noise tests. This indicates that there is greater uncertainty in the start and end times of transit compared to the white noise model. The mean posterior probability of the sections before, after and during the transit are similar in both cases at around 0.1.

The conclusion from this test is that the NGTS data contains systematic differences from a white noise lightcurve even on a short time scale. These somewhat decrease the confidence with which we can detect transits using BCP, however it should still be possible to find single transits in the real data with this technique.

5.4 Data Cleaning and Processing

5.4.1 Data Cleaning

Prior to processing the data using BCP, each lightcurve is sigma-clipped using a 5-sigma clip. This is to remove outliers which are detected as very strong change points, as the BH algorithm is not designed to take into account outliers or uncertainty in the data. This step helps to keep down the number of erroneous change points. NGTS lightcurves also contain a number of points which are set to zero, usually due to bad weather or clouds at the time. The flags are not always perfect, and so some cloudy data may remain. These points are also removed prior to processing.

5.4.2 Removing Correlated Change Points

One of the advantages of using the BCP algorithm is that the location of many change points are correlated in time. Correlated changes points are most likely due to external factors such as cloud or the moon level as opposed to intrinsic

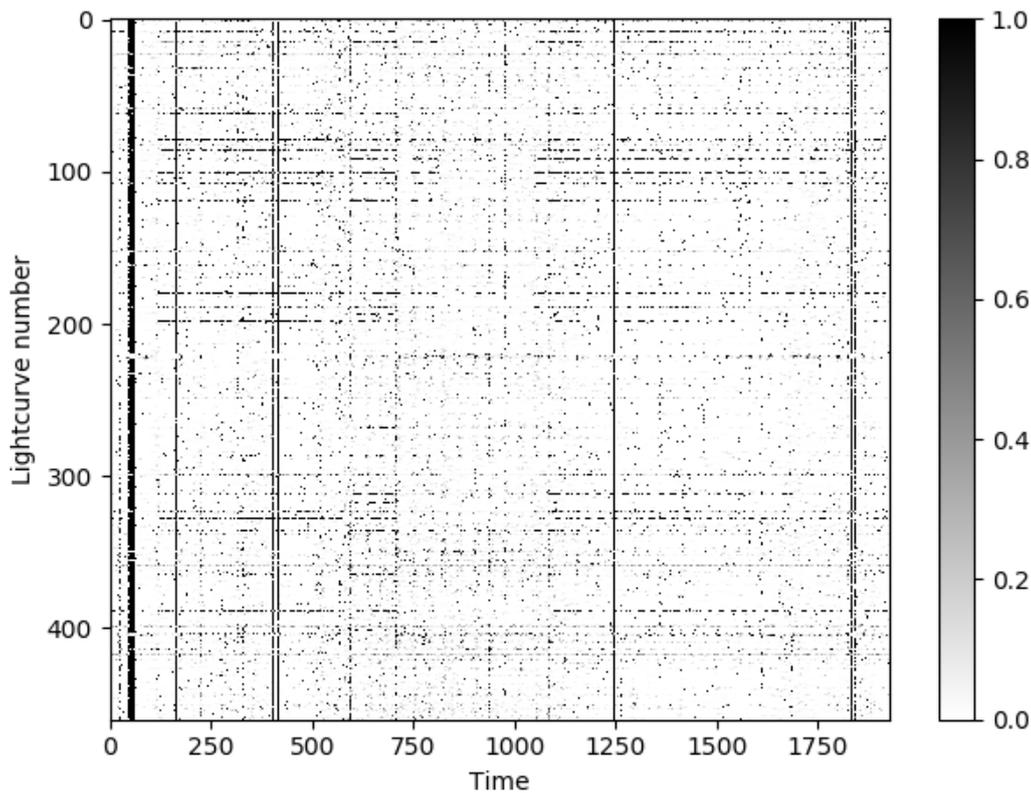


FIGURE 5.4: The posterior probability of all the lightcurves from field NG0314-2518 after processing with BCP. The flux was binned to 10 minutes, keeping night boundaries intact. Strong vertical lines are change points occurring across every lightcurve and are likely due to external factors.

changes in the stellar flux. An example of this can be seen in Figure 5.4, where the vertical lines indicate a change point which is common to many lightcurves. As such they should be removed or down-weighted and ignored. Some of these change points may even be related to the general correlated noise found in the lightcurves.

To remove the correlated change points, several simple schemes were attempted. However the most successful scheme was to remove the mean posterior probability at each point in time and to divide by the standard deviation. This can be written as

$$p_{i,t} = \frac{p_{i,t} - \bar{p}(t)}{\sigma_p(t)}, \quad (5.1)$$

where p is the posterior probability, $\bar{p}(t)$ is the mean posterior probability across all the lightcurves at that point in time and finally $\sigma_p(t)$ is the standard deviation of the posterior probability at a point in time.

As a test data set the NG0341-2518 field was used. This contains 3,650 lightcurves

with 1890 flux measurements which were binned to a cadence of 10 minutes. Each lightcurve was processed using BCP and the final result was collated and shown in Figure 5.4. The 'cleaning' scheme was applied and the result shown in figure 5.5. The change points now range in value from approximately 3.0 through to 14.0. This is a representation of the "significance the detection" and is not directly related to the significance of a detection of a transit of a transit for example. Instead it highlights the strength of a certain change point with respect to the rest of the data taken at that point in time.

From examining Figure 5.5, it is immediately apparent that most of the correlated change points have been effectively removed, however many change points remain across the dataset which have a high significance. Variable stars can be seen by eye as dashed horizontal lines going across the figure. For this field a cut-off of 7.0 was chosen and all change points below this level were discarded, leaving approximately 3% of the set. The values were then converted back to a probability to make it easier to interpret the results in terms of the simulated injection tests.

The lightcurves and their corresponding change points were inspected before and after applying the cleaning scheme. Visually the number of change points due to erroneous data points were greatly reduced. This leaves many change points which are likely due to intrinsic changes in the targets. The lightcurves still contain many change points at the beginning and end of night, however this doesn't appear to be true in all targets, the fainter targets in particular don't seem to have this issue. This might explain why not all the beginning and end of night change points were effectively removed.

In future a more rigorous test will be used to assess the cleaning scheme. Transit should be injected into some small percentage of lightcurves at random, and the ability to recover these transits will be tested with and without using the cleaning scheme. Additionally the number of false positive candidates in both cases needs to be estimated.

5.5 Results

5.5.1 Single Transit Search

At present, a very simple method is used to screen for potential transits using the detected change points in a field. After the change points are 'cleaned' using

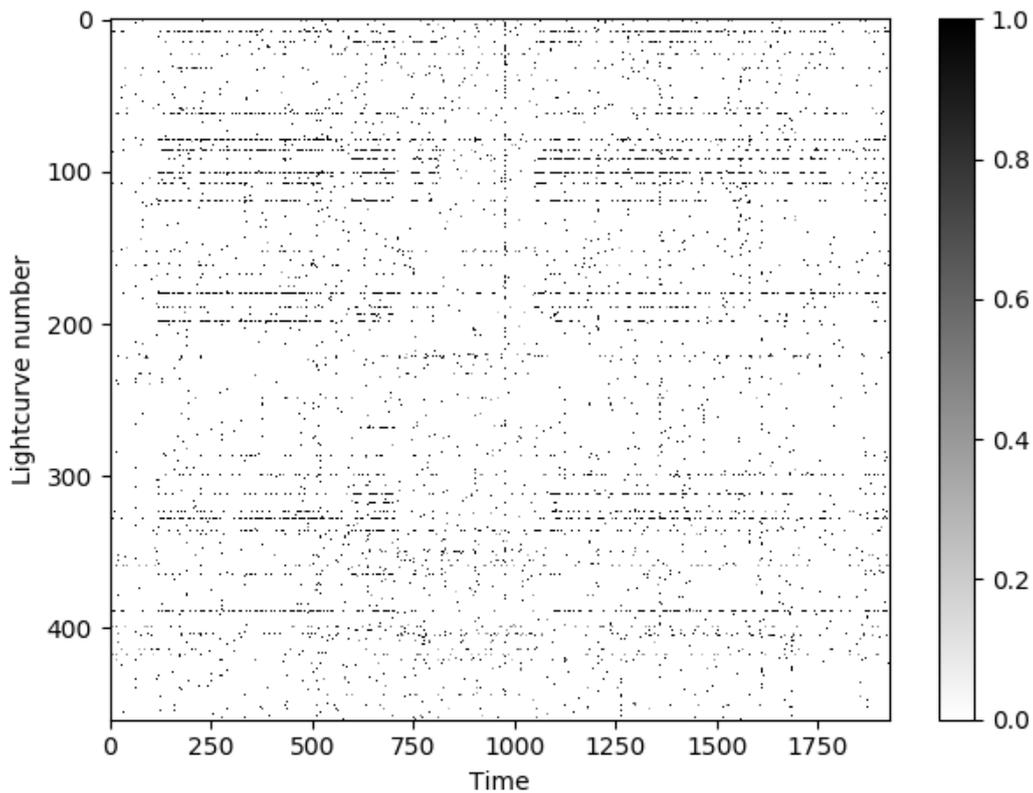


FIGURE 5.5: The posterior probability of all the lightcurves from field NG0314-2518 after it has been cleaned using the scheme described in section 5.4.2.

the scheme detailed in section 5.4.2, each lightcurve is split into segments based on the separation in time between significant change points. The segments are then filtered to find those which are below the mean level of the lightcurve for that particular night. These are then filtered based on the time elapsed between the change points, more than an hour and less than five hours as well as the difference between the segment mean and the nightly mean with the segment removed. This must be less than 2% and greater than 0.1%.

The set of lightcurves produced in this way are then filtered to remove those with more than three candidate transits per night, as well as lightcurves which are in the top 10% in terms of number of significant change points. The first step is necessary to remove candidates where many transits are detected, as these are likely picked up by the main BLS search. If more than three transit like events are detected in a lightcurve it is automatically excluded from consideration. The second step is taken to remove variable lightcurves, by making a cut on the average posterior probability per target, see section 5.5.3. The threshold is chosen such that the top 1% of lightcurves are removed. The resulting candidates

are then checked by eye. None of these parameters have been fine-tuned and the search procedure has not been carefully evaluated against injected transits. There is a likely quite a bit of scope for improvement, however for the purposes of detecting the most easily found single transits this setup is sufficient.

For the single transit search 18 NGTS fields covering the RA range from 03:00:00 through to 06:00:00 hours are processed. This amounts to a total of 140,130 targets and 335,821,116 data points (at 10 minute cadence). Or an average of 2396 points per target, which amounts to 16.6 continuous days of data spread out over an observing season. Many of the fields were still being surveyed at the time of processing and so only a partial season was observed. The data was processed using the BH Bayesian change points algorithm and cleaned as discussed in section 5.4.2. The data is then checked for single transits.

From the dataset of 140,130 targets 5,096 are selected as having a potential single transit. This is just over 3.6% of the dataset, which is a relatively large number. The transits are first sorted by magnitude (NGTS I-band) and then by the signal to noise of the transit depth. This is given by the difference between the mean of the change point segment and the rest of the lightcurve divided by the standard deviation of the lightcurve excluding the segment between the change points. This orders the list to aid the manual vetting of the candidate single transits. In future several other criteria can be used, such as the chi-squared of the fit of a transit to the lightcurve and the ratio of transit depth to the colour of the star. Many of the candidates can be excluded this way as most don't actually look like transits, but are unusual deviations in the lightcurves or occasionally due to linear trends which start halfway through the night. Many false positives are also due to beginning and end of night change points, which will be filtered out more carefully in future.

The sorting of the transit candidates is designed to reduce the eye-balling and bring the best candidates to the top of the list. In future this might be solved automatically using a decision tree based approach, such as a Random Forrest or a Gradient Boosted Decision tree. Decision tree based methods are where cuts are made on the lightcurve 'features' such as magnitude, transit depth and duration, in order to find the best set of candidates. This would be done on a simulated single transit injection first, before being applied to the full dataset once the chosen features and parameters were optimised.

The first thousand candidates were checked by eye to find promising transit signals, of these 24 candidates were selected as being promising in the first pass. A

further 5 were rejected in a second pass, leaving 17 potential candidates for vetting. The lightcurves now need to be carefully assessed for their level of systematics and out of transit variations, the transit depth and fit and any other transit like features in the lightcurves. Once this vetting has been done the candidate transits which remain can be properly fit, along with the colours of the host star to determine the potential planet size. Other vetting steps such as a centroiding check (Günther et al., 2017) also need to be undergone before the candidate is suitable for radial velocity follow-up. Here the best candidates were selected to be presented in this section as a case study for the initial vetting.

Figure 5.6 and figure 5.7 show the single transit candidate NG18598. Figure 5.6 shows a close up of the 6.8 mmag transit, which lasts 2.5 hours. The transit shape is seen both in the un-binned data as well as the 10-minute data. The rest of the lightcurve shows some variable particularly at the start and end of observations. There is a 15-20 day period in the middle of observations where the lightcurve is quiet, during which the single transit occurred. The remaining lightcurves do not show strong variability on the night of the transit making it unlikely to be due to cloud. It is unlikely that the cloud would present so consistently over a two hour period. The lightcurve is carefully checked for additional transits which are not seen. A BLS periodogram needs to be run to try and identify if there are ingress/egress of transits in the lightcurve which are consistent with the main single transit.

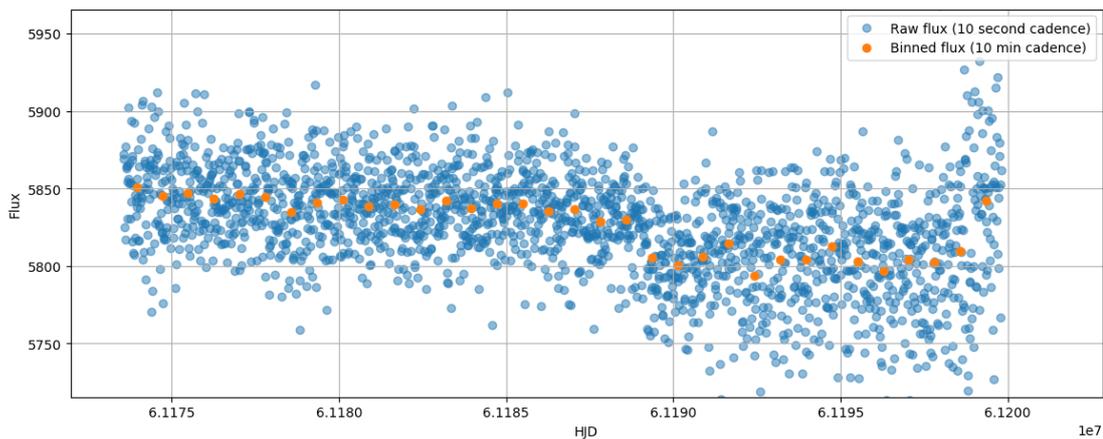


FIGURE 5.6: Plot of the night of the single transit.

NG18598 is a good candidate for further vetting, follow-up and characterisation. Future work would involve carefully fitting the transit and estimating the stellar radius in order to determine the rough size of the planet. In addition it is necessary to attempt to constrain the period in order to attempt ground based

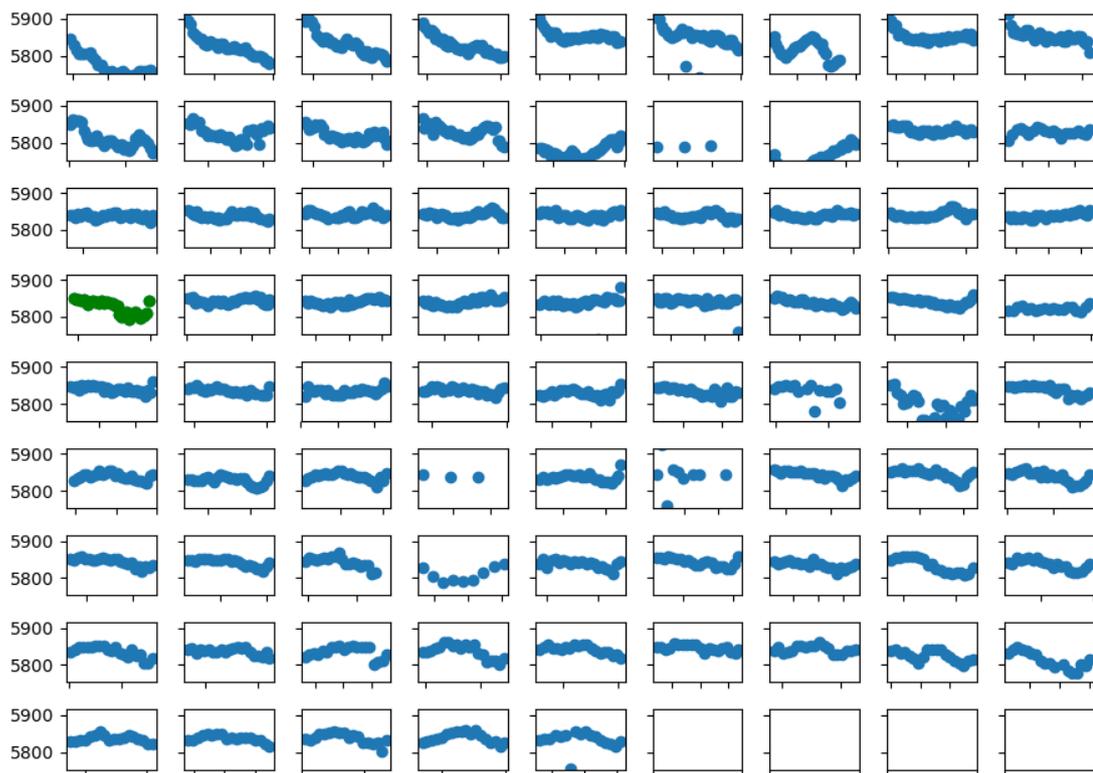


FIGURE 5.7: Single transit candidate NG18598, full lightcurve. Cadence is 10 minutes. Each panel is a separate night of observation (77 in total). Green is the night of the single transit.

transit photometry, ideally in a different pass band. NG18598 and other candidates demonstrate the potential for the use of Bayesian change point detection to find single transit candidates in NGTS and other data. Future work to carefully characterise the algorithm using transit injection and recovery tests will provide a useful comparison with existing techniques.

5.5.2 One-Day Period Search

The BLS algorithm produces a power spectrum with peaks where likely transit signals may be found. All ground based transit surveys have window functions which are due to the daily gaps between observations. (Perhaps with the exception of transit surveys taken at high latitudes where it is possible to observe continuously, like for example the ASTEP survey, see [Crouzet et al. \(2010\)](#) and [Mékarnia et al. \(2016\)](#).) Poor weather can also produce gaps in the data, but because they are not regularly spaced it does not typically influence the window function. This causes a practical difficulty as strong periodic signals are often detected in the NGTS survey at 8 hours and 1.0, 2.0 and 3.0 day periods, making

it very difficult to detect planets at those periods due to the dominant window function. The one day period is particularly strong and may mask real planetary signals close to or just below the orbital period. While the relative prevalence of ultra short period (USP) planets is relatively low (Sanchis-Ojeda et al., 2014), they are scientifically interesting population of planets due to their unusual formation history (Lopez, 2017), and may be detectable around the brightest NGTS targets.

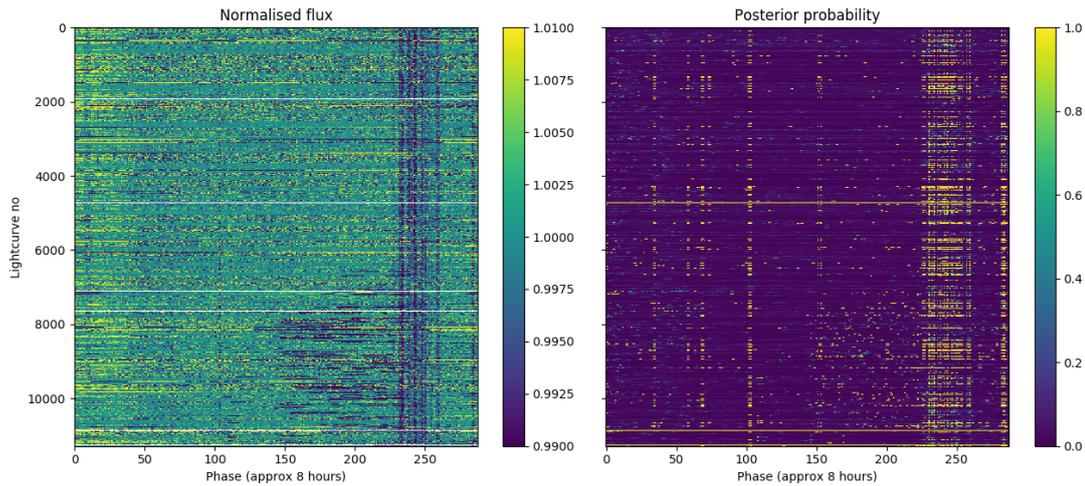


FIGURE 5.8: Comparison of the flux (left) and the posterior probability (right) of the lightcurves folded to a 1.0 day period.

As shown in section 5.3 it is possible to detect a transit survey in a single nights worth of data, and as shown in section 5.4.2 it is possible to remove spurious change points which are correlated in time. Here we investigate the possibility of detecting planets in the 1-day period range by running BCP on data which has been folded to that period. The test will be applied to the NG0603-3056 field, which contains 11,285 target and 85,730 flux measurements at a cadence of 10 seconds. The lightcurves are first phase folded to a one day test period and are then binned to two minute cadence. The Heliocentric Julian date (HJD) is used for the phase folding. The binning is performed to decrease the noise, as well as to make the folded lightcurves easier to manage in terms of memory usage and speed of processing. The final bin, which contains less than two minutes worth of data is discarded, to prevent the introduction of any systematic differences. Before binning any erroneous data points are clipped out using a 5-sigma clip. The folded lightcurves cover approximately one third of the phase for the 1-day period, due to day time gaps.

The folded and binned lightcurve are then processed using the BCP detection algorithm used for the single transit search. The resulting output is then cleaned

and searched for transits in the same way as before. Running BCP on the folded data is a more effective way to search for 1-day periods, as it would allow to better compensate for the 1-day systematics. Other signals are averaged out in the folded data. Out of the 11,285 targets 534 are selected as having a potential transit like signal. These candidates are then screened by eye, temporarily in place of a more systematic approach (see 5.5.1), to find any potential interesting transit like shapes. One candidate transit is seen, shown in figure 5.9. A 2nd-order polynomial is used to detrend the candidate lightcurve in order to better assess the transit shape. Figure 5.9 shows the full lightcurve over the course of the observing period, some variability can be seen which might discount this target as being suitable for follow-up. Close inspection of the lightcurve, night by night, does not yield any strong single transits or any strong individual events which may lead to a transit like shape in the folded lightcurve. However, the beginning of the nights do show some transit like features.

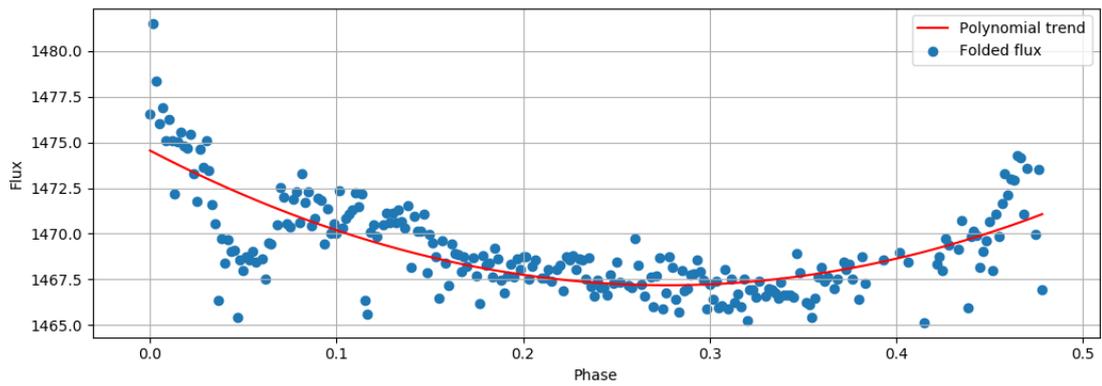
Further photometric follow-up is required to confirm or disprove the candidate. The transit depths is 2.7 mmag around a 13th magnitude star. While the star is a bit fainter, it should be possible to reach the requisite signal to noise with a 1-m follow-up telescope. Before any further follow-up is done the period of the transit and the epoch also need to be refined.

This simple case study shows that BCP may be suitable to search for transiting candidates at periods which are inaccessible to BLS. Further works is needed to decrease the number of false positives and to rank the candidates for easy vetting. This should be straightforward to optimise using synthetic transit injection and recovery tests, in a similar way to the main single transit search.

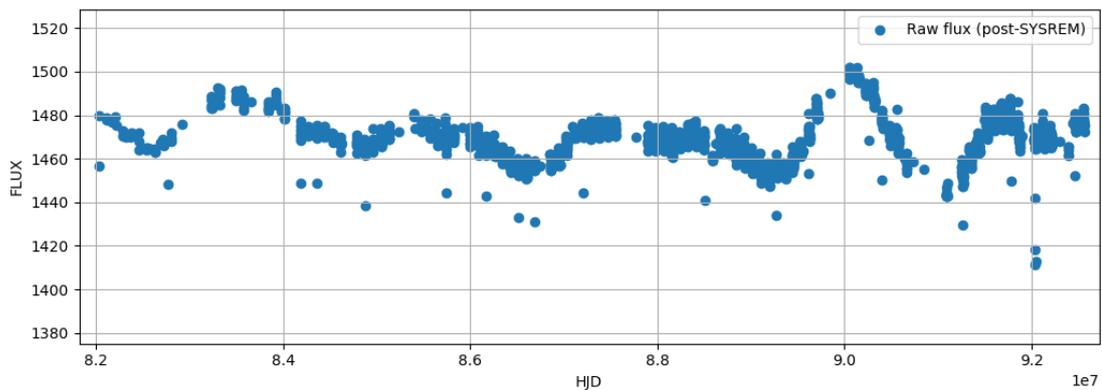
5.5.3 Variable and Transient Event Search

Variable lightcurves can be identified by looking at the average number of change points per target. Those with a large number of change points, either consistently throughout the season, or in one part of the lightcurve, may be showing some intrinsic variability. By focusing only on the number of change points, and not the spacing between them or their relative strengths, we are not excluding one type of variable for another.

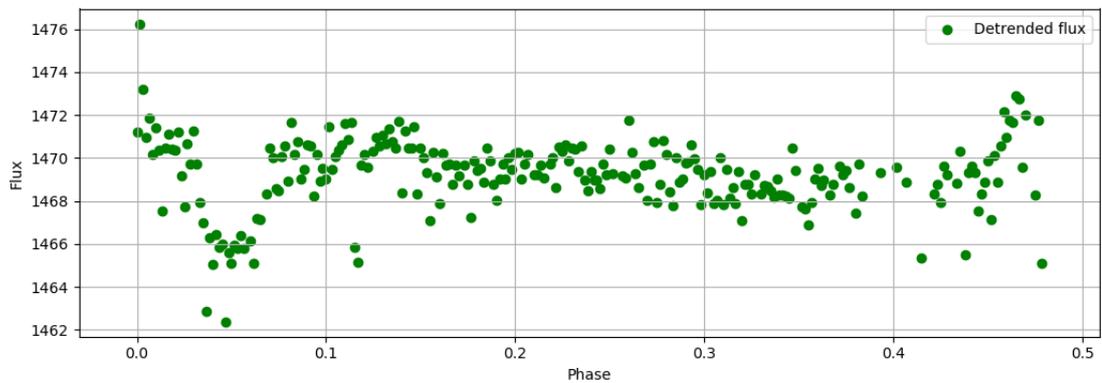
In this section a single field, NG0603-3056, will be searched to find variable light curves. The field contains 11,285 stars each with 1,431 measurements at a cadence of 10 minutes. Any short transients events, for example a flare, may be



a)



b)



c)

FIGURE 5.9: Plots of the folded flux for the transit candidate.

missed due to the binning up of the data. In future these tests may be repeated on the un-binned data set as well. The posterior probability is calculated using the same BCP algorithm and cleaned as before, prior to the variable and transient event search.

Variables

In order to search for variables, a set of lightcurves is selected which have a very large number of change points of strong significance. Specifically the average posterior probability per data point is calculated for each target and can be seen in figure 5.10. Several interesting features can be seen in the plot; the average posterior probability increases slightly with the flux of the target, there are many targets in the faint end with a large number of change points and similarly the brightest nearly saturated stars also have a large number of change points. The number of change points for the brightest stars is likely due to systematics in the lightcurves from saturating the pixels on the CCD. On the faint end this may be due to excess noise, or there may be many more true variables due to the sheer number of faint targets.

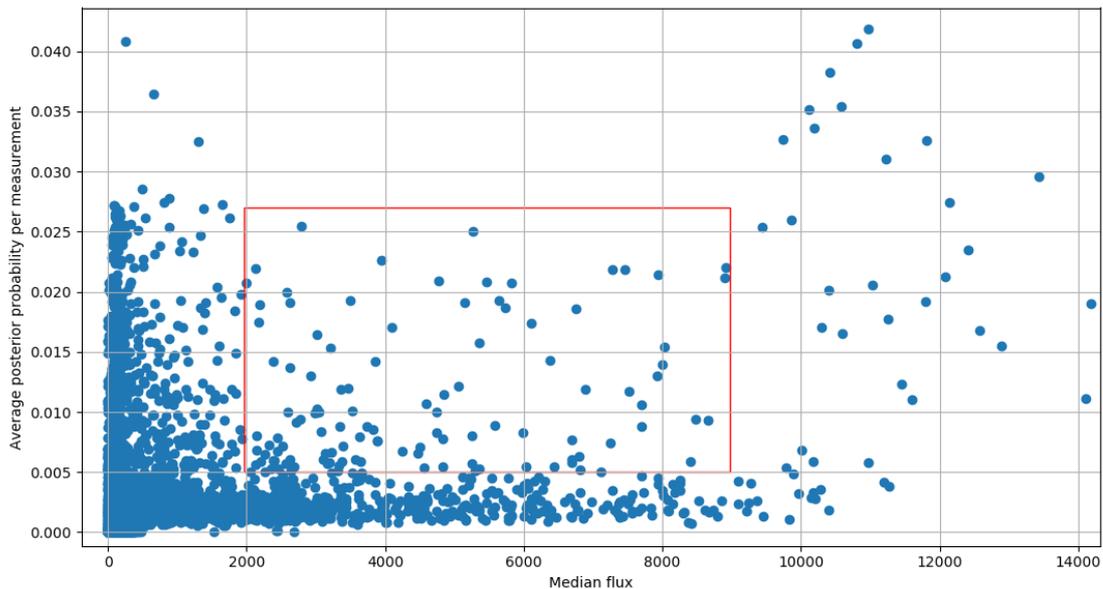


FIGURE 5.10: Plot showing the distribution of the average posterior probability compared to the mean lightcurve flux.

For this simple study the median flux of the stars will be restricted to those which are higher than 2000 ADU to reduce the number of considered stars. A selection cut is made, illustrated by the red box in figure 5.4.2 with a lower limit of 0.005 and upper limit of 0.027 (average posterior probability). Only targets in the range of 2,000 to 8,000 counts are considered, effectively removing faint or saturated stars and limiting the size of the data set. Lightcurves within this range are deemed to be variable.

The cut returns 288 potentially variable targets from a selection of 626 potential targets meeting the brightness criteria. The candidate variables were checked

by eye and only one of the 288 candidates showed clear signs of variability. The remaining 388 non-candidates were also checked by eye and none showed clear variability either, indicating that there were no false negatives. The dataset is too small to draw any conclusions from particularly given the lack of strong variability, nevertheless it is likely that a more stringent cut should be made. Raising the threshold from 0.005 to 0.015 keeps the variable (average posterior probability of 0.025) but cuts down the false positives to 31. Figure 5.11 shows a comparison between the detected variable and a false positive selected at random. Even the detected variable isn't a particularly 'clean' variable, but some sinusoidal modulations can clearly be seen. The false positive shows a lot of changes in the lightcurve but none which look like any coherent variability.

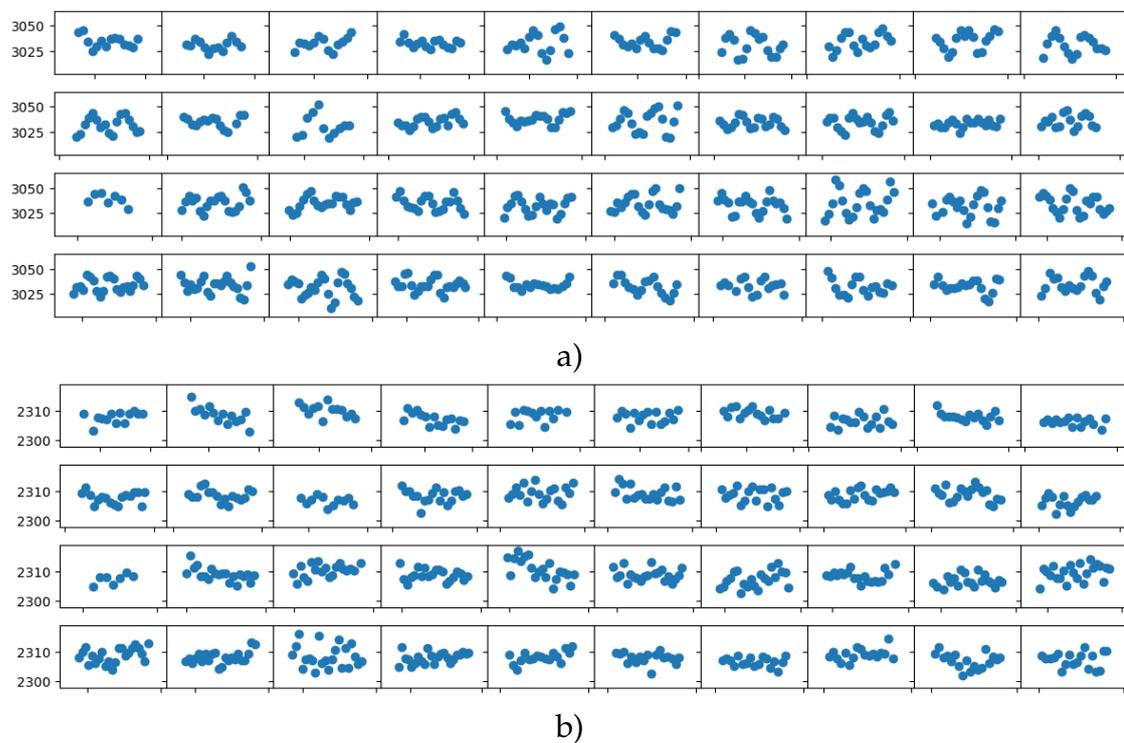


FIGURE 5.11: Plot showing several nights from a detected variable a) and a detected false positive b). The y-axis is the flux of the targets and the x-axis is time. Each panel represents a distinct night from the lightcurve.

The large number of false positives appear to be due to real changes in the lightcurve, which due to the relatively high signal to noise are being picked up as significant, often at a probability of >0.8 . Many of these change points occur at the beginning and the end of the night for each lightcurve. In future it may be possible to filter these out using a better scheme, but also to change the BCP algorithm used, in order to make it easier to identify more significant variability. Most other change point algorithms do not report a probability but a

'significance' value of some sort, which is not capped at 1.0. There may be correlations between not only the time of spurious change points but also their relative strength, something which was only indirectly accounted for in the cleaning scheme used in this work.

Transients

A similar approach is taken to identify transients in the data set. Instead of looking for a lightcurve with a large average posterior probability, those are filtered out using a threshold of 0.015 (the same one used in the variable search), leaving only relatively quiet stars. The average posterior probability of a change point is then computed for each night of the lightcurve. A cut-off is then selected to pick out nights which have unusually large posterior probabilities. Lightcurves which have such nights are then flagged up as being of interest. As a next step the lightcurve is investigated to see if the change is due to a systematic issue or actual variability. The goal of this test is to identify events which may be transient in nature and may otherwise be missed by a typical variable search. Of particular interest are 'unusual' events, for example the transit of a planet with a ring system or a dusty tail to name a few, which are rare and may not be spotted in a large dataset.

The same field, NG0603-3056 is used in the proof of concept search for variables and transients. For a first step the sum posterior probability is calculated for each lightcurve and those with values higher than 0.0015 are eliminated from the search (see Figure 5.10). For the remaining targets with a median flux greater than 500 ADU, the average posterior probability is computed per night for each object. Figure 5.12 shows a histogram of the average posterior probability per night for these field. As can be seen there are a range of values for the different nights with the vast majority being below 0.2. The periodic peaks or spikes seen in Figure 5.12 are due to a few short nights which have strongly correlated outliers due to clouds and beginning and end of night effects. Each target has a similar number of change-points for that night which leads to the unusual peaks. The process of removing removing correlated change points somewhat reduces the number of these events, but does not get rid of them entirely.

A cut-off threshold of 0.5 is chosen to screen for nights which show strong variability. This results in 241 objects which have at least one night with a posterior probability matching this criteria out of a total of 2,417 targets which meet the 500 ADU lower cut-off. Of these targets only 17 shows signs of being legitimate

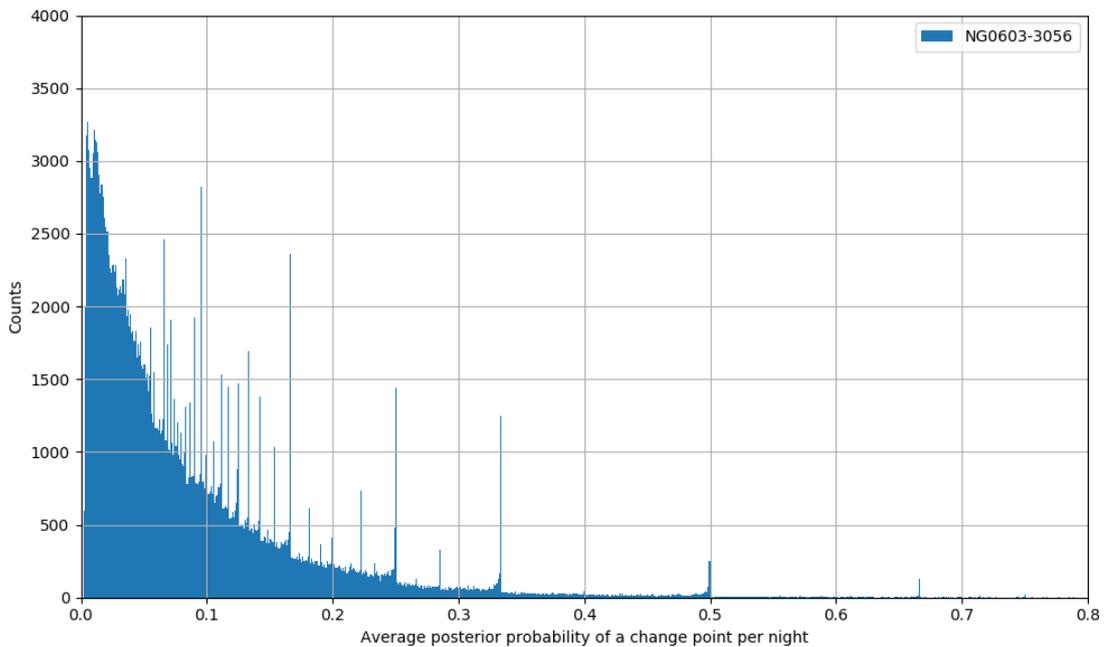


FIGURE 5.12: Histogram of the average posterior probability of the change point per night for the field of NG0603-3056. Certain nights show higher values due to intrinsic variability, while others do so due to systematics in the data. The histogram can be used to select an appropriate cut-off for screening for transients.

transients. The vast majority of which are due to systematic issues with NGTS data. Only 3 targets show real likely transient behaviour. The remaining 224 objects are composed largely of targets which were triggered by an especially short night of observations, consisting of 30 minutes of data, with two strong change points. The night is the same across many of the targets and so this is one way it can possibly be filtered out. The underlying cause of the change points are due to cloud. Applying a minimum cut-off of 2 hours for the shortest length a night can be filters the false positives down to 26 targets.

The two of the three targets, 6760 and 9051 show short flare like events of with an increase of about 10% of the flux. The events show the sudden rise and gradual comedown associated with flares. Target 9051 has two events, only the second one of which was picked up in the transient search. The threshold cut-off used will likely need better fine tuning. Target 1926 shows a particularly deep eclipse over the course of approximately three hours. Two more smaller eclipse events are seen later in the lightcurve. Further study of the target is needed to determine what may be causing this variability, however the occultations do not appear to be correlated in time with other similar events in other lightcurves. They may be due to target specific systematics.

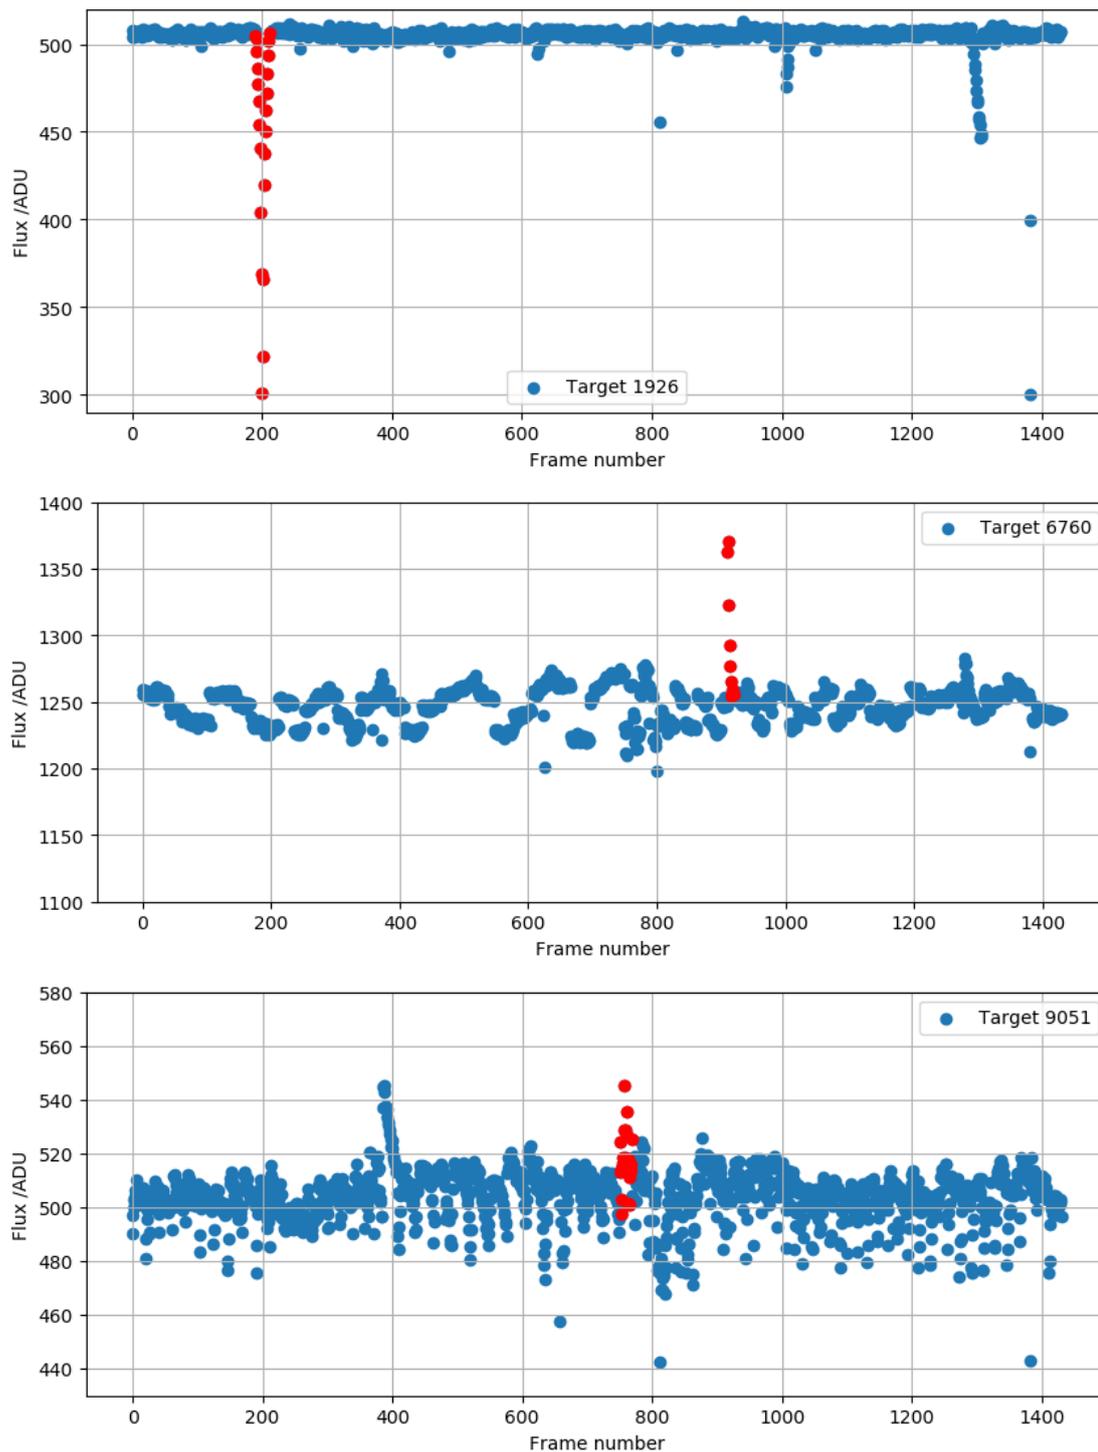


FIGURE 5.13: Plots of the three transient events detected in the field NG0603-3056. The x-axis is the binned frame number and not the Julian date. This is done to better shows the full lightcurve structure and level of variability. The top and bottom lightcurves contain two and one additional events respectively which are also detected but not marked. The data was binned to 10 minute cadence.

While this work shows that transient events may be detected using changepoint analysis, much further work is needed to fine-tune the algorithm. A very crude strategy was used and in practise better results could be derived. For example, the dataset contains shorter nights which are of the order of 1-2 hours and longer nights up to 8 hours so it might be useful to down-weight shorter nights. This will certainly cut-down the false positives. For this work the lightcurve was assessed on a night by night basis, it may be that using a rolling 3-hour window is a better choice. Though of course the choice of the window length will determine what kind of transients will be discovered.

5.6 Conclusion

Changepoint analysis was applied to NGTS data with the primary objective of showing that it was possible to detect single transits in the data with the technique. The Barry and Hartigan (1993) algorithm was chosen to be used. Transit injection and recovery tests showed that it was indeed possible to recover transits using this method and the technique was applied to several NGTS fields. A total of 24 single transit candidates in a first pass through the results. One particularly good candidate was presented.

BCP was also applied to search for transiting exoplanets with a 1-day period, in phase folded data. One candidate planet was found in the dataset, though with a very high false positive rate of 533 additional targets. Rigorous injection and recovery tests are required to properly bench mark false positive and false negative rates on single transit recovery and 1-day period searches. A more rigorous study could also be used to compare BCP to other existing techniques for single transit search.

Changepoint detection was applied to search for transients and variables in the field NG0603-3056 with mixed success. No viable variables were found in the brightness range searched, and it is unclear whether there is any advantage to using BCP for this task. The transient search returned 3 targets and 14 due to NGTS systematics along with a number of false positives.

This work showed that as a proof-of-concept these techniques can in principle work to find transits, variable and transient. Several candidates were produced for NGTS follow-up, and when applied to the full dataset it is likely that many more candidates will be found. In future changepoint detection might prove to

be powerful and useful technique for mining astronomical datasets, particularly once it have been properly fine tuned and rigorously tested.

Chapter 6

Photometric Follow Up of NGTS Candidates

6.1 Introduction

Photometric follow-up is a key part of confirming and characterising potential transiting exoplanet candidates from ground and space based missions. From the ground, a high precision transit lightcurve can confirm that the planet is not a result of instrumental systematics and can update transit parameters and ephemerides. From space, a transit lightcurve taken in multiple colours can provide a transmission spectrum of the planetary atmosphere and greatly refine planet parameters. At the current stage of the NGTS survey the most important use for photometric follow up is to confirm potential detected candidates as real and not systematics, as well as to provide a second independent transit lightcurve for known NGTS planets. This chapter discusses a follow-up program undertaken in the South African Astronomical Observatory using the SHOC camera on a 1-m telescope; covering candidate selection, planning and taking the observations, reducing the data and finally producing photometry. Observations were conducted over the course of three weeks in July and November of 2016. A total of 14 candidates were observed, allowing 13 candidates to be rejected for secondary follow-up. High precision lightcurves were produced for one candidate, which will go on to further study. Due to the high signal to noise requirements of the follow-up, a custom pipeline is developed to reduce the data and perform photometry. The development and performance of the pipeline is discussed. Finally a discussion is also included on the precision achieved with the 1-m telescope in SAAO using SHOC and the potential of the site for further follow-up work and exoplanet science.

6.2 Candidate Selection

Before each observing run an extensive candidate selection process was undertaken by the full consortium to find the best possible targets to be observed at the time. A plethora of information is available with which vetting can be done; stellar colours from previous surveys, past SuperWASP lightcurves and follow-up observations, odd-even variations, the box least squares (BLS) power spectrum, inspection of the full lightcurve and individual transit events. The first step on the vetting process was to produce a list of observable fields. This was done by checking the sidereal time of the peak altitude of each field, using the co-ordinates at the centre of the field. In the first run a total of five incomplete fields were observable, three in the first half of the night and two in the second half of the night. The fields ranged from having around 30,000 data points per object through to about 120,000. The fields were first run through a custom implementation of the BLS algorithm, called ORION (Wheatley et al., 2017), which identifies a list of potential candidates from each field. For each BLS power spectrum the top three periods are recorded and only those which have a signal detection efficiency (SDE) of greater than 10 are kept. The list includes many false positives due to eclipsing binaries, signals from the one day aliases, the observing window aliases as well as events which are unlikely to be planetary in nature and other astrophysical false positives. Figure 6.1 shows an example of the interface, called OPIS, which is used to initially vet the candidates.

The CANVAS software, developed by Lionel Métraiiller from the University of Geneva, is run through the list of potential BLS hits and filters out candidates out which have a period which is a multiple of a day, or a period which is common to many other candidates. Both of these criteria can be used to remove candidates that are likely due to systematic noise. The software further filters out candidates with large out of transit variability and strong odd-even transit differences. After running the CANVAS software the pool of candidates is reduced from a list of several hundred candidates to a list of ~ 100 or so per field. For each candidate the host star properties are estimated using an SED fit from publicly available colour information for the host. The transit is fit using the BATMAN package ¹ and can be combined with the host star information to provide an estimate of the planet radius. This estimate can then be used to prioritise smaller planets.

¹<https://www.cfa.harvard.edu/~lkreidberg/batman/>

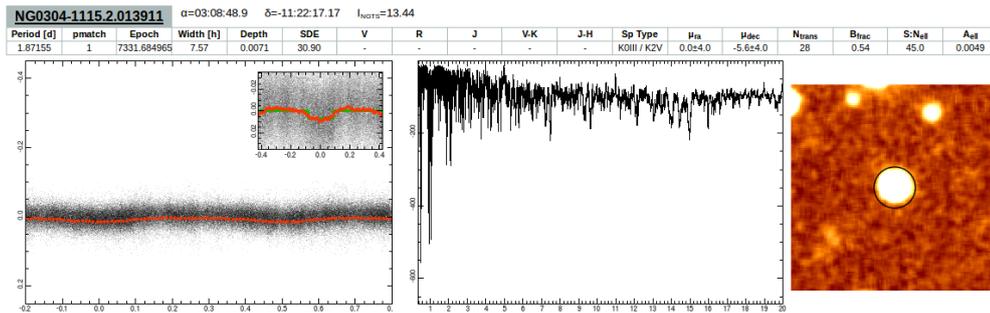


FIGURE 6.1: Example image of the interface used to perform an initial vetting on the candidate. The two main panels are the phase folded lightcurve and the BLS periodogram.

The remaining candidates are then screened by eye by checking the general quality of the lightcurve and transits as well as using the best likelihood planetary parameters derived by the CANVAS software. A postage stamp images of each star is checked to exclude blended targets, which can mimic transit signals due to changes in the point-spread function of the star. Typically after vetting a field we are left with a handful of targets which may plausibly be of planetary nature and so are worth investing follow-up resources for. The top priority candidates are scheduled first and then the remaining time was filled with secondary candidates, which were selected by myself using the same procedure.

As an addition step the centroid vetting algorithm developed by Günther et al. (2017) is used to search for correlations between the flux and the centroid of the target. As detailed by Gunther et al, due to the precision of the NGTS auto-guiding it is possible to check for systematic changes in the centroid position during transit. This is the same method used to vet planet candidates from the Kepler mission (Batalha et al., 2010). Any indication that the transit has a corresponding centroid shift will indicate that the target is an astrophysical false positive. While the centroid shift can not rule out all false positives, especially from the ground, it is a useful statistic for vetting.

For the first run ten targets were selected, which in addition to being good candidates, also covered a range of magnitudes and transit depths, from one to ten milimag. The NGTS survey had not been running for very long and due to some technical issues (rats chewing through cables (Wheatley et al., 2017)) not much data had been taken, so candidates were selected with helping to characterise the performance of the telescope and developing best practices for observing future transits in mind. The candidates were selected for host stars covering a range of magnitudes, which is useful for understanding the precision of the telescope.

For the second run four candidates were selected representing the best possible candidates available at the time, given the known limitations of the telescope.

6.3 Observations

Observations were conducted on the 1-m South African Astronomical Observatory (SAAO) telescope using the SHOC high speed camera's. The telescope was focused at the start of each night by observing a bright star at low airmass. Observations were conducted in a number of filters, typically R or z, chosen on the night to maximise the signal to noise of the transit before the formal start of observations. Exposure times were limited by the nearest bright comparison star, where one was present. Typically I strove to keep the counts below 20,000 ADU as any higher risks saturation due to focus and seeing changes throughout the night. Occasionally it was necessary to re-focus the telescope during observations due to large temperature changes between dusk and the night time. The seeing and FWHM of the stars were carefully monitored through the night, to keep the stars as focused as possible.

Bias frames were taken once at the beginning of the observing run and once at the end. Flat field frames were taken at dusk and sometimes at dawn when the sky was clear. Flats were taken by varying the exposure time and keeping the total counts at roughly 40k per pixel, high enough to maximise signal to noise, while staying below the non-linear regime of the camera. Depending on how crowded the field is it was difficult in some instances to find a suitable guide star. This means that a large exposure time had to be used on the guide camera meaning that stars may drift more than is desirable. Great care was taken to ensure that as many flats as possible were obtained per passband to minimise the flat as a potential source of correlated noise. No dark frames were taken, as the dark current in the SHOC cameras is minimal. The camera is cooled to -50C during operations.

The first observing run was conducted in July from the Southern hemisphere, which meant that the nights were very long and often two candidates could be observed each night (one from the early rising fields and one from the late rising fields). The second run was conducted in November, with much shorted summer nights, and so typically only one transit was observed per night. For each transit the target was observed for a period of two hours before and two hours after expected transit. This value is based on a simple estimate of the

Target	Field	I_{NGTS}	Transit Depth [mmag]	Transit Duration [h]	Period [days]
004320	NG1416-2518	9.12	4.22	1.88	0.452
029667	NG1416-2518	12.11	2.70	1.82	2.220
063395	NG2028-2518	9.35	1.20	1.03	0.375
020085	NG1947-4200	10.88	9.20	2.72	1.061
003528	NG1428-2518	11.92	2.10	4.81	5.239
023276	NG1947-4200	11.30	5.30	2.88	0.923
015761	NG1135-2518	13.35	1.00	3.65	0.943
026905	NG2028-2518	11.74	2.50	2.00	3.166
041317	NG2028-2518	14.63	8.57	6.86	0.871
004858	NG1428-2518	11.08	1.10	2.00	1.357

TABLE 6.1: Transit parameters for candidates selected in the first observing run in April of 2016.

Target	Field	I_{NGTS}	Transit Depth [mmag]	Transit Duration [h]	Period [days]
020500	NG0304-1115	12.09	1.2	3.21	1.778
022940	NG0304-1115	9.37	1.4	3.55	2.789
020279	NG0409-1145	14.65	100	1.34	0.8737
044284	NG0612-2518	13.53	13.9	1.09	0.76689

TABLE 6.2: Transit parameters for candidates selected in the second observing run in November of 2016.

ephemerides error and the number of transits since last observation, which for most targets in the first and second run was in the range of 30-60 minutes. In future an MCMC fit may be performed in each transit to better derive a formal error on the ephemerides which may be better propagated to ensure maximally efficient scheduling. Where time tabling allowed targets were observed for longer, until it was necessary to finish observations for the night or change over to another target for the secondary programme of the observations.

6.4 Reduction and Photometry

The photometric images were processed and reduced using a custom built pipeline name SAFPhot², based on the SEP photometric package³. SEP is a photometric package to perform source extraction and photometry based on the widely used Source Extractor software⁴. SAFPhot is publicly available software jointly developed by the author and Liam Raynard. The purpose of this section is to describe the methods used in SAFPhot to reduce the images and produce high quality lightcurves for NGTS candidate follow up. Other packages were available with which the data could be reduced, however it was felt at the time that having custom software built with the use case of SAAO based NGTS follow-up in mind would ensure that we were getting the best out of the data and may help to overcome any potential quirks of the SHOC camera or the SAAO 1-m telescope. While SAFPhot is publicly available, it is still currently being beta tested by other members of the NGTS consortium in Leicester and will be released formally at a later date.

The SAFPhot pipeline starts by first reducing the images using the standard procedure after producing a master flat field and bias image. For each observing run one master image is produced for each filter in which flats were taken as well as a single bias image for the camera which was used. Figure 6.2 shows the calibration frames used to reduce the data for the first observing run in July. A clear difference can be seen between the flat fields taken in different filters. The flats taken in bluer filters show more structure than those taken in a red filter, a clear example of the blue diamond effect. The flat taken in white light shows a mixture of the features from the red and blue flats. Dust can also clearly be seen in the flats, which needs to be carefully flat fielded out. The strength of

²https://github.com/apchsh/saa0_scripts

³<https://github.com/kbarbary/sep>

⁴<https://www.astromatic.net/software/sextractor>

Target	Field	I_{NGTS}	Filter	Exposure Time [s]	Number of Exposures [s]
004320	NG1416-2518	9.12	z'	2.0	5730
029667	NG1416-2518	12.11	z'	10.0	1560
063395	NG2028-2518	9.35	z'	2.0	9450
020085	NG1947-4200	10.88	R	2.0	11250
003528	NG1428-2518	11.92	R	12.0	1500
023276	NG1947-4200	11.30	R	30.0	630
015761	NG1135-2518	13.35	R	5.0	2040
026905	NG2028-2518	11.74	R	2.0	9600
041317	NG2028-2518	14.63	R	5.0	3842
004858	NG1428-2518	11.08	R	2.0	9122

TABLE 6.3: Observational parameters for candidates selected in the first observing run in April of 2016.

Target	Field	I_{NGTS}	Filter	Exposure Time [s]	Number of Exposures [s]
020500	NG0304-1115	12.09	z'	10.0	1520
022940	NG0304-1115	9.37	z'	2.0	5400
020279	NG0409-1145	14.65	z'	10.0	1230
044284	NG0612-2518	13.53	z'	30.0	630

TABLE 6.4: Observational parameters for candidates selected in the second observing run in November of 2016.

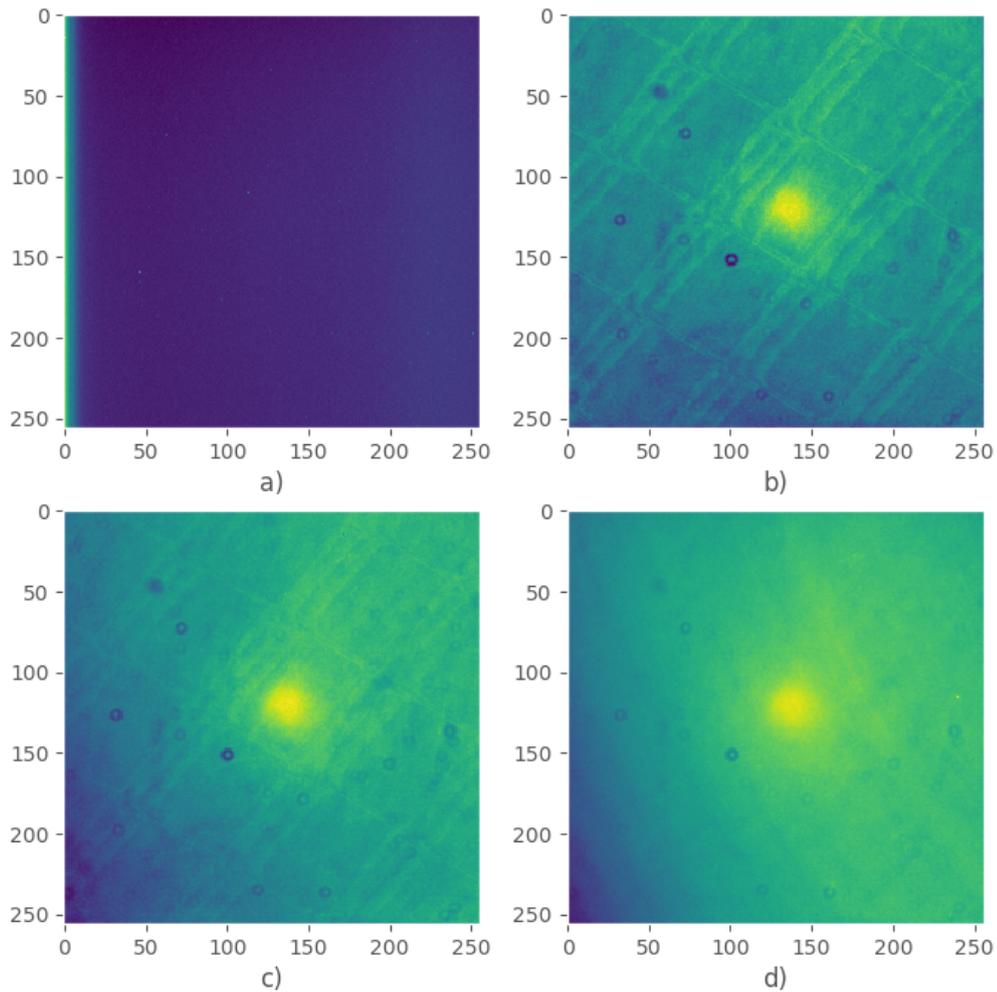


FIGURE 6.2: A plot showing the master calibration frames for the first observing run. From top left to bottom right: a) master bias frame, b) master flat in V, c) master flat in white light and finally d) master flat in R.

the dust features is also filter dependent and somewhat worse in blue light. The bias frame is relatively uniform across, with a small gradient near the left hand edge of the frame. Features like this are typical. Several white spots can be seen which are hot pixels and should be masked out of the images before photometry.

After the reduction, the images are saved individually and then processed by the SAFPhot photometry script. Each image is background subtracted using the SEP background function which produces a background map by splitting the image in box sizes of 32 by 32 pixels and convolving the pixels with a 3x3 pixel smoothing filter to remove any stars. The median value of each box is then calculated and used to create the background map by bilinear interpolation with the other image boxes. For a high-spatially varying background the box sizes should be made smaller to more accurately track changes in the background.

However, in this case, SHOC has such a small field of view that this is typically not an issue and so the box sizes are kept relatively big with respect to the image size (256 by 256 pixels) to minimise shot noise and systematic noise from the stars in the background estimate. (Lower box sizes can lead to an over estimate of the background levels if stars are not properly removed.)

After the background is removed source detection was then run on each image, to find the target star and all the comparisons. The first image from each series of observations was used as the reference image for the source catalogue. Each star is numbered according to its position. An example can be seen in Figure 6.3. Once the stars are found a windowed centroiding algorithm is run on the image to further refine the stellar positions. For each new image, stars are matched to the original reference image by finding the nearest star to these catalogue reference. In some cases the telescope can move enough so that a star shifts off-frame or a new one appears. So long as it is not more than a few tens of pixels the SAF-Phot pipeline can deal with changes in position. Aperture photometry is then run on each frame using aperture sizes of 3px through to 12px in increments of 0.1px. This allows us to choose the optimal aperture size once the photometry is complete and is only marginally slower than running the photometry for a single aperture. Typically a single run can be reduced and processed automatically in less than half an hour, where a run contains up to ten thousand images or so.

After aperture photometry is complete for each field we perform differential photometry. Comparison stars are selected by hand from those available in the image and mean combined and normalised to produce an ensemble reference star. The target lightcurve is normalised and divided by the reference star photometry to leave a normalised differential lightcurve. Normalised lightcurves are used as we are looking for transits and so are interested only in relative changes. This process is performed for every aperture radius choice and the signal to noise is calculated in each case, by dividing the mean of the lightcurve by its standard deviation. The final aperture is chosen to be the one which has the best signal to noise. Typically it is a trade-off between including as many photons as possible on the bright stars (which have optimal apertures in excess of 10px) and keeping the apertures on the faint stars small so as to not introduce too much error from the background estimation. By performing these processes end to end before photometry we ensure that we get the best target and comparison lightcurve as is possible. Aperture sizes are kept the same for all of the stars in the image, throughout the night, to prevent the introduction of systematic noise.

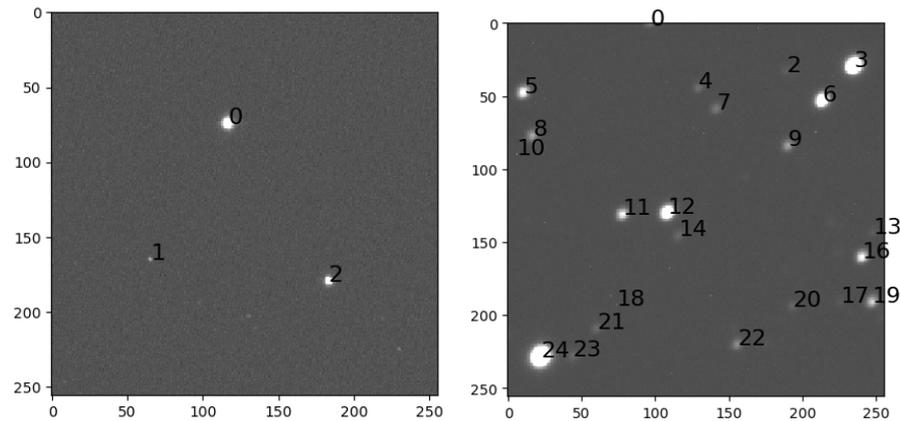


FIGURE 6.3: Two example images from SAAO data shown after source detection has been performed. The stars are annotated with numbers which are based on a cross-match with the stars in the reference image.

Finally we compare the differential lightcurve of the target star using the mean reference star and every reference star individually to check that none of the reference stars are variable. Each variation of the target lightcurve is then inspected for quality. The standard deviation is also computed as an additional check. If a star is suspected to be variable it is removed from the mean comparison star to mitigate against introducing spurious photometric signals which may mask a real transit or introduce a false positive. The resulting target lightcurves are then inspected by hand for quality or evidence of any systematic issues due to the pipeline. No major problems were present in lightcurves from observing runs one and two.

6.5 Results

6.5.1 Transit Confirmations

The SAFPhot pipeline was successfully run over all the data. All candidates were inspected by eye for evidence of a transit. Of these only NG0612-2518.44284 showed clear evidence of a deep (1.8%) transit, while the other targets showed no evidence of a transit. This likely indicates that these candidates were the result of instrumental systematics in the NGTS lightcurves and as such can be safely excluded from secondary follow-up. Figure 6.4 shows the lightcurves of

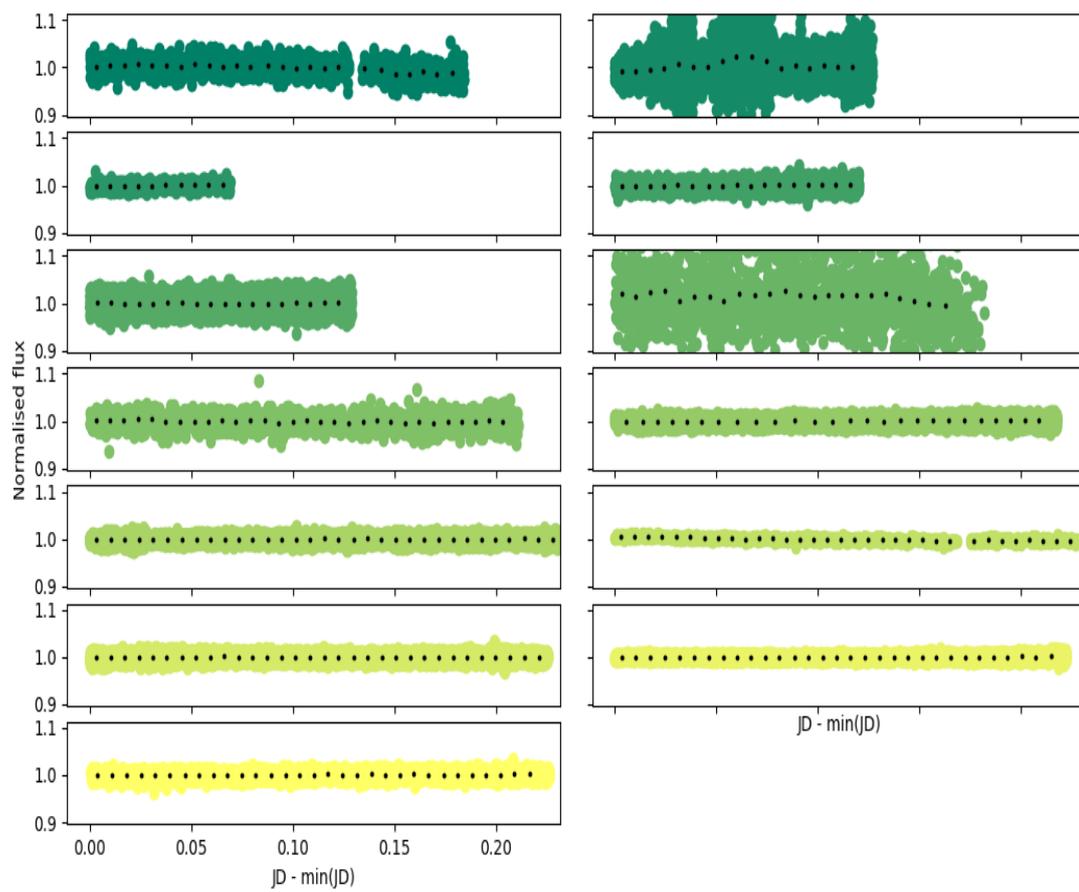


FIGURE 6.4: Lightcurves of the thirteen targets observed in the first run which show no sign of a transit. Black dots represent 10-minute cadence.

the targets, with a 10 minute cadence. While no transits were observed, the quality of the lightcurves suggest that in future follow-up taken with the SAAO 1-m using SHOC will yield valuable observations.

6.5.2 NG0612-2518.44284

The transit of NG0612-2518.44284 was observed in the R-band. The target was located next to a nearby bright object, and as such it was suspected that the real transit was much deeper and had been made shallower in NGTS photometry by blending. The target was observed with NGTS to confirm the final transit depth and check whether it matches the NGTS photometry or not. As can clearly be seen in Figure 6.5 the transit depths match in both SAAO and NGTS photometry in both pass-bands. While it is unlikely that the transit it blended out to exactly the same depth in both datasets, we nevertheless run an additional test to verify that the transit depth is stable and not a function of the aperture radius (which would indicate blending). The photometry is repeated for apertures with radii 3.0px through to 13px in increments of 0.1px. For each aperture radius we compute the depth of the transit in the normalised lightcurve. The results can be seen in Figure 6.6. While there is some dependence on the transit depth, perhaps indicating diffuse light from the nearby companion, it is not significant and in fact the transit depth stabilises and then decreases around a radius of 5px. Secondly the target is more than 30 pixels away from the bright companion significantly limiting the chances of contamination.

NG0612-2518.44284 is an intriguing candidate. The similar transit depth in both z' and R may indicate that the transit is of a planetary nature. A second scenario is that the target is a binary, where both companions are of a similar spectral type, accounting for the consistent transit depth and lack of secondary transit in the lightcurve.

6.5.3 Performance of the SAAO 1-m

Figure 6.7 shows a plot of the standard deviation, after binning to ten minutes, of each lightcurve observed in run one which does not show evidence of a transiting planet. Three of the targets show excessively high noise, the first due to a lack of bright comparison stars with which to perform differential photometry and the second and third due to high cloud which was present on the night. Overall the lightcurves obtained are of really high quality, having a noise level

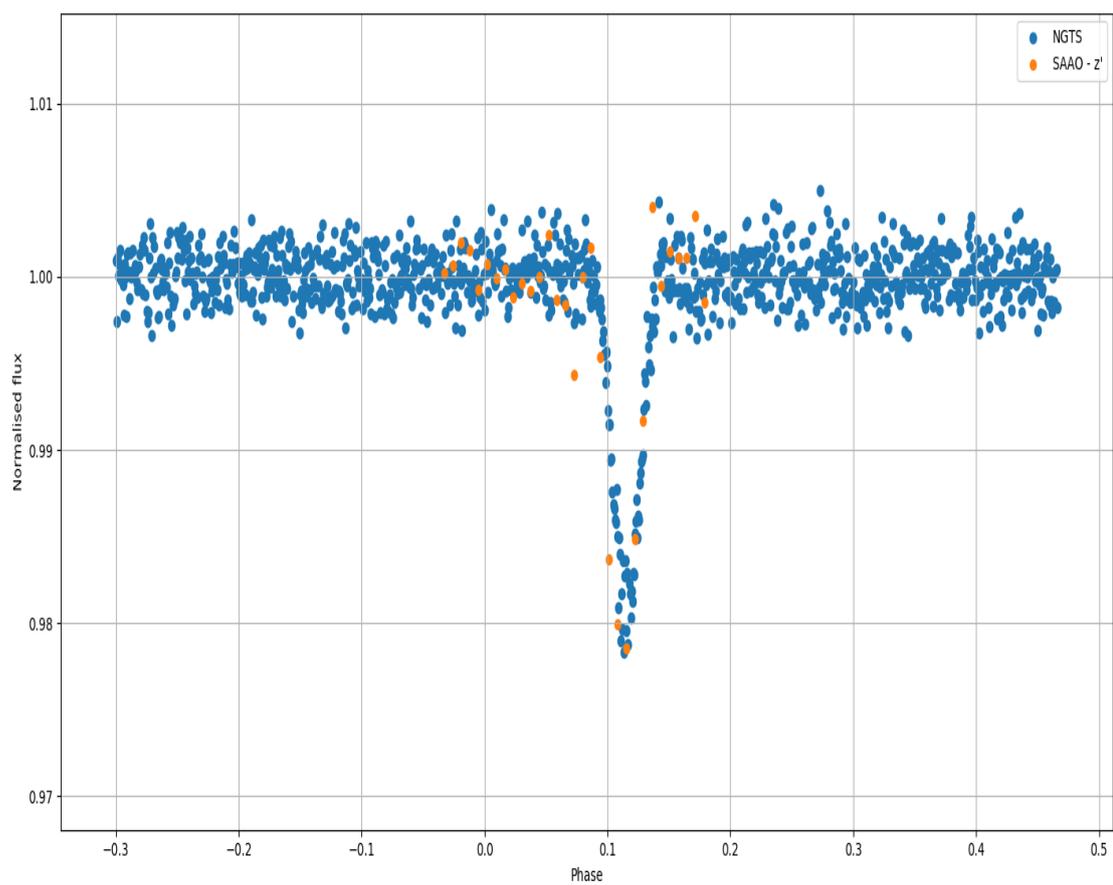


FIGURE 6.5: Comparison between SAAO photometry (orange) and NGTS photometry for NG44284. SAAO photometry is binned to 10 minuted cadence.

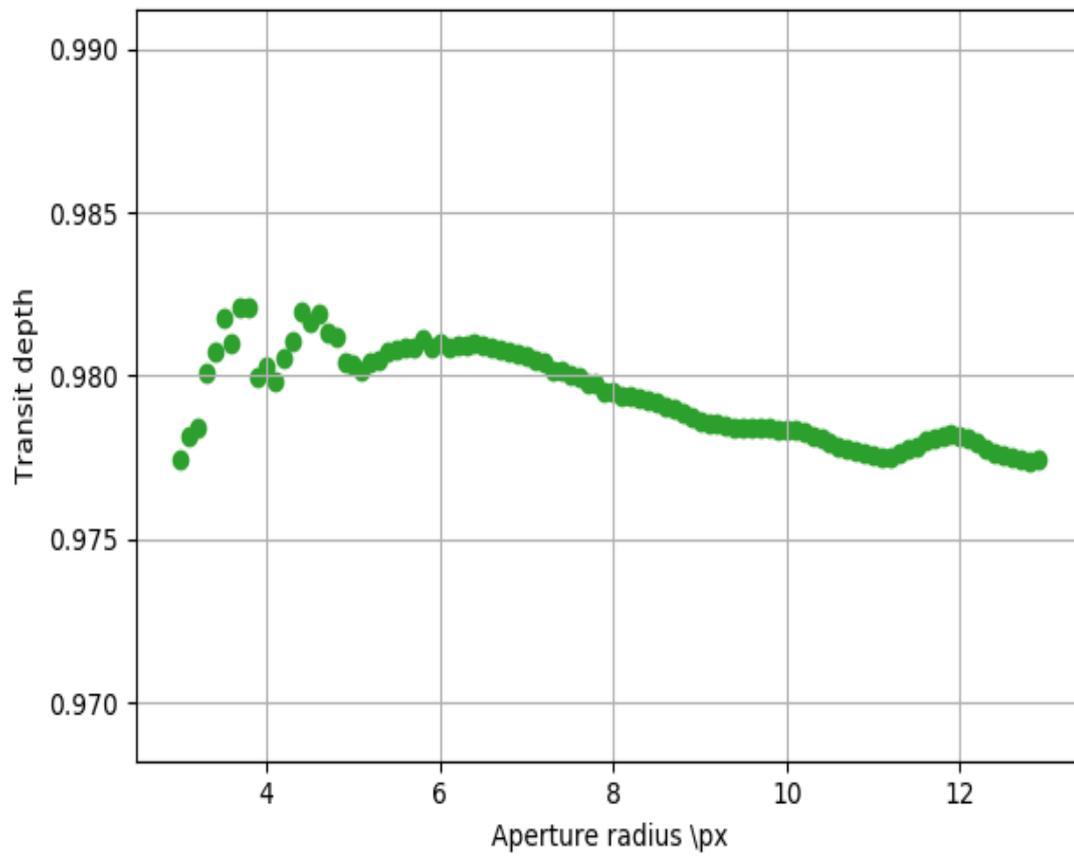


FIGURE 6.6: Comparison of the transit depth of NG44284 in the SAAO data using different aperture radii.

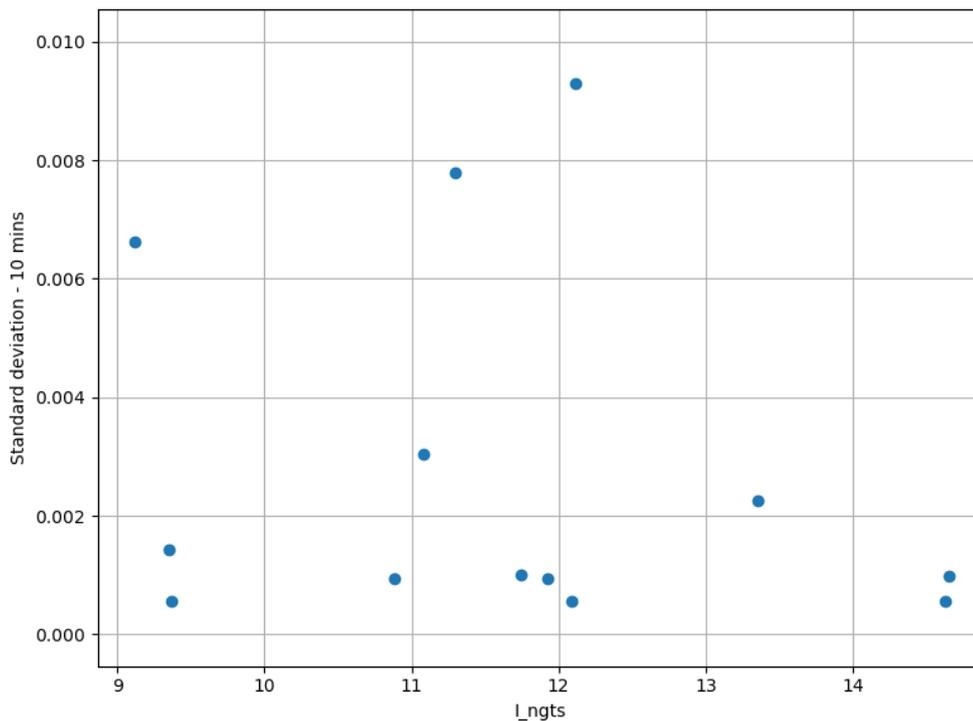


FIGURE 6.7: Plot of lightcurve standard deviation versus NGTS magnitude. Only targets which showed no transit were used.

of less than 2mmag in 10 minutes, with some as low as 0.5mmag. This is an impressive result to achieve with a single night’s worth of data and further demonstrates the potential of a 1-m telescope to perform high quality follow-up work and transit vetting for the NGTS survey. Even potential candidates in the 1 mmag regime can be vetted from South Africa, in some cases in a single night. Furthermore the quality of the final 10 minute binned photometry does not show a strong dependence on the magnitude of the NGTS target in the NGTS pass band. The quality of the binned lightcurves is fairly similar despite the initial starting point (see Figure 6.4 for the lightcurves) and so follow-up observations can be conducted for the full range of NGTS targets in our parameter space of I_{mag} of 8 to 13 (Wheatley et al., 2017).

6.6 Conclusion

Over the course of three weeks of total observing time, 14 potential transiting candidates from the NGTS survey were observed using the SHOC camera on the 1-m SAAO telescope. The objects were observed for five to eight hours each,

depending on the transit duration and the observation constraints for that particular night. The data was reduced using the SAFPhot pipeline, a custom built pipeline designed to produce the best possible lightcurves for the vetting of shallow transiting candidates. Thirteen candidates can be eliminated for secondary follow-up with one candidate remaining. NG0612-2518.44284 is likely a binary star or of planetary nature, however further follow-up is needed. The quality of the lightcurves indicates that the 1m SAAO telescope in combination with SHOC will prove to be a very useful instrument for future follow-up work.

Chapter 7

Conclusion

The Next Generation Transit Survey is poised to discover super-Earth and Neptune sized worlds around bright host stars in the over the next couple of year. The first NGTS planet discovery (Bayliss et al., 2017) has already been published, an exciting discovery of a system with the largest planet to star ratio currently known, and many more planets are in the pipeline. In additions, the survey has produced hundreds of thousands of lightcurves of exceptional quality opening up the opportunity for additional, serendipitous science.

Discussion

Finding the smallest planets requires extremely precise photometry, a key part of which is the image reduction process. The NGTS reduction pipeline is discussed at length, detailing the steps necessary to produce the high signal to noise calibration frames required for the survey. The on-sky calibration masters are compared to flat fields made in the laboratory and to a night flat produced from stacking together science frames, verifying the high quality of the calibration frames. There frames are then used in the main production pipeline to produce the best precision lightcurves possible.

Another key limiting factor to finding small planets is the red or correlated noise present in the survey. The NGTS data set is analysed to measure the correlated noise components of the lightcurves. The correlated noise is then studied in comparison with other properties of the data to determine if any single component may be causing a large contribution. Finally, the background subtraction is investigate to see if the correlated noise could be minimised by tuning the parameters. The results of this optimisation are inconclusive and may point to an instrumentation origin for the correlated noise.

With the high precision NGTS lightcurves there are opportunities to search for single transit events and to study the variable sky. To this end Bayesian change point analysis is applied to the dataset to search for these types of events. The algorithm produces two strong candidate transits, one from a single transit event and another a regular transiting planet with a 1-day period. The change point analysis also produces some variable and transient events, though with relatively high false positives. The method is successful as a proof-of-concept for single transit search and for identifying variables.

Finally detailed photometric follow-up of NGTS candidates is performed using the 1m telescope of the South African Astronomical Observatory. A custom pipeline is implemented to maximise the signal to noise of the lightcurves. One interesting candidate is discussed in detail.

Future Work

Much work remains to be done to fully understand the correlated noise in the data. A key question is what is the part which the SysRem algorithm plays in the detrending process. Does it do so optimally? Are there any limitations to how well SysRem performs on real world data. Future work to understand whether the detrending process can be improved will be crucial to maximising the yield of planets from NGTS.

Even though the correlated noise may be greatly reduced, there is a chance that some will remain. As such it is key to understand how best to discover transiting exoplanets in the low signal to noise regime. Machine learning techniques are playing an increasingly important role in astronomy, and the area of exoplanet detection is no exception. Neural networks should be explored as an option to detect transits in data with correlated noise, due to the relative flexibility of such models and their ability to learn features from the data.

Final Outlook

The goal of this thesis was to maximise the yield of planets from NGTS by thoroughly understanding the noise properties of the data and by finding planets which would otherwise have been missed. Change point analysis is already producing candidates for study, and once applied to the full dataset will likely yield many more. The reduction pipeline has helped to minimise the noise in the

lightcurves by producing high precision calibration master frames, while the exploration of the correlated noise has helped to constrain which components may be adding to the problem. Taken together, this study of the NGTS data and noise will hopefully contribute to the wider goals of the exoplanet community in some small way.

Bibliography

- Alard, C. (2000). "Image subtraction using a space-varying kernel". In: *AAPS* 144, pp. 363–370. DOI: [10.1051/aas:2000214](https://doi.org/10.1051/aas:2000214).
- Alard, C. and R. H. Lupton (1998). "A Method for Optimal Image Subtraction". In: *APJ* 503, pp. 325–331. DOI: [10.1086/305984](https://doi.org/10.1086/305984). eprint: [astro-ph/9712287](https://arxiv.org/abs/astro-ph/9712287).
- Armstrong, D. J. et al. (2016). "K2 variable catalogue - II. Machine learning classification of variable stars and eclipsing binaries in K2 fields 0-4". In: *MNRAS* 456, pp. 2260–2272. DOI: [10.1093/mnras/stv2836](https://doi.org/10.1093/mnras/stv2836). arXiv: [1512.01246](https://arxiv.org/abs/1512.01246) [[astro-ph](https://arxiv.org/abs/astro-ph).SR].
- Bakos, G. et al. (2004). "Wide-Field Millimagnitude Photometry with the HAT: A Tool for Extrasolar Planet Detection". In: *PASP* 116, pp. 266–277. DOI: [10.1086/382735](https://doi.org/10.1086/382735). eprint: [astro-ph/0401219](https://arxiv.org/abs/astro-ph/0401219).
- Bakos, G. Á. et al. (2002). "System Description and First Light Curves of the Hungarian Automated Telescope, an Autonomous Observatory for Variability Search". In: *PASP* 114, pp. 974–987. DOI: [10.1086/342382](https://doi.org/10.1086/342382). eprint: [astro-ph/0206001](https://arxiv.org/abs/astro-ph/0206001).
- Barry, D. and J. A. Hartigan (1992). "Product Partition Models for Change Point Problems". In: *Ann. Statist.* 820.
- (1993). "A Bayesian Analysis for Change Point Problems". In: *JASA* 88, pp. 309–319.
- Batalha, N. M. et al. (2010). "Pre-spectroscopic False-positive Elimination of Kepler Planet Candidates". In: *ApJ* 713, pp. L103–L108. DOI: [10.1088/2041-8205/713/2/L103](https://doi.org/10.1088/2041-8205/713/2/L103). arXiv: [1001.0392](https://arxiv.org/abs/1001.0392) [[astro-ph](https://arxiv.org/abs/astro-ph).EP].
- Bayliss, D. et al. (2017). "NGTS-1b: A hot Jupiter transiting an M-dwarf". In: *ArXiv e-prints*. arXiv: [1710.11099](https://arxiv.org/abs/1710.11099) [[astro-ph](https://arxiv.org/abs/astro-ph).EP].
- Bertin, E. and S. Arnouts (1996). "SExtractor: Software for source extraction." In: *A&AS* 117, pp. 393–404. DOI: [10.1051/aas:1996164](https://doi.org/10.1051/aas:1996164).
- Bramich, D. M. (2008). "A new algorithm for difference image analysis". In: *MNRAS* 386, pp. L77–L81. DOI: [10.1111/j.1745-3933.2008.00464.x](https://doi.org/10.1111/j.1745-3933.2008.00464.x). arXiv: [0802.1273](https://arxiv.org/abs/0802.1273).

- Brink, H. et al. (2013). “Using machine learning for discovery in synoptic survey imaging data”. In: *MNRAS* 435, pp. 1047–1060. DOI: [10.1093/mnras/stt1306](https://doi.org/10.1093/mnras/stt1306). arXiv: [1209.3775](https://arxiv.org/abs/1209.3775) [astro-ph.IM].
- Brown, D. J. A. and D. J. Armstrong (2015). “Constraints on circumbinary planet orbits from Kepler single transit events”. In: *European Planetary Science Congress 10*, EPSC2015-588, EPSC2015–588.
- Crouzet, N. et al. (2010). “ASTEP South: an Antarctic Search for Transiting Exoplanets around the celestial south pole”. In: *A&A* 511, A36, A36. DOI: [10.1051/0004-6361/200913629](https://doi.org/10.1051/0004-6361/200913629). arXiv: [0912.2644](https://arxiv.org/abs/0912.2644) [astro-ph.IM].
- Dittmann, J. et al. (2017). “Identifying Long-period Planets from Single Transit Events with the MEarth Project”. In: *American Astronomical Society Meeting Abstracts*. Vol. 229. American Astronomical Society Meeting Abstracts, p. 415.01.
- Foreman-Mackey, D., D. W. Hogg, and B. Schölkopf (2015). “The search for single exoplanet transits in the Kepler light curves”. In: *IAU General Assembly 22*, 2258352, p. 2258352.
- Gillon, M. et al. (2012). “The TRAPPIST survey of southern transiting planets. I. Thirty eclipses of the ultra-short period planet WASP-43 b”. In: *A&A* 542, A4, A4. DOI: [10.1051/0004-6361/201218817](https://doi.org/10.1051/0004-6361/201218817). arXiv: [1201.2789](https://arxiv.org/abs/1201.2789) [astro-ph.EP].
- Gillon, M. et al. (2017). “Seven temperate terrestrial planets around the nearby ultracool dwarf star TRAPPIST-1”. In: *Nature* 542, pp. 456–460. DOI: [10.1038/nature21360](https://doi.org/10.1038/nature21360). arXiv: [1703.01424](https://arxiv.org/abs/1703.01424) [astro-ph.EP].
- Günther, M. N. et al. (2017). “Centroid vetting of transiting planet candidates from the Next Generation Transit Survey”. In: *MNRAS* 472, pp. 295–307. DOI: [10.1093/mnras/stx1920](https://doi.org/10.1093/mnras/stx1920). arXiv: [1707.07978](https://arxiv.org/abs/1707.07978) [astro-ph.EP].
- Howard, A. W. (2013). “Observed Properties of Extrasolar Planets”. In: *Science* 340, pp. 572–576. DOI: [10.1126/science.1233545](https://doi.org/10.1126/science.1233545). arXiv: [1305.0542](https://arxiv.org/abs/1305.0542) [astro-ph.EP].
- Irwin, J. et al. (2007). “The Monitor project: data processing and light curve production”. In: *MNRAS* 375, pp. 1449–1462. DOI: [10.1111/j.1365-2966.2006.11408.x](https://doi.org/10.1111/j.1365-2966.2006.11408.x). eprint: [astro-ph/0612395](https://arxiv.org/abs/astro-ph/0612395).
- Irwin, J. M. et al. (2015). “The MEarth-North and MEarth-South Transit Surveys: Searching for Habitable Super-Earth Exoplanets Around Nearby M-dwarfs”. In: *18th Cambridge Workshop on Cool Stars, Stellar Systems, and the Sun*. Ed. by G. T. van Belle and H. C. Harris. Vol. 18. Cambridge Workshop on Cool Stars, Stellar Systems, and the Sun, pp. 767–772. arXiv: [1409.0891](https://arxiv.org/abs/1409.0891) [astro-ph.EP].
- Irwin, M. J. (1997). “Detectors and data analysis techniques for wide field optical imaging.” In: *Instrumentation for Large Telescopes*. Ed. by J. M. Rodríguez Espinosa, A. Herrero, and F. Sánchez, pp. 35–74.

- Kenworthy, M. A. and E. E. Mamajek (2015). “Modeling Giant Extrasolar Ring Systems in Eclipse and the Case of J1407b: Sculpting by Exomoons?” In: *ApJ* 800, 126, p. 126. DOI: [10.1088/0004-637X/800/2/126](https://doi.org/10.1088/0004-637X/800/2/126). arXiv: [1501.05652](https://arxiv.org/abs/1501.05652) [astro-ph.SR].
- Killick, R., P. Fearnhead, and I.A. Eckley (2012). “Optimal detection of change-points with a linear computation cost”. In: *JASA* 107, pp. 1590–1598.
- Kovács, G., G. Bakos, and R. W. Noyes (2005). “A trend filtering algorithm for wide-field variability surveys”. In: *MNRAS* 356, pp. 557–567. DOI: [10.1111/j.1365-2966.2004.08479.x](https://doi.org/10.1111/j.1365-2966.2004.08479.x). eprint: [astro-ph/0411724](https://arxiv.org/abs/astro-ph/0411724).
- Kovács, G., S. Zucker, and T. Mazeh (2002). “A box-fitting algorithm in the search for periodic transits”. In: *A&A* 391, pp. 369–377. DOI: [10.1051/0004-6361:20020802](https://doi.org/10.1051/0004-6361:20020802). eprint: [astro-ph/0206099](https://arxiv.org/abs/astro-ph/0206099).
- Lopez, E. D. (2017). “Born dry in the photoevaporation desert: Kepler’s ultra-short-period planets formed water-poor”. In: *MNRAS* 472, pp. 245–253. DOI: [10.1093/mnras/stx1558](https://doi.org/10.1093/mnras/stx1558). arXiv: [1610.01170](https://arxiv.org/abs/1610.01170) [astro-ph.EP].
- Mayor, M. and D. Queloz (1995). “A Jupiter-mass companion to a solar-type star”. In: *Nature* 378, pp. 355–359. DOI: [10.1038/378355a0](https://doi.org/10.1038/378355a0).
- McCormac, J. et al. (2013). “DONUTS: A Science Frame Autoguiding Algorithm with Sub-Pixel Precision, Capable of Guiding on Defocused Stars”. In: *PASP* 125, p. 548. DOI: [10.1086/670940](https://doi.org/10.1086/670940). arXiv: [1304.2405](https://arxiv.org/abs/1304.2405) [astro-ph.IM].
- Mékarnia, D. et al. (2016). “Transiting planet candidates with ASTEP 400 at Dome C, Antarctica”. In: *MNRAS* 463, pp. 45–62. DOI: [10.1093/mnras/stw1934](https://doi.org/10.1093/mnras/stw1934).
- Nutzman, P. and D. Charbonneau (2008). “Design Considerations for a Ground-Based Transit Search for Habitable Planets Orbiting M Dwarfs”. In: *PASP* 120, p. 317. DOI: [10.1086/533420](https://doi.org/10.1086/533420). arXiv: [0709.2879](https://arxiv.org/abs/0709.2879).
- Ofir, A. et al. (2010). “The SARS algorithm: detrending CoRoT light curves with Sysrem using simultaneous external parameters”. In: *MNRAS* 404, pp. L99–L103. DOI: [10.1111/j.1745-3933.2010.00843.x](https://doi.org/10.1111/j.1745-3933.2010.00843.x). arXiv: [1003.0427](https://arxiv.org/abs/1003.0427) [astro-ph.EP].
- Osborn, H. P. et al. (2016). “Single transit candidates from K2: detection and period estimation”. In: *MNRAS* 457, pp. 2273–2286. DOI: [10.1093/mnras/stw137](https://doi.org/10.1093/mnras/stw137). arXiv: [1512.03722](https://arxiv.org/abs/1512.03722) [astro-ph.EP].
- PAGE, E. S. (1957). “On problems in which a change in a parameter occurs at an unknown point”. In: *Biometrika* 44.1-2, pp. 248–252. DOI: [10.1093/biomet/44.1-2.248](https://doi.org/10.1093/biomet/44.1-2.248). eprint: [/oup/backfile/content_public/journal/biomet/44/1-2/248](https://www.jstor.org/stable/2342248). URL: [+http://dx.doi.org/10.1093/biomet/44.1-2.248](http://dx.doi.org/10.1093/biomet/44.1-2.248).

- Parviainen, H. (2015). “PYTRANSIT: fast and easy exoplanet transit modelling in PYTHON”. In: *MNRAS* 450, pp. 3233–3238. DOI: [10.1093/mnras/stv894](https://doi.org/10.1093/mnras/stv894). arXiv: [1504.07433](https://arxiv.org/abs/1504.07433) [astro-ph.EP].
- Pepper, J. et al. (2017). “KELT-11b: A Highly Inflated Sub-Saturn Exoplanet Transiting the $V = 8$ Subgiant HD 93396”. In: *AJ* 153, 215, p. 215. DOI: [10.3847/1538-3881/aa6572](https://doi.org/10.3847/1538-3881/aa6572). arXiv: [1607.01755](https://arxiv.org/abs/1607.01755) [astro-ph.EP].
- Pollacco, D. L. et al. (2006). “The WASP Project and the SuperWASP Cameras”. In: *PASP* 118, pp. 1407–1418. DOI: [10.1086/508556](https://doi.org/10.1086/508556). eprint: [astro-ph/0608454](https://arxiv.org/abs/astro-ph/0608454).
- Pont, F., S. Zucker, and D. Queloz (2006). “The effect of red noise on planetary transit detection”. In: *MNRAS* 373, pp. 231–242. DOI: [10.1111/j.1365-2966.2006.11012.x](https://doi.org/10.1111/j.1365-2966.2006.11012.x). eprint: [astro-ph/0608597](https://arxiv.org/abs/astro-ph/0608597).
- Prescott Adams, R. and D. J. C. MacKay (2007). “Bayesian Online Change-point Detection”. In: *ArXiv e-prints*. arXiv: [0710.3742](https://arxiv.org/abs/0710.3742) [stat.ML].
- Richards, J. W. et al. (2011). “On Machine-learned Classification of Variable Stars with Sparse and Noisy Time-series Data”. In: *ApJ* 733, 10, p. 10. DOI: [10.1088/0004-637X/733/1/10](https://doi.org/10.1088/0004-637X/733/1/10). arXiv: [1101.1959](https://arxiv.org/abs/1101.1959) [astro-ph.IM].
- Sanchis-Ojeda, R. et al. (2014). “A Study of the Shortest-period Planets Found with Kepler”. In: *ApJ* 787, 47, p. 47. DOI: [10.1088/0004-637X/787/1/47](https://doi.org/10.1088/0004-637X/787/1/47). arXiv: [1403.2379](https://arxiv.org/abs/1403.2379) [astro-ph.EP].
- Scott, A.J and M. Knott (1974). “A Cluster Analysis Method for Grouping Means in the Analysis of Variance”. In: *Biometrics* 30, pp. 507–512.
- Smith, A. M. S. et al. (2006). “The impact of correlated noise on SuperWASP detection rates for transiting extrasolar planets”. In: *MNRAS* 373, pp. 1151–1158. DOI: [10.1111/j.1365-2966.2006.11095.x](https://doi.org/10.1111/j.1365-2966.2006.11095.x). eprint: [astro-ph/0609618](https://arxiv.org/abs/astro-ph/0609618).
- Southworth, J. et al. (2017). “Detection of the Atmosphere of the 1.6M Exoplanet GJ 1132 b”. In: *ApJ* 153.4, p. 191. URL: <http://stacks.iop.org/1538-3881/153/i=4/a=191>.
- Surma, P. (1993). “Shutter-free flatfielding for CCD detectors”. In: *Astronomy and Astrophysics* 278, pp. 654–658.
- Tamuz, O., T. Mazeh, and S. Zucker (2005). “Correcting systematic effects in a large set of photometric light curves”. In: *mnras* 356, pp. 1466–1470. DOI: [10.1111/j.1365-2966.2004.08585.x](https://doi.org/10.1111/j.1365-2966.2004.08585.x). eprint: [astro-ph/0502056](https://arxiv.org/abs/astro-ph/0502056).
- Tsiaras, A. et al. (2016). “Detection of an Atmosphere Around the Super-Earth 55 Cancri e”. In: *ApJ* 820, 99, p. 99. DOI: [10.3847/0004-637X/820/2/99](https://doi.org/10.3847/0004-637X/820/2/99). arXiv: [1511.08901](https://arxiv.org/abs/1511.08901) [astro-ph.EP].
- Udalski, A. (2003). “The Optical Gravitational Lensing Experiment. Real Time Data Analysis Systems in the OGLE-III Survey”. In: *Acta Astron.* 53, pp. 291–305. eprint: [astro-ph/0401123](https://arxiv.org/abs/astro-ph/0401123).

- Wakeford, H. R. et al. (2017). “HAT-P-26b: A Neptune-mass exoplanet with a well-constrained heavy element abundance”. In: *Science* 356, pp. 628–631. DOI: [10.1126/science.aah4668](https://doi.org/10.1126/science.aah4668). arXiv: [1705.04354](https://arxiv.org/abs/1705.04354) [astro-ph.EP].
- Welch, D. L. and P. B. Stetson (1993). “Robust variable star detection techniques suitable for automated searches - New results for NGC 1866”. In: *AJ* 105, pp. 1813–1821. DOI: [10.1086/116556](https://doi.org/10.1086/116556).
- Wheatley, P. J. et al. (2013). “The Next Generation Transit Survey (NGTS)”. In: *European Physical Journal Web of Conferences*. Vol. 47. European Physical Journal Web of Conferences, p. 13002. DOI: [10.1051/epjconf/20134713002](https://doi.org/10.1051/epjconf/20134713002). arXiv: [1302.6592](https://arxiv.org/abs/1302.6592) [astro-ph.EP].
- Wheatley, P. J. et al. (2017). “The Next Generation Transit Survey (NGTS)”. In: *ArXiv e-prints*. arXiv: [1710.11100](https://arxiv.org/abs/1710.11100) [astro-ph.EP].
- Zucker, S. and R. Giryes (2017). “Shallow Transits - Deep Learning I: Feasibility Study of Deep Learning to Detect Periodic Transits of Exoplanets”. In: *ArXiv e-prints*. arXiv: [1711.03163](https://arxiv.org/abs/1711.03163) [astro-ph.IM].