

University of Warwick institutional repository: <http://go.warwick.ac.uk/wrap>

A Thesis Submitted for the Degree of PhD at the University of Warwick

<http://go.warwick.ac.uk/wrap/60463>

This thesis is made available online and is protected by original copyright.

Please scroll down to view the document itself.

Please refer to the repository record for this item for information to help you to cite it. Our policy information is available from the repository home page.

Library Declaration and Deposit Agreement

1. STUDENT DETAILS

Please complete the following:

Full name: SIMON WALKER

University ID number: 0501572

2. THESIS DEPOSIT

2.1 I understand that under my registration at the University, I am required to deposit my thesis with the University in BOTH hard copy and in digital format. The digital version should normally be saved as a single pdf file.

2.2 The hard copy will be housed in the University Library. The digital version will be deposited in the University's Institutional Repository (WRAP). Unless otherwise indicated (see 2.3 below) this will be made openly accessible on the Internet and will be supplied to the British Library to be made available online via its Electronic Theses Online Service (EThOS) service.

[At present, theses submitted for a Master's degree by Research (MA, MSc, LLM, MS or MMedSci) are not being deposited in WRAP and not being made available via EThOS. This may change in future.]

2.3 In exceptional circumstances, the Chair of the Board of Graduate Studies may grant permission for an embargo to be placed on public access to the hard copy thesis for a limited period. It is also possible to apply separately for an embargo on the digital version. (Further information is available in the *Guide to Examinations for Higher Degrees by Research*.)

2.4 If you are depositing a thesis for a Master's degree by Research, please complete section (a) below. For all other research degrees, please complete both sections (a) and (b) below:

(a) Hard Copy

I hereby deposit a hard copy of my thesis in the University Library to be made publicly available to readers (please delete as appropriate) EITHER immediately ~~OR after an embargo period of~~
~~months/years as agreed by the Chair of the Board of Graduate Studies.~~

I agree that my thesis may be photocopied. YES / ~~NO~~ (Please delete as appropriate)

(b) Digital Copy

I hereby deposit a digital copy of my thesis to be held in WRAP and made available via EThOS.

Please choose one of the following options:

EITHER My thesis can be made publicly available online. YES / ~~NO~~ (Please delete as appropriate)

OR My thesis can be made publicly available only after....[date] (Please give date)
~~YES~~ / NO (Please delete as appropriate)

OR My full thesis cannot be made publicly available online but I am submitting a separately identified additional, abridged version that can be made available online.
~~YES~~ / NO (Please delete as appropriate)

OR My thesis cannot be made publicly available online. ~~YES~~ / NO (Please delete as appropriate)

3. GRANTING OF NON-EXCLUSIVE RIGHTS

Whether I deposit my Work personally or through an assistant or other agent, I agree to the following:

Rights granted to the University of Warwick and the British Library and the user of the thesis through this agreement are non-exclusive. I retain all rights in the thesis in its present version or future versions. I agree that the institutional repository administrators and the British Library or their agents may, without changing content, digitise and migrate the thesis to any medium or format for the purpose of future preservation and accessibility.

4. DECLARATIONS

(a) I DECLARE THAT:

- I am the author and owner of the copyright in the thesis and/or I have the authority of the authors and owners of the copyright in the thesis to make this agreement. Reproduction of any part of this thesis for teaching or in academic or other forms of publication is subject to the normal limitations on the use of copyrighted materials and to the proper and full acknowledgement of its source.
- The digital version of the thesis I am supplying is the same version as the final, hard-bound copy submitted in completion of my degree, once any minor corrections have been completed.
- I have exercised reasonable care to ensure that the thesis is original, and does not to the best of my knowledge break any UK law or other Intellectual Property Right, or contain any confidential material.
- I understand that, through the medium of the Internet, files will be available to automated agents, and may be searched and copied by, for example, text mining and plagiarism detection software.

(b) IF I HAVE AGREED (in Section 2 above) TO MAKE MY THESIS PUBLICLY AVAILABLE DIGITALLY, I ALSO DECLARE THAT:

- I grant the University of Warwick and the British Library a licence to make available on the Internet the thesis in digitised format through the Institutional Repository and through the British Library via the EThOS service.
- If my thesis does include any substantial subsidiary material owned by third-party copyright holders, I have sought and obtained permission to include it in any version of my thesis available in digital format and that this permission encompasses the rights that I have granted to the University of Warwick and to the British Library.


5. LEGAL INFRINGEMENTS

I understand that neither the University of Warwick nor the British Library have any obligation to take legal action on behalf of myself, or other rights holders, in the event of infringement of intellectual property rights, breach of contract or of any other right, in the thesis.

Please sign this agreement and return it to the Graduate School Office when you submit your thesis.

Student's signature: 

Date: 16/02/2014



**Analysis and optimisation of ground based
transiting exoplanet surveys**

by

Simon Robert Walker

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Astronomy and Astrophysics

September 2013

THE UNIVERSITY OF
WARWICK

Contents

List of Tables	v
List of Figures	vi
Acknowledgments	x
Declarations	xi
Abstract	xii
Abbreviations	xiii
Chapter 1 Introduction	1
1.1 Extrasolar planets	1
1.2 Detection	2
1.2.1 Radial velocity	3
1.2.2 Transits	6
1.2.3 Microlensing	9
1.2.4 Direct imaging	11
1.2.5 Pulse timing	13
1.2.6 Astrometry	14
1.3 Planet characterisation	14
1.4 Projects	18
1.4.1 WASP	18
1.4.2 Kepler	21
1.4.3 HARPS	24
1.5 Planetary formation	25
1.5.1 Core accretion	25
1.5.2 Gravitational collapse	27
1.6 Migration	28

1.6.1	Disc migration	28
1.6.2	Dynamical scattering	31
1.7	Competing theories	32
1.8	Overall properties of exoplanet populations	34
1.9	Selection biases	37
1.10	Instrument technology	37
1.10.1	Charge Coupled Devices	37
1.10.2	Calibration procedure	38
1.10.3	Estimating the brightness of an object	39
1.10.4	Sources of uncertainty	40
1.11	Thesis structure	42
Chapter 2	Quantifying the WASP selection effects	43
2.1	Motivation	43
2.2	WASP project description	47
2.2.1	Hardware	47
2.2.2	Data reduction pipeline	48
2.3	Calculating the selection effects	51
2.3.1	Overview	51
2.3.2	Simulation parameters	52
2.3.3	Other parameters	53
2.3.4	Lightcurve synthesis	54
2.3.5	Testing the transit synthesis method	57
2.3.6	Model rejection	59
2.3.7	Implementation of the data modification process	60
2.3.8	Testing and verification	62
2.4	Computing the sensitivity map	64
2.4.1	Acceptance	66
2.4.2	Selection cuts	67
2.4.3	Period matching	70
2.5	Sensitivity maps	71
2.5.1	Planet trends	71
2.5.2	Shaping the sensitivity map	73
2.5.3	Combining individual sensitivity measurements	76
2.6	Summary	80

Chapter 3	Determining the underlying population of hot Jupiters using the WASP survey	81
3.1	Modifications for analysing the full WASP stellar sample	82
3.1.1	Composing a sample of stars	82
3.1.2	Calculating unknown stellar parameters	84
3.1.3	Analysis pipeline modifications	86
3.2	Shaping the sensitivity map	87
3.3	Sensitivity maps	89
3.3.1	Incorporating probability of transit	95
3.4	Comparing input systems to detections	95
3.5	Constraining the underlying hot Jupiter population	97
3.5.1	Joint constraint with Kepler	99
3.5.2	A new model for the underlying period distribution	105
3.5.3	Investigating the radius distribution	110
3.6	Discussion	112
Chapter 4	NGTS: Design and prototype	117
4.1	Introduction	117
4.2	Achieving the targets	120
4.2.1	Design	120
4.2.2	Prototype testing	123
4.3	La Palma prototype	123
4.3.1	Aperture photometry implementation	126
4.3.2	Removing trends	133
4.3.3	Limitations	134
4.4	Noise model	135
4.5	Prototype results	138
4.5.1	Precision	138
4.5.2	Noise colour	142
4.5.3	PSF sensitivity	144
4.5.4	Blending	147
4.6	Summary	151
Chapter 5	NGTS final instrument and planet catch simulations	154
5.1	The NGTS project	154
5.2	Reapplication of the noise analysis	155
5.2.1	Updated noise model	156
5.2.2	Photometric quality	158

5.3	Camera testing	162
5.3.1	Streak characterisation	163
5.3.2	Dark current measurement	165
5.4	Optimising the observing strategy	167
5.4.1	Saturation levels	168
5.4.2	Exposure time optimisation	171
5.5	Planet catch	173
5.6	Summary	180
Chapter 6 Conclusions and future work		183
6.1	Determining the underlying population of hot Jupiters with the WASP project	183
6.2	The Next Generation Transit Survey	184
6.3	Future work	185
Appendix A Sensitivity map trends		187

List of Tables

2.1	Literature hot Jupiter occurrence rates.	46
2.2	Limb darkening coefficients for three example temperatures	56
3.1	Parameter cuts used to restrict the WASP stellar sample.	84
3.2	Properties of the Kepler stellar and planetary sample.	100
3.3	Best fit parameters of the period model. Taken from Howard et al. [2012].	101
3.4	Best fit parameters for the power law model with Gaussian excess calculated from minimising the least squares.	107
3.5	Best fit parameters for the model (Eq. 3.8). The previous values from Table 3.4 have been repeated for comparison.	109
3.6	Cash best fit radius parameters	110
4.1	Comparison of detector features	125
4.2	Sky background values for days since New Moon, for La Palma . . .	137
4.3	Coordinates of the fields used for the crowding analysis	150
5.1	Comparison of NGTS detector features	155
5.2	Saturation coefficients for a given exposure time	170
5.3	Coordinates used for the NGTS planet catch simulations	175
5.4	Coefficients used for the Kepler occurrence rate	175
5.5	Predicted planet catch for NGTS	181

List of Figures

1.1	Orbital separation and planetary mass for the known exoplanets as of December 2011.	4
1.2	Radial velocity examples	5
1.3	Discovery transit data for HD 209458 b	6
1.4	Transit geometry schematic	8
1.5	Microlensing examples	10
1.6	Image of Fomalhaut	13
1.7	Rossiter-McLaughlin example	16
1.8	Distribution of spin orbit alignments	16
1.9	Examples of atmospheric characterisation	17
1.10	The WASP-South instrument	19
1.11	Transit of HD 209458 b from the WASP0 prototype	20
1.12	Transit of WASP-1 b taken with the WASP instrument	21
1.13	Examples of Kepler transits	23
1.14	Simulations of a protoplanetary disk showing gravitational instability	28
1.15	Simulation of Type II migration	30
1.16	Distribution of known planetary orbital periods	34
1.17	Orbital eccentricity and metallicity measurements for known exoplanets	35
1.18	Mass-radius relations	36
1.19	Analogy of a CCD	38
1.20	Geometry of a photometric aperture with its annulus.	40
2.1	Hot Jupiter distributions	45
2.2	Geometry of limb darkening	55
2.3	Synthetic stellar profiles	56
2.4	Demonstration of the transit removal and synthesis process	58
2.5	Initial test results for only varying the planetary radius of WASP-12 b	59
2.6	Schematic diagram of the lightcurve alteration process.	63

2.7	Examples of transit shapes	64
2.8	Statistics of the proposed synthetic systems	65
2.9	$\Delta\chi^2$ periodogram for planet WASP-12 b	67
2.10	Input period against recovered period for every accepted synthetic object	70
2.11	Sensitivity maps for WASP-12 b and WASP-7 b	72
2.12	Sensitivity maps generated for the WASP stars ordered by V magnitude	74
2.13	Orion detection map	75
2.14	Maps made from individual selection cuts	77
2.15	Combined sensitivity map for the planet hosting stars	78
2.16	One dimensional sensitivity profiles marginalised over the second axis	79
3.1	Dwarf probability used in the classification of stars for the WASP reduced stellar input catalogue.	83
3.2	Distributions of the stellar parameters used to reduce the input sample	85
3.3	Distribution of the number of synthetic transiting systems inserted per lightcurve.	88
3.4	ORION detection map	89
3.5	Sensitivity maps made from each selection cut	90
3.6	Sensitivity maps created from the two different WASP stellar samples	91
3.7	Ratio of the sensitivity maps made from the two stellar samples . . .	92
3.8	Histogram of values in the map ratio > 0	92
3.9	Flattened sensitivity profile.	93
3.10	Sensitivity maps built up by splitting the magnitude range of study into three	94
3.11	Sensitivity map including the probability of transit	96
3.12	Comparison of the proposed and best fit planetary radius	97
3.13	Comparison of the input and detected impact parameters from the analysis	98
3.14	Correction factors calculated by inverting the WASP planets through the sensitivity map.	99
3.15	Observed and underlying hot Jupiter distributions	100
3.16	Confirmed WASP planet, and Howard et al. [2012] Kepler candidate distributions	102
3.17	Kepler candidate planet occurrence as a function of planet radius and orbital period for $P < 50$ days	103

3.18	Occurrence rate in planets per star showing the Kepler and WASP results	104
3.19	The occurrence rate in planets per star after increasing the bin resolution	105
3.20	Howard et al. [2012] occurrence rate models for orbital period	106
3.21	Modelling the occurrence rate using the Howard et al. [2012] model .	107
3.22	Results for the new occurrence rate model for giant planets	108
3.23	Fits to the occurrence rate of giant planets considering Poisson statistics	109
3.24	Result of characterising the radius distribution	111
3.25	Rotation rates for T Tauri stars in the Orion Nebula Cluster	113
3.26	Planetary mass against orbital period for planets $R_p \geq 8 R_\oplus$	113
3.27	Corrected Roche limit ratio distribution	114
3.28	Radius for hot Jupiter planets with $R_p \geq 8R_\oplus$	115
4.1	Planet radius histogram for the known exoplanets	118
4.2	Planet radii and host star radii for planets detected from the ground	119
4.3	Throughput of the NGTS instrument	121
4.4	A complete NGTS unit, assembled at the Geneva Observatory. . . .	122
4.5	Weather quality for Paranal	123
4.6	Computer generated renders of the NGTS facility	124
4.7	Images of the NGTS prototype installed in La Palma	125
4.8	Graphic representing the nights for which NGTS prototype data exists	126
4.9	An example NGTS prototype image, taken on 2009-10-07	128
4.10	Example aperture used by the photometry pipeline	128
4.11	QE quoted from the Andor brochure	129
4.12	Histogram of the estimated zero points for the field of the night of 2009-10-07	130
4.13	Geometry of the light path length through the atmosphere, from Birney et al. [2006].	131
4.14	Extinction behaviour for a group of stars	132
4.15	Extinction measurements from the Carlsberg Meridian Telescope . .	133
4.16	SYSREM coefficients for 2009-11-19	134
4.17	Example lightcurve from 20010-02-04	136
4.18	Sky background measurement from the NGTS prototype	137
4.19	Noise model for the NGTS prototype	139
4.20	Properties of the NGTS prototype dataset	140
4.21	Fractional rms for the field collected on 2009-11-19, pre-meridian . .	141

4.22	Measured fractional noise values for the night of 2009-11-19	143
4.23	Binned median precision of bright stars	144
4.24	Example data from a single lightcurve collected during 2009-10-07. .	145
4.25	Examples of the focus levels shown in Fig. 4.20	146
4.26	PRF responses	148
4.27	NGTS yearly coverage	149
4.28	Histogram of the fraction of objects which are un-blended	150
4.29	Transits observed by the NGTS prototype	152
5.1	Comparison of noise models for the prototype and Geneva instruments	157
5.2	Binned noise models for NGTS for a range of stellar magnitudes . .	159
5.3	Binned noise models for NGTS for a range of exposure times	160
5.4	Fractional rms for the data collected at Geneva on 2012-08-17	161
5.5	Binned fractional rms of stars with magnitudes $11 \leq I \leq 8$	162
5.6	Example streaks from on-sky data	163
5.7	Example of the two sided CCD sensitivity	165
5.8	Streak examples	166
5.9	Illuminated region of the CCD tested at Leicester	166
5.10	Median dark current values at different temperatures	168
5.11	Schematic of a Gaussian centred at (0.3, 0.3), and a FWHM of 1.5 pixels	169
5.12	Distribution of pixel flux fractions from a Monte Carlo simulation . .	170
5.13	Saturation behaviour with exposure time	171
5.14	The calculated high precision region	172
5.15	Analysis of optimising the exposure time	174
5.16	Assumed period window function for the NGTS project.	176
5.17	Planet catch simulations in units of velocity semi major amplitude .	179
5.18	Planet catch simulation results in units of planet radius	180
A.1	Sensitivity maps in order of T_{eff}	188
A.2	Sensitivity maps in order of Metallicity.	189
A.3	Sensitivity maps in order of R_{\star}	190
A.4	Sensitivity maps in order of M_{\star}	191

Acknowledgments

I'd like to thank my supervisor Dr. Peter Wheatley, without whom this thesis definitely would not be possible. His support and assistance during the course of my PhD were invaluable. Thanks to everyone in the Astronomy and Astrophysics group at Warwick, who have provided a constant stream of support and entertainment. I want to thank both the WASP and NGTS consortiums for making this work possible, especially the members who I have worked directly with: Dr. Richard West and Prof. Don Pollacco.

I'd like to thank my parents for both supporting me during both the happy and stressful times, and for being wonderful. I would not be where I am today without them. Without good office mates, a person may go slightly insane. Alternatively *with* good office mates a person may go slightly insane. Whichever is more appropriate I've had fantastic office mates, and special thanks goes to Jon, Lieke, Tom and João. In particular, without João I could not have achieved quite so much, and would not have discovered some fantastic music. I'd like to thank my friends and fellow year-mates Rachel, Matt, Steve and Pete for countless hours of laughter. And finally I'd like to thank my girlfriend Zoe who has been wonderful, managed to not get *too* infuriated with me.

Declarations

I declare that the work presented in this thesis is my own except where stated otherwise, and was carried out entirely at the University of Warwick, during the period October 2009 to May 2013, under the supervision of Dr. Peter Wheatley. Work in Chapters 4 and 5 were carried out as a member of the NGTS consortium. The research reported here has not been submitted, either wholly or in part, in this or any other academic institution for admission to a higher degree.

Contributions based on this thesis are:

- National Astronomy Meeting, Manchester, UK, March 2012, Poster presentation: *Selection effects of the SuperWASP project.*
- International Astronomical Union Symposium 299, Victoria, Canada, June 2013, Poster presentation: *Determining the population and evolution of hot Jupiters by quantifying WASP selection effects.*

Abstract

One of the most surprising aspects of the exoplanet population is the existence of Jupiter sized planets orbiting close to their parent stars. It is currently uncertain how these planets reached such small separations, and they are thought to be markers for the dominant migration mechanism. The Wide Angle Search for Planets (WASP) project is ideally suited for studying these planets, as it has detected the largest number of hot Jupiters to date. I have inverted the observed sample of WASP planets to calculate the underlying population of hot Jupiters through a quantitative study of the selection biases in the WASP project. To achieve this, I synthesised transiting systems and inserted them into WASP data to calculate the probability of detection. The observed population of WASP planets is then corrected through application of this probability to determine the underlying population. I find a clear pile up in the underlying population at orbital periods between 3 to 5 days, and apply a joint constraint with the underlying population measurement from the Kepler project to propose a new model for the underlying population of giant planets. I propose a model consisting of a rising power law in period with index 1.0 ± 0.3 , with a Gaussian excess at 3.7 ± 0.1 days to model the period pile up. The observed period pile up places crucial constraints on models of hot Jupiter migration.

The Next Generation Transit Survey (NGTS) is a new transiting exoplanet survey designed to find Neptunes and super Earths around nearby stars. These stars will be bright, allowing the characterisation of the bulk densities and atmospheric compositions by current and next generation instruments. These planets are numerous but they cause shallow transits, so we must achieve a higher level of precision than has been previously achieved by a wide angle ground based survey. To test the design characteristics and that the required high levels of precision would be achieved, prototype instruments were constructed on La Palma and Geneva. We found that we were able to reach $< 0.1\%$ precision on an ensemble of bright stars on the typical transit timescales, indicating that the instrument will be capable of detecting smaller planets. We use Monte Carlo simulations coupled with a detectability analysis to predict that NGTS will detect 200 Neptunes and 30 super Earths with the next generation instrument ESPRESSO.

Abbreviations

WASP Wide Angle Search for Planets	RV Radial Velocity
NGTS Next Generation Transit Search	KIC Kepler Input Catalogue
PSF Point Spread Function	CTE Charge Transfer Efficiency
PRF Pixel Response Function	R_{\odot} Solar radius, 7.955×10^5 km
CCD Charge Coupled Device	R_J Jupiter radius, 71 492 km
FWHM Full Width Half Maximum	M_{\odot} Solar mass, 1.989×10^{30} kg
ESO European Southern Observatory	M_J Jupiter mass, 1.898×10^{27} kg
RMS Root Mean Squared	M_{\oplus} Earth mass, 5.972×10^{24} kg
ADU Analogue-to-Digital Unit	AU Astronomical Unit, 1.496×10^8 km
FOV Field Of View	

Chapter 1

Introduction

1.1 Extrasolar planets

Planets that exist outside our Solar System and therefore do not orbit the Sun are classed as extra-solar planets, or exoplanets. Exoplanets had been predicted [e.g. Newton, 1726, and before], but the first discovery of an exoplanet was in 1992 by Alex Wolszczan and Dale Frail around the millisecond pulsar PSR1257+12 [Wolszczan & Frail, 1992]. Since then 986 exoplanets have been confirmed as of 2013/9/30¹.

With increasing numbers of exoplanets and exoplanetary systems comes a greater understanding of their population and formation, but they are full of surprises. The first exoplanet detected around a Sun-like star 51 Peg, by Mayor & Queloz [1995] was found to be a half Jupiter mass planet orbiting at a distance seven times closer than Mercury to the Sun. This class of planet has no analogue in the Solar System and was a curious case.

The hot Jupiter class of planet that includes this example is loosely defined as planets with similar mass to Jupiter, orbiting their parent star in less than 10 days. These planets were first to be discovered due to their large radial velocity signal and then large transit signature (see Section 1.2), and are strongly selected for in radial velocity and especially transit surveys.

The exoplanet sample now includes similar systems to 51 Peg b, but also systems similar to the Solar System with multiple planets of different sizes, some orbiting in resonance with each other. Improved instrumentation has enabled the detection of near-Earth sized planets (e.g. CoRoT-7 b, [Queloz et al., 2009]) at increasing orbital periods. The huge diversity of the exoplanet population and in some cases the lack of similarity with the Solar System is extremely interesting, and

¹<http://exoplanet.eu>

requires development of theories of planetary system formation to encompass both singular hot Jupiter systems and multi-planet systems observed.

As well as studying the population, individual systems are analysed to ascertain properties of the planet and its composition. With high-precision space telescopes the planetary atmospheres are probed, determining the atmospheric chemical composition, enabling the search for dis-equilibrium chemistry and potential biomarkers. Combining detection techniques allows for planetary density measurements to be made, resulting in some planets which are more dense, and some planets that are less dense than *a priori* assumptions suggest.

In the following section the methods used to detect exoplanets are described. The detection techniques often complement each other providing more information than each individual technique alone. These same techniques may also be used for planet characterisation, and some uses of the detection techniques beyond inferring the presence of the planet are described in Section 1.3. Some example projects of exoplanet surveys involved in the detection or characterisation of exoplanets are described in Section 1.4. We use the observed population of exoplanets as evidence of the formation and subsequent evolution of planets. The two primary theories for the formation of planets out of a disk of material orbiting a young star are given in Section 1.5. The existence of hot Jupiters at distances much closer than is possible to form such large planets is compelling evidence that planets undergo migration from their initial formation location to their observed positions. Possible mechanisms for this process are described in Section 1.6, and a discussion of the corresponding evidence for possible formation and migration methods is given in Section 1.7. With a statistical sample of exoplanets the population can be studied, where common properties or behaviours can be determined. This is discussed in Section 1.8 with a corresponding caution about studying the observed population as biases exist. These biases are due to the sensitivities exhibited by the various detection techniques, and may cause incorrect conclusions to be drawn about population (Section 1.9). The process of performing astronomical observations, technologies used and uncertainties therein are described in Section 1.10.

1.2 Detection

Multiple methods exist for detecting exoplanets. Different detection methods have different regions of this parameter space where they are sensitive to detecting planets. In this section some of the different detection techniques are described. Figure 1.1 summarises the exoplanet population in terms of orbital separation and

planet mass, and shows the detection limits for various detection methods.

1.2.1 Radial velocity

A body orbiting a star causes a reflex motion in the star around the centre of mass of the system which, depending on the masses and orientations involved, may be detectable with a high-precision spectrograph.

Light from a moving object changes frequency due to the Doppler effect. The projection of this motion along the line of sight to the object causes blue and red shifting of the light received. In the classical form the radial velocity v_r is observed to be

$$v_r \approx \left(\frac{\Delta\lambda}{\lambda_{em}} \right) c \quad (1.1)$$

where $\Delta\lambda$ is the change in wavelength observed from rest frame observations λ_{em} , and c is the speed of light.

This signal is detected through observing an object with a high-resolution spectrograph. Multiple spectra are taken of the star and radial velocity signals are measured from the time-varying spectra. The presence of a massive body in the system displaces the measured spectral lines from their rest wavelength in a periodic manner. To increase the precision of the radial velocity measurement, the spectral lines are typically cross correlated with a template spectral line. This is possible because each line is displaced by the same amount.

The eccentricity e and argument of periastron ω are measured from the shape of the radial velocity signal. A circular orbit causes a sinusoidal signal, whereas elliptical orbits cause deviations from sinusoidal behaviour, and this deviation can be modelled through varying the ellipticity and argument of periastron. Another key feature of the radial velocity signal is the semi-amplitude defining maximum radial velocity signal observed. The measured semi-amplitude of the periodic signal K_\star along with the independently determined eccentricity and orbital period allow the calculation of the mass ratio of the planetary system, with an uncertainty by an unknown factor due to the system inclination. The velocity semi-major amplitude is given by

$$K_\star = \left(\frac{2\pi G}{P} \right)^{1/3} \frac{m_p \sin i}{(M_\star + m_p)^{2/3}} \frac{1}{\sqrt{1 - e^2}} \quad (1.2)$$

[Cumming et al., 1999] where m_p and M_\star are the masses of the planet and star, e is the eccentricity, P the orbital period, G the gravitational constant ($G = 6.67 \times 10^{-11} \text{ m}^3 \text{ kg}^{-1} \text{ s}^{-2}$) and i is the inclination. The orientation of the system and therefore the inclination i cannot be determined from radial velocity measure-

Exoplanets: 149+463+14+13+26=665 (Dec 2011)

Orbit Period (Earth Years – assuming solar mass star)




Figure 1.1: Orbital separation and planetary mass for the known exoplanets as of December 2011. Points mark the detections and are coloured by the detection technique used to discover them. Lines mark the predicted sensitivity limits for the detection techniques. Image courtesy of Keith Horne.




Figure 1.2: Left: discovery data for 51 Peg b [Mayor & Queloz, 1995]. Right: example radial velocity measurements with best fit solution for HD 100777 b. The top panel shows the measured radial velocity profile, the bottom panel shows the residuals to the best fitting solution. Image from Naef et al. [2007].

ments alone, so planetary systems which are only detectable by the radial velocity signal do not have a unique mass value, only the projected mass $m_p \sin i$ is available. Figure 1.2a shows the discovery data for 51 Peg b showing smooth sinusoidal variations characteristic of a massive planet on a circular orbit. Figure 1.2b shows an example dataset collected by HARPS for the star HD 100777 around which a planet with mass 1.16 times the mass of Jupiter ($M_J = 1.90 \times 10^{27} \text{ kg}$) orbits on a 384-day orbit. The deviation from a sinusoidal shape implies an eccentric orbit; from Keplerian modelling the eccentricity was measured to be $e = 0.36 \pm 0.02$ [Naef et al., 2007].

The velocity semi-amplitude is proportional to $m_p/P^{1/3}$ (from Eq. 1.2, assuming $m_p \ll M_\star$) so the radial velocity technique excels at finding large mass planets in short period orbits around their host stars. The technique does not require a full period to be observed and so is also successful at detecting massive planets relatively far from their host star, as shown in Fig. 1.1 by the large population of planets with masses $m_p \gtrsim 100 M_\oplus$ ($M_\oplus = 5.97 \times 10^{24} \text{ kg}$) at separations $a \sim 1 \text{ AU}$ ($\text{AU} = 1.496 \times 10^{11} \text{ m}$). Eccentricities $e \gtrsim 0.6$ also pose difficulties for detecting radial velocity signals at shorter periods due to the sparse sampling of the periastron passage [Cumming, 2004]. This bias towards more significant signals suggests that the detected population is incomplete as less significant signals from




Figure 1.3: Discovery transit data for HD 209458 b consisting of two epochs of observations overlaid in time since the mid-transit point [Charbonneau et al., 2000].

true planets may have been missed.

1.2.2 Transits

Light from a star is attenuated when a planet crosses its disk. This decrease in flux may be observed and the presence of the planet inferred. Figure 1.3 shows the transit data for the first exoplanet detected using this method: HD 209458 b. The maximum loss of light from a transit is

$$\delta_{tra} \approx \left(\frac{R_p}{R_\star} \right)^2 \left[1 - \frac{I_p}{I_\star} \right] \quad (1.3)$$

[Seager, 2011] where R_p is the planetary radius, R_\star is the stellar radius, I_p is the flux emitted by the planet and I_\star the flux emitted by the star. Limb darkening on the star causes a non-uniformity across the stellar disk and changes the shape of the transit. As the deviation from normal stellar background is proportional to the ratio of the object areas the transit signal is typically small, with even the largest planets only causing a transit signal of a few percent. For example Jupiter crossing the Sun creates a transit signal of $\delta_{tra} = (R_J/R_\odot)^2 = 1.045\%$ where $R_\odot = 7.00 \times 10^5$ km.

Similarly when the planet passes behind the star a secondary eclipse occurs with observable depth

$$\delta_{occ} \approx \left(\frac{R_p}{R_\star} \right)^2 \frac{I_p}{I_\star}. \quad (1.4)$$

The deviation from predicted secondary eclipse timing provides a tight constraint on the eccentricity of the planetary orbit [Perryman, 2011].

Observing a transit requires a close alignment between the orbital plane of the planetary system and the line of sight. Transits can only be detected within a narrow region of orbital inclinations $i \geq i_c$ where

$$\sin i_c \geq \frac{(R_p + R_\star)}{a} \quad (1.5)$$

[Seager, 2011]. A lightcurve of a transiting system provides information about the objects contained. The orbital period of the system is measured from the time between transit events, from which the orbital separation a is calculated by substituting Kepler's third law [Haswell, 2010]:

$$\frac{a^3}{P^2} = \frac{G(M_\star + m_p)}{4\pi^2} \quad (1.6)$$

Generally $m_p \ll M_\star$, so given an estimate of the stellar mass M_\star (e.g. from stellar spectroscopy, or asteroseismology) the separation can be determined [Haswell, 2010]:

$$a \approx \left(GM_\star \left(\frac{P}{2\pi} \right)^2 \right)^{1/3} \quad (1.7)$$

Figure 1.4 illustrates the geometry of the transiting system, with impact parameter b showing that the planet is not passing across the centre of the star, but is offset by projected distance b in units of stellar radii. There are four observable quantities which characterise the duration and profile of the transit: the orbital period P given by the spacing between transits, the transit depth δ_{tra} , the interval between first and fourth contacts T_{tot} , and the interval between the second and third contacts T_{full} [Perryman, 2011]. These observables are used in the following geometric equations [Seager, 2011]:

$$\Delta F = \left(\frac{R_p}{R_\star} \right)^2 \quad (1.8)$$

$$b^2 = \frac{(1 - \sqrt{\delta_{tra}})^2 - (T_{full}/T_{tot})^2(1 + \sqrt{\delta_{tra}})^2}{1 - (T_{full}/T_{tot})^2} \quad (1.9)$$

$$\frac{R_\star}{a} = \frac{\pi}{2\delta_{tra}^{1/4}} \frac{\sqrt{T_{tot}^2 - T_{full}^2}}{P} \left(\frac{1 + e \sin \omega}{\sqrt{1 - e^2}} \right), \quad (1.10)$$

where ω is the argument of periastron, and other parameters have been defined.




Figure 1.4: Illustration of a transit showing the geometry discussed in Section 1.2.2. Above is a schematic diagram of the transiting system seen as the planet crosses the disk of the star. Below is an idealised transit lightcurve across a uniform brightness stellar disk. Specific contact points are labelled. From Seager [2011].

With an independent estimate of the stellar radius (e.g. from stellar models) the planetary radius can be directly obtained from Eq. 1.8. Recent development in Bayesian techniques allow the direct estimation of the orbital eccentricity from the transiting lightcurve by using prior constraints on the host star density and studying the posterior eccentricity distribution [Dawson & Johnson, 2012].

The transit detection process starts with searching the lightcurve of a star for a repeated box-like feature. Various algorithms for detecting transits exist [Enoch et al., 2012], the box-least-squares (BLS) algorithm [Kovács et al., 2002] is the most common and was developed to perform this and uses a grid search in orbital period, transit depth and epoch and compares the generated box feature with the lightcurve to calculate goodness of fit. If a significant transit-like signal is found then the planet is commonly confirmed with the radial velocity technique, providing both an independent confirmation of the validity of the system and with the full set of system parameters solved, the candidate planet can be confirmed as a real planet or false positive. Common false positives which give transit-like signals include shallow eclipsing binary stars, where a grazing transit can produce a $\sim 1\%$ repeated dip; an eclipsing binary nearby to the target star, which injects a transit-like signal into the lightcurve; or sunspots on the target star which also cause dips in flux though these are typically not on the same timescales.

The significance of a transit-like signal is dominated by the noise level of the lightcurve. Projects built to search for transiting planets are designed to produce the highest precision lightcurves possible. The probability that a planet will cross the disk of a star along the line of sight from the observer is given as

$$p = \frac{R_{\star}}{a} \quad (1.11)$$

calculated from the solid angle on the sphere swept out by the planet's shadow [Perryman, 2011] and is generally very small even for hot Jupiters (e.g. a planet orbiting at 0.1 AU around a $1R_{\odot}$ star has a transit probability of 4.7%) so transiting exoplanet surveys typically aim to observe as many stars as possible simultaneously.

1.2.3 Microlensing

Under General Relativity the presence of mass deforms light travel paths through spacetime. Similar to optical systems when an object passes behind a gravitational lens the received image is distorted and magnified. The gravitational lens in this case is a massive object such as a star. The characteristic length scale for this distortion is the Einstein radius R_E such that

$$R_E = \left[2R_S \frac{D_L D_{LS}}{D_S} \right]^{1/2} \quad (1.12)$$

where R_S is the Schwartzchild radius

$$R_S = 2GM_L/c^2 \quad (1.13)$$

and the D terms the distances from the observer to the lens (D_L), the observer to the source D_S and the distance from the lens to the source D_{LS} . This radius characterises the Einstein ring around the lens star (see Fig. 1.5a for a schematic of the geometry.) A lens star passing in front of a background star will cause magnification and distortion of the background star, provided the projected angular separation between the two objects is small. The distortion causes a deviation in the source image path, a secondary image of the source to appear on the opposite side of the Einstein ring (I_- in Fig. 1.5a) and for both images to appear distorted from their original shape. Magnification occurs because the flux from each image is the product of the (constant) source brightness, and solid angle subtended by each image which increases (Fig. 1.5a) [Perryman, 2011]. For single star lenses the region of maximum magnification (*caustic* region) of the background star is directly

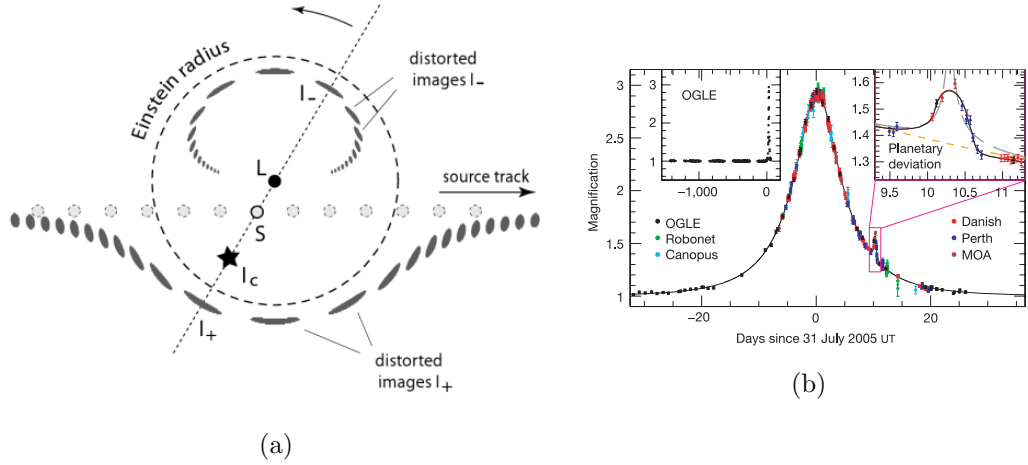


Figure 1.5: Left: geometry of the microlensing system, projected into the sky plane of the observer and reference frame of the lens object L . The source object S travels across the system along the source track causing the images I_+ and I_- to appear. The images rotate whilst remaining on opposite sides of L as marked by the diagonal dashed line. Image adapted from Paczynski [1996]. Right: microlensing lightcurve for OGLE-2005-BLG-390L b a $5.5 M_{\oplus}$ planet [Beaulieu et al., 2006]. Inset: the microlensing signal caused by the presence of the planet.

behind the lens along the observers line of sight. At this point the magnification of the source is formally infinite but the alignment is never perfect as the two objects are not points. The magnification can be large, the largest to date is a magnification of 3000 ± 1100 times [Dong et al., 2006].

A third body such as a planet orbiting the central lens will distort the caustic region to contain a region behind the planet, where the particular shape is determined by the orbital separation of the lens star and its planet. As the planet orbits, this region coincides with the background star to create divergence from smooth magnification profile of the background star, and adds a feature to the otherwise smooth magnification event. Figure 1.5b shows the lightcurve for OGLE-2005-BLG-390L where the inset shows the effect of the orbiting planet, a $5.5 M_{\oplus}$ planet. The overall smooth shape shown in the main figure is caused by the background star passing behind the lens, and the imperfection is caused as the position of the planet during its orbit causes a caustic region to fall on the path of the background star inducing extra magnification.

Microlensing events are rare, and occur only once per target star, so projects provide a constant monitoring of the galactic bulge where the galactic stellar density is the highest. Microlensing is sensitive to orphan planets that do not orbit host stars. Should the path of an orphan planet cross the path of a background star

a microlensing event is detectable. Sumi et al. [2011] report the detection of 10 possible rogue planet events with 6 confirmed by the OGLE project and quote an abundance rate of these objects as $1.8_{-0.8}^{+1.7}$ planets per star.

1.2.4 Direct imaging

Repeated high angular resolution images of a star are taken. Planets in the system may be detected, either through reflected light from the host star or directly detecting thermal emission from the planet. For reflected light of wavelength λ the planet/star flux ratio can be written

$$\frac{f_p(\alpha, \lambda)}{f_\star(\lambda)} = p(\lambda) \left(\frac{R_p}{a} \right)^2 g(\alpha) \quad (1.14)$$

where $p(\lambda)$ is the geometric albedo and $g(\alpha)$ is a phase dependant function of α , the angle between observer and star subtended at the planet [Perryman, 2011].

The planet can be approximated with a blackbody and so emits a thermal spectrum given by the Planck function

$$B_\lambda(T) = \frac{2hc^2}{\lambda^5} \frac{1}{e^{\frac{hc}{\lambda k_B T}} - 1} \quad (1.15)$$

where $B_\lambda(T)$ is spectral radiance in units of $\text{W sr}^{-1} \text{m}^{-3}$ at wavelength λ , and h is Planck's constant, c is the speed of light, k_B is Boltzmann's constant, T is the temperature of the blackbody. The wavelength of maximum emission from a blackbody is related to the temperature of the body by Wien's displacement law

$$\lambda_{max} T = 2.897 \times 10^{-3} \text{ m K} \quad (1.16)$$

and is derived from Eq. 1.15. This relation shows that the peak emission of a cooler body is at longer wavelength. The planetary effective temperature is lower than the stellar effective temperature so the flux ratio between planet and star is maximised at longer wavelengths.

Combining the two effects increases the contrast ratio between planet and star, but this ratio is small as the star is typically many orders of magnitude brighter than the planet even at favourable wavelengths, for example the Jupiter/Sun flux ratio is $\sim 10^{-9}$. To reduce the stellar flux and increase the contrast ratio further, coronagraph masks are employed to block the light from the central star. A physical mask is inserted into the optical path which minimises the stellar flux, and self-cancelling interference is employed to reduce the diffraction patterns induced by the

central mask.

Planets can only be detected through direct imaging at large separations as the glare from the star masks the flux from the planet. Atmospheric refraction and seeing limit the minimum separation a planet is detectable. To reduce the extent of the stellar glare a high angular resolution is required, making the planet visible. Adaptive optics, or taking observations from space e.g. the Hubble Space Telescope (HST) are used for direct imaging observations to minimise or negate the effects of the atmospheric refraction and seeing. The wavefront from a bright reference star or synthetic laser guide star is sampled at timescales on the order 1 ms, and correction actuators apply the inverse wavefront to a deformable mirror in the camera system. By correcting for atmospheric turbulence the angular resolution can be increased to the diffraction limit of the telescope. The angular resolution is approximated by

$$R \approx 1.22 \frac{\lambda}{D} \quad (1.17)$$

where D is the diameter of the telescope aperture. By using large aperture telescopes the angular resolution can be increased further.

Images are plagued with speckle noise due to random intensity patterns by the interference of incoming wavefronts. This interference is caused by atmospheric effects and instrumental imperfections, adding noise which does not reduce with increasing exposure time. To reduce the speckle noise a technique called angular difference imaging is employed. The speckle noise is correlated between exposures, with a slowly changing intensity pattern in the image plane. This orientation of this pattern is fixed with the rotation of the telescope. By rotating the telescope between exposures the speckle noise is averaged but any background objects such as planets remain visible, but their positions offset. The images are then stacked in the original orientation to reduce speckle noise.

To reject false positive objects such as background objects multi-epoch measurements are made. The presence of a planet is inferred from common proper motion objects, and the orbit is estimated from the residual small positional shifts between the host star and planet.

Fomalhaut b was detected in 2008 through the direct imaging method (see Fig. 1.6). It orbits at a distance of 113 AU from its host star and has a mass constrained to $< 3M_J$ in a system with a clear dust ring. The mass limit is inferred through modelling of the stability of the dust ring with the presence of Fomalhaut b, a larger mass would disrupt the dust ring. A debate is ongoing as to the validity of Fomalhaut b as a planet as Kalas et al. [2008] report non-detections at longer




Figure 1.6: Image of Fomalhaut using coronagraphic observations from the HST showing the dust ring and (inset) locations of Fomalhaut b in two epochs. Image courtesy of NASA.

wavelengths where a young planet would emit the most flux, but a re-analysis of the data confirms the presence of Fomalhaut b [Galicher et al., 2013]. HR 8799 is a star hosting a debris disk and four planets with mass $m_p > 5 M_J$ and wide orbits (> 10 AU) [Marois et al., 2010]. The planets are co-planar and pose questions for planet formation theorists: did they form *in situ* or form closer in and migrate out together?

Direct imaging can detect planets at much larger separations than other techniques, and directly measures the light from the planet’s surface allowing analysis of the atmospheric composition.

1.2.5 Pulse timing

Millisecond pulsars are old ($\sim 10^9$ yr) rapidly rotating neutron stars in a binary system, where accretion from the secondary spins up the rotation of the primary until its rotational period is 1ms or shorter [Wolszczan & Frail, 1992]. The rotation of the primary star coupled with the change in period are predictable to a very high accuracy. Observing these systems provides a highly sensitive way of identifying perturbing elements in the orbit. The first exoplanetary system discovered was found around the millisecond pulsar PSR1257+12 and contains two planets of mass $M \sin i = 3.4 M_\oplus$ and $2.8 M_\oplus$. Wolszczan & Frail [1992] were able to measure

changes in the orbital period of the pulsar corresponding to radial velocities of 1 m s^{-1} from the two planets. Since the discovery a third massive body has been suggested [Wolszczan et al., 2000] to explain a further 25.3 day periodicity. Twelve planets around variable stars have been detected², five of which orbit a pulsar while the rest orbit pulsating stars, or stellar systems in which the presence of a planet has been inferred from deviations from a periodic signal. For example the eclipsing white dwarf binary NN Ser [e.g. Beuermann et al., 2010; Marsh et al., 2014] in which the planet was detected through transit timing variations of the eclipses. Since variable objects have a high sensitivity to perturbing bodies, and therefore the low sample size suggests that they are rare.

1.2.6 Astrometry

This technique relies on the reflex motion of a star due to the presence of a massive body, similar to the radial velocity technique but observes the stellar motion projected on the plane of the sky. This method has not yielded any new planets to date as the positional change due to this perturbation is on the order $1 \mu\text{as}$ to 1 mas [Perryman, 2011] but astrometric signals of known planets have been detected [e.g. Benedict et al., 2006]. The ESA mission GAIA [Perryman et al., 2001] is expected to detect around 2500 planets with semi-major axes of 3 - 4 AU, provided the target astrometric precision of $12 \mu\text{as}$ is achieved [Casertano et al., 2008].

1.3 Planet characterisation

The transit method coupled with radial velocity measurements allow the radius and mass of the planet to be determined, giving constraints on the density of exoplanets. This proves that hot Jupiters are gas giants and do not have a rocky composition. Many planets have been discovered which are inflated beyond expected values (E.g. WASP-57 b, HAT-P-32 b. This is thought to be due to irradiation from the host star, or tidal heating as the planetary orbit is circularised. Over-dense planets have also been detected, for example HAT-P-20 b, suggesting a high metal content [Bakos et al., 2011]. The densities of planets can be compared to compositional models [e.g. Fortney et al., 2007; Seager et al., 2007] and the likely internal composition determined.

Asteroseismology is a method to determine pulsation and oscillation modes of stars providing information about the internal structure of stars. This information

²<http://exoplanet.eu>

applies further observational constraints on the stellar structure, types and properties of many stars. When applied to exoplanet host stars, these constraints allow high precision measurements of the exoplanet properties, as the stellar parameters are known to a high precision. This technique has been applied to data obtained by Kepler to constrain properties of more than 500 main-sequence and sub-giant stars, test theories of stellar evolution which improves our understanding of exoplanetary host stars [Chaplin et al., 2011].

Transiting planets allow the study of the Rossiter-McLaughlin effect of planetary systems to determine the projection of the orbital misalignment between the planet’s orbit and the stellar spin axis onto the sky. It is a deviation from expected behaviour of the radial velocity profile of an orbiting planet only visible in transiting planets. An example is shown in Fig. 1.7. The stellar disk is split by the stellar rotation axis into one half approaching and one half receding. As the planet occults the disk the ratio of these areas changes, causing a deviation from the bulk radial velocity caused by the planet. The observed profile depends on whether the area obscured is approaching or receding relative to the star’s bulk motion, the projected stellar rotation rate $v \sin i$, and the mutual inclination of stellar spin axis and planet planetary orbit ϕ [Queloz et al., 2000].

The orientation of the planetary orbital rotation axis relative to the stellar spin axis can be inferred through the measurement of the Rossiter-McLaughlin effect, and has consequences for theories of planet formation. Figure 1.8 shows the measurements of the absolute projected spin orbit alignment angle $|\lambda|$ for planets where this property has been estimated. A strong peak at $|\lambda| = 0$ is visible showing that the bulk of the measured planets have an aligned orbit, but large tails of the distribution suggest that a significant fraction of stars have misaligned planets. A smaller peak is arguably apparent at $|\lambda| = 180$ indicating that a subset of planets are in retrograde orbits, orbiting opposite to the rotation direction of their host star.

Transmission spectroscopy is the method of determining the atmospheric composition for a planet by observing the transit at multiple wavelengths. During the transit of a planet across the disk of a star the planet blocks a fraction of the light, but the planet is not an opaque disk. The composition of the planetary atmosphere partially absorbs the flux corresponding to the excitation energy of the molecules in the atmosphere or scattering from dust or clouds, leading to a varying opacity with wavelength. At the wavelength of strong molecular absorption the atmosphere appears more opaque and the effective silhouette appears larger causing a different transit depth to be measured. The planet can be treated as an opaque disk with a thin atmosphere ring around the circumference, with characteristic scale height H

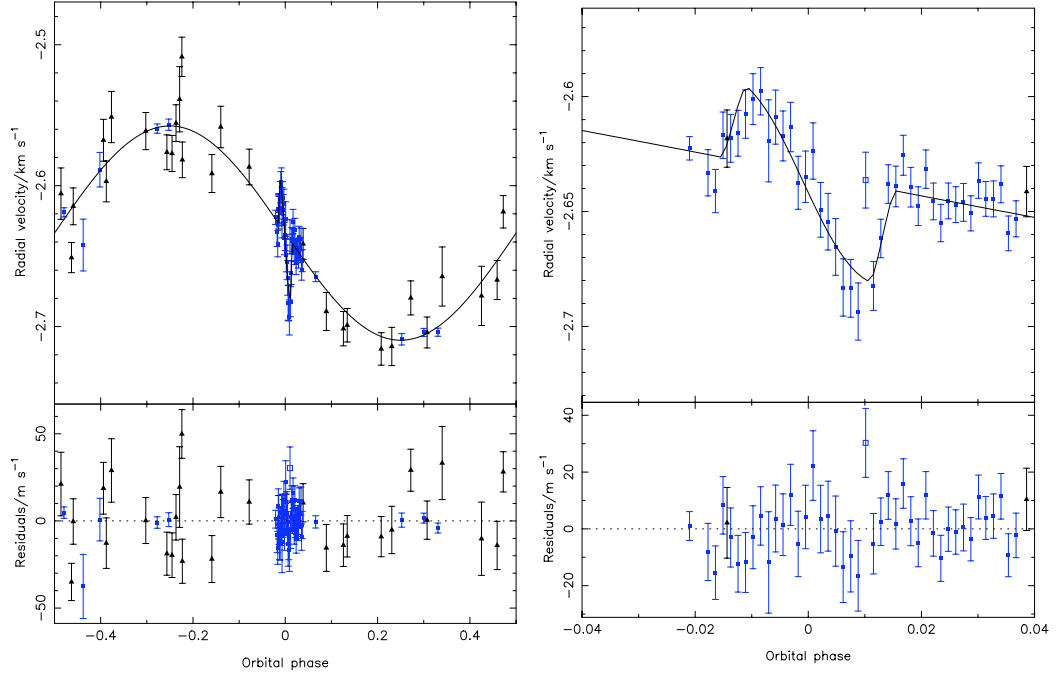


Figure 1.7: Radial velocity measurements for WASP-25 b showing (left) the full orbit and (right) the radial velocity measurements during and around the transit at phase 0. Black points represent data taken from CORALIE and blue points from HARPS. The right figure shows the characteristic defect of the radial velocity profile caused by the Rossiter-McLaughlin effect. From [Brown et al., 2012].




Figure 1.8: Spin orbit alignment values where an estimation has been performed.




Figure 1.9: Left: transmission spectroscopy of HD 189733 b. Coloured lines show an atmospheric prediction. A comparison haze-free model is shown in grey [Pont et al., 2013]. Right: Spitzer measurements of the planet to star flux ratio for HD 209458 b, with predicted emission spectrum from Burrows et al. [2006]. Reproduced from Knutson et al. [2008].

given by

$$H = \frac{k_B T}{\mu_m g} \quad (1.18)$$

where T is the temperature, μ_m is the mean molecular mass, g is the surface gravity of the planet and k_B is Boltzmann's constant. The extra decrease in flux from just the atmosphere is given as

$$\Delta\delta_{tra} \approx 2N_H\delta_{tra} \left(\frac{H}{R_p} \right) \quad (1.19)$$

where N_H is the depth of the atmosphere in scale heights, typically of order unity [Seager, 2011].

Emission spectroscopy is a complementary technique where the light from the planet is directly detected and characterised. Just before the planet passes behind the star, the day side of the planet is visible, and when the planet is completely occulted by the star during the secondary eclipse only the stellar flux is observed. The difference in these provides a measure of the thermal emission and associated spectral features of the planet, and also the reflected emission especially in the optical. The secondary eclipse depth is typically very small ($\sim 10^{-3} - 10^{-4}$ for hot Jupiters depending on the wavelength). By assuming the star and planet are both blackbody radiators the secondary eclipse depth is derived from Eq. 1.4:

$$\delta_{occ}(\lambda) = \delta_{tra} \frac{B_\lambda(T_p)}{B_\lambda(T_\star)} \quad (1.20)$$

where $B_\lambda(T)$ is the Planck function (Eq. 1.15). By integrating Eq. 1.20 accounting for the spectral sensitivity of the observing instrument and the spectral energy distribution of the star, the brightness temperature of the planet is estimated.

Hot Jupiters are excellent candidates for atmospheric characterisation. The low surface gravity and high temperatures lead to a larger atmospheric scale height (Eq. 1.18) leading to a larger $\Delta\delta_{tra}$, and the short periods facilitates easier organisation of observations. The two most often studied planets for atmospheric work are HD 209458 b and HD 189733 b as they are both hot Jupiters with large atmospheric scale heights orbiting bright stars. Figure 1.9 shows two results from studying planetary atmospheres. The first (Fig. 1.9a) suggests that the transmission spectrum of HD 189733 b is dominated by Rayleigh scattering over the whole visible and near-infrared indicating a cloud with grain sizes increasing linearly with pressure and an opaque cloud deck [Pont et al., 2013]. A similar analysis was performed for HD 209458 b and found no such haze feature [Désert et al., 2008] suggesting that the two most studied hot Jupiters are remarkably different in their atmospheric composition. Figure 1.9b shows the emission spectrum for HD 209458 b. Knutson et al. [2008] infer that the atmosphere does not follow traditional models of hot Jupiter atmospheres [e.g. Seager et al., 2005], and that a temperature inversion layer high in the atmosphere is required to explain the observed excess emission at $5.8\text{ }\mu\text{m}$.

1.4 Projects

In this section some major planet detection projects are introduced. I discuss WASP, the most successful ground based transiting survey, Kepler the most successful space-based transiting survey and HARPS a likely follow up instrument for the Next Generation Transit Survey (NGTS, discussed in Chapters 4 and 5).

1.4.1 WASP

The Wide Angle Search for Planets (WASP) project is a UK led initiative to find transiting exoplanets from the ground suitable for spectroscopic confirmation. It consists of two sites, one in the Northern hemisphere on Roque de los Muchachos on La Palma, and one in the Southern hemisphere at the South African Astronomical Observatory (SAAO). These two sites cover declinations ranging from -90 degrees to 60 degrees, a huge fraction of the sky. Each site houses eight telescopes each with a field of view of 60 square degrees [Pollacco et al., 2006]. The South African instrument is shown in Fig. 1.10. Over 37 million stars have been observed with WASP and over 100 confirmed exoplanets have been discovered, though some are



Figure 1.10: A photograph of the WASP-South instrument. The site in the north houses an identical instrument. Reproduced from Pollacco et al. [2006], image credit: David Anderson.




Figure 1.11: Observations of HD 209458 b from the WASP0 prototype.

unpublished. The planets discovered orbit bright ($9 \leq V \leq 13$) solar-type stars which are observable using ground based spectroscopy and high precision photometry, validating the planets discovered.

A prototype instrument was installed on La Palma from June - August 2000 [Kane et al., 2004] and observed the Pegasus field. The known transiting planet HD 209458 b was observed to transit as shown in Fig. 1.11 proving that the instrument could detect transiting planets.

The final instruments are similar in design with f/1.8 200 mm telephoto lenses, 4 megapixel CCDs with $13.5 \mu\text{m}$ pixels, and a wide spectral response (400 nm - 700 nm) to maximise the light collected. Figure 1.12 shows the phase folded lightcurve of the first WASP planet, WASP-1 b. WASP is ideal for detecting hot Jupiters as the large sky area observed maximises the chances of finding these rare objects. The project has produced the lowest density planets detected from the ground (WASP-31 b and WASP-57 b at 0.132 g cm^{-3} and 0.12 g cm^{-3} respectively), the hottest planets detected from the ground (WASP-33 b and WASP-12 b at 2463 K and 2363 K respectively) and the largest planet at the time of discov-




Figure 1.12: Phase folded lightcurve of WASP-1 b. Crosses denote 2004 season WASP photometry, and filled circles show the Volunteer Observatory light curve of 2006 October 1 [Collier Cameron et al., 2007b].

ery (WASP-17 b at $1.93^{+0.052}_{-0.1} R_J$).³ Brown et al. [2011a] study WASP-18 b and WASP-19 b and constrain the stellar and planetary tidal quality factors Q'_s and Q'_p through Monte Carlo fitting and propose that WASP-19 b may have a remaining lifetime of $0.0067^{+1.1073}_{-0.0061}$ Gyr suggesting rapid infall. The large errorbars for this result are due to large uncertainties on the stellar age estimation.

The project excels at detecting hot Jupiters due to their large size and relatively deep transit depth, but the project's strengths are the large number of stellar targets observed and the brightness of the targets. WASP planets are often good candidates for atmospheric studies as they transit frequently, typically orbit bright host stars, and often have large atmospheric scale heights (WASP-17 b and WASP-39 b in particular). These hot Jupiters are rare [e.g. Howard et al., 2012] but relatively easy to detect. By observing a large stellar sample, the population of these unusual objects can be understood allowing WASP to lead the analysis of hot Jupiter populations, provided the selection effects can be understood (Chapters 2 and 3).

1.4.2 Kepler

The Kepler mission was designed to determine the frequency of Earth-sized planets in and near the habitable zone of Sun-like stars [Borucki et al., 2010]. The other scientific goals include determining the radius and semi-major axis distributions of the planets and to estimate and characterise the multi-planet systems [Borucki et al., 2009]. By observing from space the mission is not hindered by observing through the Earth's atmosphere allowing for much higher precision measurements of the stellar

³Data taken from <http://exoplanets.org>

flux and detecting smaller transit depths than is possible from the ground. The target precision was such that the transit of an Earth-sized planet around a 12th magnitude G2 star would be detected at 4σ [Borucki et al., 2010]. The spacecraft was launched on March 6 2009 into an Earth-trailing orbit continuously observing a sample of 150000 stars for transit events.

To date 136 confirmed or validated planets have been discovered, with over 3500 *candidate* planets awaiting validation. The planet candidate hosting stars that are typically observed with Kepler are too faint to perform radial velocity analysis of, so statistical vetting is used to argue the validity of the planet candidates detected. Planets are often validated through BLENDER estimation of the false positive chance [Torres et al., 2004, 2005]. Multi-planet systems are often confirmed through dynamical estimation of the masses in the system based on orbital solutions constrained by the observed transit timing variations (TTVs) where differences in the mid-points of the transits allows the presence of another massive body in the system to be inferred. This method requires co-planarity of the systems which places constraints on planetary migration methods.

The Kepler mission has found a wealth of interesting individual planets, from sub-Earth sized planets (e.g. Kepler-37 b, Kepler-62 c and Kepler-42 d) to the most dense planet to date Kepler-68 c. The extremely high precision of the photometric measurements allows the detection of very small planets, especially around the later type stars available in the Kepler field of view. Though the Kepler instrument was designed to search for small planets, the project has detected Jupiter class planets with exquisite photometric quality (e.g. Kepler-12 b [Fortney et al., 2011], Kepler-17 b [Désert et al., 2011]). Example Kepler lightcurves for a hot Jupiter and super Earth are shown in Fig 1.13. In the case of Kepler-17 b the starspots and high stellar rotation rate allow a limit to be placed on the orbital obliquity of $< 15^\circ$.

Kepler has shown that multi-planet systems are common and are in stable coplanar orbits [Lissauer et al., 2011]. The distribution of observed period ratios shows that the vast majority of candidate pairs are neither in or near low-order mean motion resonances, though a non-negligible sample are, especially near the 2:1 resonance. Resonant orbits of multi-planet systems are thought to be an indicator of smooth disk migration, and the co-planarity supports this.

Kepler has discovered planets in binary star systems, both wide binaries and close binaries where the planet orbits both stars. Kepler-16 b was discovered around an M1III detached binary. The orbital solution of the three body system allowed the mass to be determined at $0.33 \pm 0.02 M_J$ and the transit depth gave a planetary radius of $0.754^{+0.003}_{-0.002} R_J$. The majority of stars in the galaxy are in binaries, so this

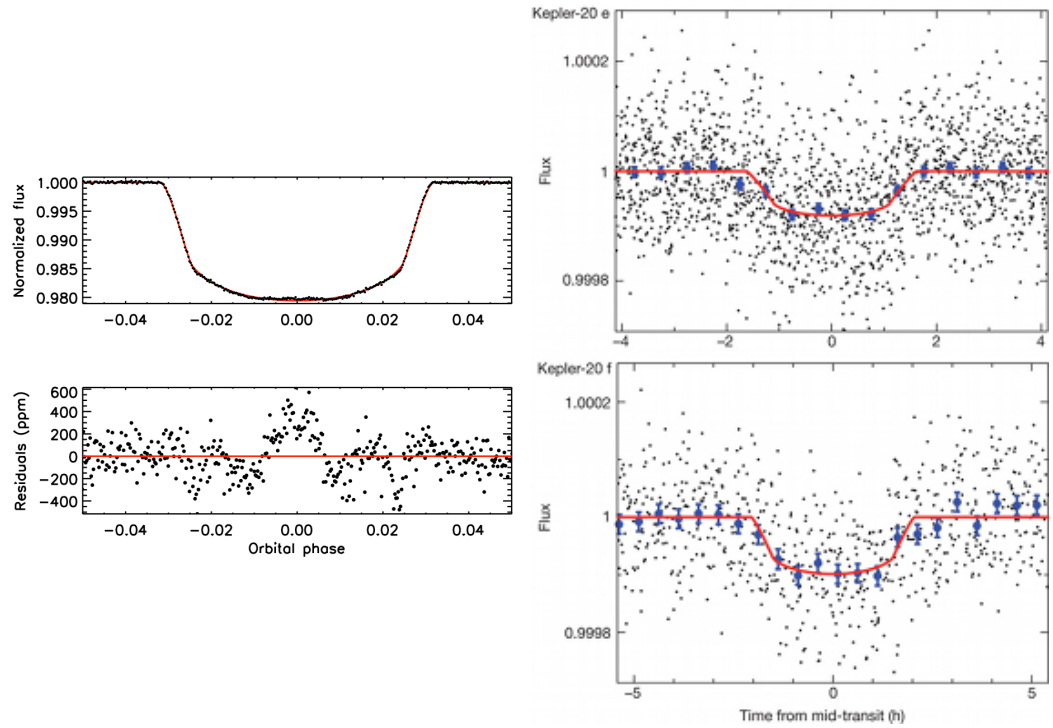


Figure 1.13: Left: Kepler lightcurve for Kepler-17 b. Upper: phase folded lightcurve. Lower: residuals to the best fit model, the systematic noise is due to stellar spots [Désert et al., 2011]. Right: Kepler lightcurve for (upper) Kepler-20 e and (lower) Kepler-20 f [Fressin et al., 2012].

discovery is encouraging as we do not need to rule out a large number of stars when considering abundance rates of planets.

Howard et al. [2012] use the Kepler *candidates* to constrain the underlying distribution of planets within 0.25 AU of their host stars. The Kepler input catalogue is restricted to a solar subset of relatively bright ($K_p < 15$)⁴ solar type main sequence GK stars ($4100 \leq T_{eff} \leq 6100\text{K}$ and $4.0 \leq \log g \leq 4.9$). By correcting for the detection efficiency of each star and the geometrical transit probability, the underlying occurrence rate is estimated. The underlying occurrence rate is estimated considering orbital period and planetary radius separately. The functional form of the period distribution consists of a power law with an exponential cutoff towards shorter periods, noting that the occurrence rate drops off at shorter periods. The radius distribution is fitted as a simple power law with index -1.92 ± 0.11 reflecting that larger planets are rarer. The overall occurrence rate of hot Jupiters around Solar type stars orbiting within 0.25 AU was calculated to be 0.165 ± 0.008 planets per star.

1.4.3 HARPS

The High Accuracy Radial velocity Planet Searcher (HARPS, Mayor et al. [2003]) is a high precision echelle spectrograph devoted to searching for exoplanets. Installed in 2003, it is mounted on the 3.6 m telescope at La Silla covering the spectral domain of 380 nm to 690 nm with a resolution of $R = 115000$ [Mayor et al., 2003]. To date HARPS has detected over 140 planets. The extremely high long term precision of 0.8 m s^{-1} [Dumusque et al., 2012] is provided by a vacuum chamber for the spectrograph, that provides high levels of atmospheric stabilisation with the pressure kept below 0.01 mbar and rms temperature variations of 0.001 K [Mayor et al., 2003]. This level of stability has enabled the detection of the lowest mass radial velocity planets to date: α Centauri B b with $M \sin i = 1.127 \pm 0.096 M_{\oplus}$ and GJ 581 e with $M \sin i = 1.95 \pm 0.22 M_{\oplus}$.

The closest star to the Solar System α Centauri is a binary star consisting of a G2V primary and K1V secondary, and the planet α Centauri B b orbits the secondary with a period of 3.24 days [Dumusque et al., 2012]. The star is quiet compared to other HARPS target stars, but the high precision measurements required to detect a low mass planet required accurate modelling of the stellar activity including stellar oscillations, granulation and starspots.

GJ 581 e is the fourth planet in a multi-planet system. Detected two years

⁴Kepler magnitude covering 400 nm - 900 nm, unique to the Kepler instrument but conversions to SDSS g' and r' filters are possible, <http://keplergo.arc.nasa.gov/CalibrationZeroPoint.shtml>

after the other members of the system the planet as its velocity semi-amplitude is smaller. The planet orbits on a 3.15 day orbit around an M3V star. The radial velocity technique is sensitive to the mass ratio of the objects in a system, so observing M dwarfs decreases the minimum mass detection limit. A dedicated M dwarf survey consisting of 300 stars including GJ 581 is being performed with HARPS, though observing M dwarfs presents extra challenges such as increased magnetic activity [Mayor et al., 2009]. Other programmes being undertaken by HARPS are a metal-poor stellar survey targeting low metallicity stars, a low mass planet search where the highest precision achievable is employed to search for low mass planets, and a volume limited search of 850 stars in the local neighbourhood.

1.5 Planetary formation

Planets form out of stellar material from the initial molecular cloud remaining after a star has formed. A star forms through gas collapse until the density is high enough to start nuclear fusion, so the outward pressure from the interior balances the gravitational collapse. Gas which has too much angular momentum to collapse onto the star forms an envelope around the star [Armitage, 2010]. The timescale for accretion onto the star is much longer than the orbital timescale of the envelope so material orbits around the star. Due to low disk temperatures and pressures, the gaseous envelope collapses under gravity to form a protoplanetary disk.

In this protoplanetary disk two major formation theories have been proposed for the formation of planets from the disk material: core accretion and gravitational collapse.

1.5.1 Core accretion

Dust particles in the protoplanetary disk form planetesimals through pairwise collisions. This is efficient for micron-sized grains with relative disk velocities of 1 m s^{-1} or slower [Barnes, 2010] where the collision velocities restrict further growth at around 1 m. In a rotating gas disk the massive bodies formed through pairwise collision are decoupled from the gas disk. Gas pressure causes the gas to rotate at sub-Keplerian velocities, slower than the material. The material suffers resistance to its motion from the drag of the slow moving gas, loses orbital energy to the disk and starts to spiral inwards towards the host star. Bodies larger than 10 m have enough inertia to resist the slowing force and so the maximum orbital decay occurs for 1 m objects, where the infall timescale of around 100 yr is much shorter than the typical growth timescales [Barnes, 2010]. This metre-sized problem requires extremely

rapid growth from 1 m to 10 m, possibly explained by local over-density in the disk allowing increased accretion rates [Armitage, 2010], though a complete explanation is uncertain [Perryman, 2011].

Once the rock is larger than 10 m further collisions cause continued growth to around 10 km in 10^4 to 10^5 years. An object 1 km and above is termed a planetesimal, loosely defined as an object whose internal strength is dominated by self-gravity and whose orbital dynamics are not affected by gas drag [Perryman, 2011]. Planetary growth continues to form terrestrial planets by further collision. The increasing mass causes the collisional cross section (extended by gravitational focussing) of the objects to increase much larger than their physical cross section which accelerates planetary growth further, until the *isolation mass* is reached and the planet has accreted all other bodies in its vicinity.

After this stage the planet undergoes oligarchic growth where the growth rate slows. Larger bodies tend to grow more slowly than smaller ones while most planetessimals remain small. The final stage of terrestrial planet formation involves giant impacts among the protoplanets, for example the formation of the Moon in the Solar System. Giant planets form through gas accretion whilst accreting planetessimals [Kokubo & Ida, 2002; Ida & Lin, 2004].

The temperature of water condensation in a typical low pressure protoplanetary disk is 150 K - 170 K [Armitage, 2010]. The temperature profile for an optically thin disk is given by

$$T = 2.8 \times 10^2 \left(\frac{a}{1\text{AU}} \right)^{-1/2} \left(\frac{L_\star}{L_\odot} \right)^{1/4} \quad (1.21)$$

where L_\star is the luminosity of the central star [Kokubo & Ida, 2002]. This gives a typical separation at the water condensation point of 2.7 AU for the solar system. Beyond this snow line water ice can be present and increases the surface density of solid material in the protoplanetary disk by a factor of 3-4 [Perryman, 2011] causing increased material accretion. This has the effect of forming planets at a higher rate than inside the snow line where water remains in its gaseous form.

Terrestrial planets typically form inside the snow line, as here the isolation mass is less than the mass required to form a gaseous envelope. Outside the snow line where the disk surface density is much higher the planetesimals can undergo the formation of a gaseous envelope. The envelope starts to form once the planet has enough mass to gravitationally capture the surrounding gas in the disk. The envelope and core evolve simultaneously; initially the core accretion is faster than the gas accretion until the solid material in the disk depletes, where the gas accretion

rate exceeds the solid accretion rate [Armitage, 2010]. Once the core mass and envelope mass are equal, a runaway phase of gas accretion occurs where the accretion rate is no longer limited by the cooling rate of the envelope, but is limited by the hydrodynamic interaction between planet and disk [Armitage, 2010], a much faster rate. Once no more gas is present in the local region of the disk around the planet, the gas accretion stops and the planet cools and contracts.

1.5.2 Gravitational collapse

Instead of forming planets by growing material from dust grains to planets, the gravitational collapse model proposes local regions of high density in a high mass protoplanetary disk will gravitationally collapse into planetesimals [e.g. Baruteau et al., 2011]. This process is much quicker than the planetesimal formation through core accretion and so would side-step the lengthier formation timescale demanded by core accretion [Perryman, 2011]. It is also a proposed explanation of how planets can exist at extremely large orbital separations (> 10 AU), though core accretion can explain this also [Baruteau et al., 2011].

Some protoplanetary disks are likely to be gravitationally unstable during their evolution. The self gravity in the disk alters its structure and evolution creating density perturbations, which grow and develop into spiral arms transferring angular momentum outwards and mass inwards. Under some circumstances the disk may break up into self-bound clumps and further collapse directly into giant planets [Perryman, 2011]. Gravitational instability occurs when the Toomre parameter $Q \geq 1$ where

$$Q = \frac{c_s \Omega}{\pi G \Sigma}, \quad (1.22)$$

c_s is the sound speed in the disk, Ω is the angular velocity and Σ the surface density. For a disk with aspect ratio $h/r = 0.05$ at 10 AU around a $1 M_\odot$ ($M_\odot = 1.989 \times 10^{30}$ kg) star the surface density needs to be $\Sigma \sim 10^4 \text{ kg m}^{-2}$, on order twice the density of the minimum disk mass required to build the solar system planets (the minimum mass solar nebula) at the same radial distance, suggesting that the mechanism operates preferentially at early epochs where the disk mass is high [Barnes, 2010]. Figure 1.14 shows some hydrodynamical disk simulations with a range of disk parameters, where clumps of local density can be seen, especially in the lower right image. Mayer et al. [2007] find a correlation between the chances of fragmentation of the disk increase with mean molecular weight, and that only disks with a metallicity comparable with solar or higher can fragment. Dodson-Robinson et al. [2009] find that the most likely formation channel for wide giant planets (such




Figure 1.14: Protoplanetary disk simulations showing logarithmic surface density after 1300 years of evolution. The panels represent different mean molecular weights and simulated particle orientations. From Mayer et al. [2007].

as the HR 8799 system) is gravitational collapse, neither dynamical scattering to wide orbits or outward migration can explain such distant massive systems.

1.6 Migration

For gas giant formation both methods described above require large quantities of gas along with temperatures low enough for ices to condense and form heavier cores. The observed hot Jupiter population and other short period planets did not form at their current separations, as the temperature in the protoplanetary disk would have been too high for planet formation [Rice et al., 2012]. Two main methods for planetary migration have been proposed: disk migration and dynamical scattering.

1.6.1 Disc migration

A disk of material forms around the protostar out of which the planets are formed (Section 1.5). The gas in the disk interacts with the forming planets by exchanging angular momentum, changing the orbital separations of the star and planet. The planet moving through the disk excites spiral density waves, which in turn exert torques on the planet [Perryman, 2011].

Two classes of resonant location exist: co-rotation resonance, where the angular frequency Ω_p of the planet is equal to the angular frequency of the rotating gas at radius r

$$\Omega_p = \Omega(r) \quad (1.23)$$

co-orbital with the planet, and Lindblad resonances at locations

$$r_L = \left(1 \pm \frac{1}{m}\right)^{2/3} a \quad (1.24)$$

where a is the separation of the planet, and m is an integer [Armitage, 2010]. The gas inside the orbit of the planet increases the angular momentum of the planet and causes migration outwards. Conversely gas inside the orbit of the planet decreases the angular momentum of the planet and causes migration inwards. The ratio of the strength of the respective torques from gas at resonant locations determines whether the planet migrates inward or outward: torques exerted by the Lindblad resonances interior to the planet exert a positive torque, and Lindblad resonances exterior to the planet exert a negative torque [Goldreich & Tremaine, 1979]. The typical timescale for inward migration of an Earth sized planet at 5 AU is $1 - 10 \times 10^5$ years, and the timescale for migration decreases for more massive planets [Tanaka et al., 2002] which is at odds with the required planet formation timescales of 10^7 years, calculated from the measured anti-correlation of infra-red excesses with stellar age in star-forming regions. The migration of the planet may decrease the required formation timescale, as it prevents the severe depletion of the feeding zone of the planet for accretion providing a higher mass accretion rate than calculated in *in situ* calculations [Alibert et al., 2005].

Disk migration is generally split into two classes: Type I and Type II, characterised by whether a gap is formed in the disk around the planet. In the Type I regime the planet is not massive enough to open a gap leaving the disk surface density approximately unperturbed. Material is present at all resonant locations, and the outer Lindblad usually dominate due to a steep radial pressure gradient of the disk causing the planet to migrate inwards [Armitage, 2010]. Type II migration occurs when the planet's Hill sphere (the region of space around which satellites are bound to a massive object)

$$R_H = \left(\frac{m_p}{3M_\star}\right)^{1/3} a \quad (1.25)$$

is larger than the scale height of the disk h , leading to the condition for a gap to




Figure 1.15: Two-dimensional hydrodynamic simulations depicting the interaction between a planet and viscous protoplanetary disk. Left panel: example of Type I migration where the interaction is weak enough that the local surface density of the disk is effectively unperturbed. Right panel: example of Type II migration where a gap is opened up in the disk in the vicinity of the planet. Reproduced from Armitage [2010].

open

$$\frac{m_p}{M_\star} \geq 3 \left(\frac{h}{a} \right)^3. \quad (1.26)$$

For typical protoplanetary disks the aspect ratio $h/a \approx 0.05$ setting the mass ratio limit between Type I and Type II migration at $m_p/M_\star \sim 4 \times 10^{-4}$, i.e. for planets between the mass of Saturn and Jupiter around a solar type star [Perryman, 2011]. Figure 1.15 shows example simulations of disk under Type I and Type II migration regimes, showing a clear gap present in the Type II case.

Lin et al. [1996] theorise that disk migration is the cause of the small separation of 51 Peg b, where the migration stopped due to tidal interactions with the host star, or truncation of the inner circumstellar disk by the star’s magnetosphere. Rice et al. [2012] find a current observed population of exoplanets consistent with simulated populations of exoplanets which underwent Type II migration followed by a pile-up due to a truncation at the inner disk, providing evidence that Type II migration is a viable option for explaining the observed population of exoplanets.

It is believed that disk migration is too efficient at migrating planets, as the timescale for terrestrial planet migration is significantly shorter than the formation timescales of gas giants and that of typical protoplanetary disks, which would lead to a significant loss of planets into the star [Perryman, 2011]. A slowing effect, possibly due to a steep surface density in the inner disk [Morbidelli et al., 2008] or non-isothermal disks [Kley & Crida, 2008] may be required.

1.6.2 Dynamical scattering

Dynamical scattering is the inward migration of planetary systems through gravitational interactions with planetessimals or other forming planets. Giant planets form beyond the snow line a few AU from their host stars in approximately circular orbits. Collisions or gravitational interactions with other planets in the system alter the orbital paths of the affected members, and may induce changes in separation, eccentricity and inclination. The primary migration path for such systems is a multi-planet system interacting gravitationally, which excites the eccentricity and inclination of the planets, causing at least one planet to obtain a near hyperbolic orbit around the star. Often planets are thought to be ejected from the systems contributing to the population of free floating planets [e.g. Lucas & Roche, 2000], or interact with the central star causing tidal disruption of the planet. The wide range of orbital separations of the known exoplanets can be explained by disk migration, but the large range in observed eccentricities is still a matter of debate [Barnes, 2010].

A proposed method of exciting the forming planets into highly eccentric orbits is through the Kozai mechanism, where a mutually inclined third body at large orbital separations causes oscillations in the eccentricity and inclination of the inner planet, shrinking the periastron distance. Angular momentum exchange between the two orbiting bodies causes the eccentricity and orbital inclination to oscillate out of phase with each other i.e. increasing one at the expense of the other, whilst conserving the integral of motion (the Delaunay quantity)

$$H_K = (1 - e^2)^{1/2} \cos i \quad (1.27)$$

[Kozai, 1962; Perryman, 2011]. To excite these oscillations the only requirement is the mutual inclination between the star-planet system and the third outer body must be above a critical value

$$i_c = \arcsin \sqrt{2/5} = 39.23 \text{ degrees} \quad (1.28)$$

Innanen et al. [1997]. Originally formalised for inclined asteroids under the gravitational influence of Jupiter this process has been applied to binary stars and subsequently exoplanets [e.g. Wu et al., 2007; Triaud et al., 2010; Plavchan & Bilinski, 2013] to explain the initial induced eccentricity, and subsequent measurements of misaligned planets.

Many planets exhibit completely circular orbits, especially the ones nearest

their host stars, so a recircularisation process must occur to reduce the eccentricity of these orbits to match the observed population. The recircularisation of these high eccentricity planets that remain bound to the system is performed by the planet inducing tides on the planet converting the planetary orbital angular momentum into tidal heating of the planet. The gravitational pull from the planet also raises tides in the star. The net torques from the planet transfer angular momentum from the star to the planet causing the stellar rotation and planetary orbit to synchronise. These torques also align the rotation axes of the star and planet, and circularise the planetary orbit [Perryman, 2011]. Typically the stellar rotation period is longer than the planetary orbital period so the tides slow the orbit of the planet causing gradual in-spiral [Barnes, 2010]. The tidal effects are strongest during the periastron passage, where the eccentric orbit comes closest to the star, and each successive orbit decreases the eccentricity and circularises the orbit [Socrates et al., 2012]. The tidal interactions can be quantified by the tidal dissipation parameter Q , the ratio of the available energy to amount dissipated by frictional losses during each orbital cycle [e.g. Hellier et al., 2009].

1.7 Competing theories

Evidence for both aspects of the planetary formation and migration exist. It is commonly thought that planets within 10 AU formed through core accretion and migrated inwards, but distant giant planets currently detectable only through direct imaging techniques were formed through gravitational instability as the timescale for core accretion at those distances is significantly longer further away from the parent star. An alternative proposal follows the expansion of the orbits of Neptune and Uranus during the solar system formation out to similar distances which allows core accretion to remain the primary formation mechanism for such distant objects [Thommes et al., 2002; Gomes et al., 2005]. This hypothesis is currently favoured for Neptune as it proposes an explanation for the populations of small objects in the Kuiper Belt.

Ford & Rasio [2006] hypothesise that if dynamical migration coupled with tidal circularisation were the dominant effect, a cutoff would be apparent in the current population with planets orbiting at twice their Roche limit

$$a_R = \frac{R_p}{0.462\mu^{1/3}} \quad (1.29)$$

where $\mu = m_p/M_\star$ from conservation of angular momentum and an initially eccentric orbit. Conversely if smooth disk migration were the dominant migration path then this would cause an inner edge *at* the Roche limit. Evidence for such an occurrence was not found by [Husnoo et al., 2012] who propose post-migration evolution of the planet may explain the discrepancy (see also Baraffe et al. [2004]; Pont et al. [2011]). By considering the excess rotation of stars hosting exoplanets, Pont [2009] find evidence for well-defined limits in parameter space beyond which the orbits of planets are circularized, and beyond which the host star shows sign of excess spin due to tidal effects. Observed eccentric orbits and fast rotating host stars for transiting planets are compatible with tidal evolution [Pont, 2009]. Triaud et al. [2010] propose that most hot Jupiters are misaligned, concluding that observations and predictions using the Kozai mechanism match well implying that the main hot Jupiter formation method is through dynamical scattering.

By proposing initial distributions of planets consistent with Type II migration and evolving them through tidal interactions with their host stars, Rice et al. [2012] find that the resulting population is consistent with the observed population. A modest pile up at $a \sim 0.05$ AU and a peak eccentricity distribution at $e = 0$ was required, along with a stopping mechanism where the migration timescale increases greatly inside the 2:1 resonance with the inner disk edge. They acknowledge that dynamical scattering must occur to form misaligned systems, and that this preferentially occurs for more massive stars ($M_\star > 1.25 M_\odot$) where the disk lifetimes are too short for Type II migration to be effective.

By studying the obliquity of a selection of planetary systems, a correlation between obliquity and expected tidal timescale was found, where low obliquity (well aligned) systems have a low expected tidal timescale, and high obliquity systems have a high expected timescale which points to a dynamical migration method [Albrecht et al., 2012]. Wright et al. [2009] find that multiple planet systems tend to exhibit smaller eccentricities and a more even distribution in orbital period than the single planet systems, suggesting that the single planet systems could be the remnants of the aftermath of dynamical scattering where we only see the planets remaining in the system after their migration phase.

It is possible that the two migration mechanisms are acting in parallel. Matsuura et al. [2010] simulate an N-body system with dissipating gas disk and find that the distribution of orbital separations is largely determined by the gas disk, while the eccentricity distribution is determined after the disk dissipation. Dawson & Murray-Clay [2013] find with 99.1% confidence that giant planets orbiting between 0.1 and 1 AU orbiting metal poor stars ($[\text{Fe}/\text{H}] < 0$) exhibit lower eccen-




Figure 1.16: Distribution of known planetary orbital periods, for planets detected through the transit method (red) and radial velocity (blue). The black line indicates the combined distribution and contains 715 objects. Data from <http://exoplanets.org>.

tricities than those orbiting metal rich stars. It is possible that only high metallicity disks have enough disk mass to form multiple giant planets which dynamically scatter, whereas low metallicity disks cannot form as many large planets so undergo smooth disk migration [Dawson, 2013].

There is still a lot of uncertainty in determining the dominant formation and migration mechanisms. We use the observed population of exoplanets to constrain the possible outcomes of planet evolution theories, but selection biases must be considered.

1.8 Overall properties of exoplanet populations

With a sample of exoplanets the statistical properties of the ensemble can be made, providing information on the bulk properties of the *population*. Figure 1.1 shows the mass and separation of the known exoplanets. Three distinct classes are observed: the massive but distant objects, from the mass of Saturn up to massive objects close to the deuterium burning limit of $13M_J$; the hot Jupiters orbiting closely to their host stars, and which are massive; and the warm Neptunes at small separations and low masses.

Care should be taken when drawing conclusions from the detected population of exoplanets as clear biases exist for the detection methods used. Figure 1.16 shows the distribution of orbital periods for the known exoplanets. A clear bimodal distribution is seen with a strong peak at around 4 days, and another wider peak




Figure 1.17: Left: the orbital eccentricity of the known exoplanets for which a measurement exists. The grey line marks where the periastron distance equals the Roche limit for a 1 R_J, 1 M_J planet around a 1 M_☉ star. Right: metallicity distribution for the total (black), radial velocity (blue) and transiting (red) exoplanets. Data from <http://exoplanets.org>.

at around 600 days which roughly correspond to the separate groups in Fig. 1.1 but accounting for the stellar mass. The pile up in orbital period at around 4 days has been noted [e.g. Fressin et al., 2007; Cumming et al., 2008; Wu & Lithwick, 2011; Hellier et al., 2012], and was first observed in radial velocity studies, and still appears in the blue distribution of Fig. 1.16. The different detection methods match these two peaks well: the radial velocity peak dominates the long period objects, whereas the short period peak is more balanced but mostly defined by the transiting planets. The range of observed periods encompasses the periods of the solar system planets, but extends much shorter than the orbital period of Mercury.

Figure 1.17a shows the distribution of orbital eccentricities for the known exoplanets and plots this against the distance from the host star. Planets exist at almost all possible allowed eccentricities which places constraints on proposed formation mechanisms. Figure 1.17b shows the distribution of stellar metallicity for the known planets. The radial velocity planets show a bias towards high metallicities which was suggested to be a selection bias [Boss, 2002] but it is now thought that planets are more likely to form around high metallicity stars [Santos et al., 2003; Fischer & Valenti, 2005] due to a higher abundance of material that can grow into planetary cores.

Figure 1.18a shows stellar evolution models extended into the planetary regime. The relationship is linear for stars with $R \propto M$ as the gravitational collapse of material onto the core is balanced by outward gas and thermal pressure from the

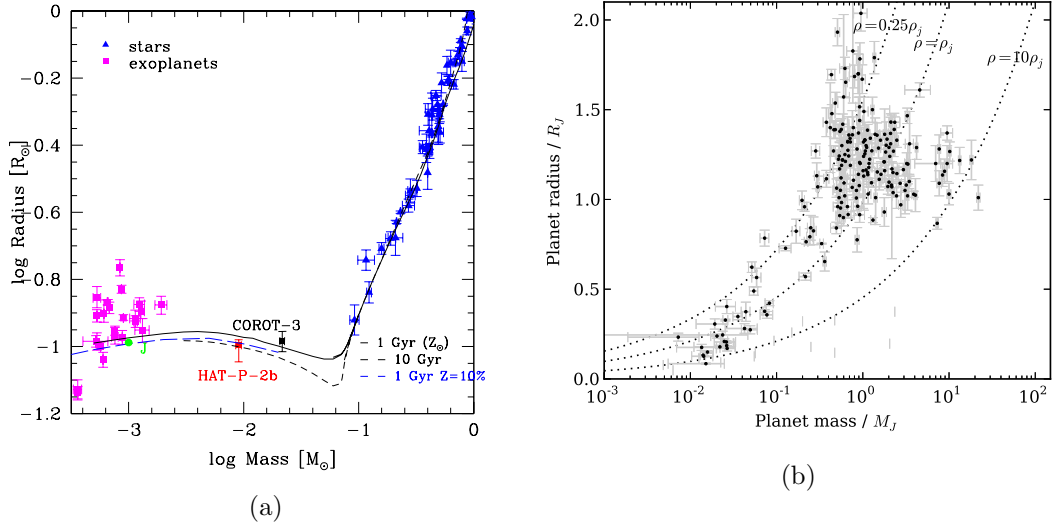


Figure 1.18: Left: mass radius relation for gas giants (pink) and stars (blue) showing stellar composition models for two isochrones (black) and 10 % mass fraction of heavy elements (blue) [Chabrier et al., 2009]. Right: mass radius relation for just the planets including terrestrial planets. Dashed lines represent lines of constant density.

heat of nuclear fusion. In the brown dwarf regime ($0.012 \lesssim M \lesssim 0.06 M_{\odot}$) the mass is not sufficient to start nuclear fusion through hydrogen burning, and gravitational collapse is instead balanced by electron degeneracy pressure. As mass decreases further towards the planetary regime, the gravitational collapse of material is further supported by electrostatic Coulomb pressure and the radius increases, towards $M \approx 0.004 M_{\odot}$ where the forces are comparable and the radius is independent of mass. The electrostatic force dominates the giant planet mass-radius relation until the terrestrial planet regime where the radius depends on the compressibility of material in the core and the stratified layers above. The gas giants ($M \sim 10^{-3} M_{\odot}$) show significant deviation from the predictions, where a wide range of planetary densities are required to explain the distribution of planetary parameters (Fig. 1.18b). Inflated planets such as HAT-P-32 b or WASP-57 b are much larger than expected through an external heat source (for example stellar irradiation [e.g. Fortney et al., 2006] or tidal recircularisation [e.g. Miller et al., 2009]), but are relatively easy to detect so they are likely selected for and so are most likely rare.

1.9 Selection biases

Each detection technique, and each project applies biases to the sample of exoplanets detected. When drawing conclusions regarding the underlying population based on the detected sample, these biases must be corrected for. The observed sample represents a subset of the true population, which may not show the same behaviour. Typically the biases are due to detection limits where the sample of detected planets is composed of objects which cause a relatively large observable signal. As technology and methods advance these biases are shifted, but still require thought.

Radial velocity surveys are biased towards high mass planets and shorter periods (see Eq. 1.2, with $P^2 \propto a^3$ from Kepler's law $\sigma_{detect} \propto m_p/\sqrt{a}$). This expected behaviour is clearly visible in Fig. 1.1: planets orbiting with longer periods need to be more massive to produce the same detection signal. We are therefore biased against detecting the less massive planets at longer periods. The transit technique is biased towards large transit depths, either due to large planets or smaller stars (see Eq. 1.8). Similarly at least three transits are required to constrain an orbital period so surveys are very strongly biased towards shorter periods. Real objects which have a transit depth too small to detect with high confidence are missed, causing the observed population to over emphasise the number of larger planets. Direct imaging surveys are biased against short orbital separations as the flux from the planet is negligible compared to the flux from the star. They are also biased towards young and hot planets and massive planets. These biases can lead to incorrect conclusions being drawn about the sample, therefore understanding the selection biases in a survey is required to study the population.

1.10 Instrument technology

1.10.1 Charge Coupled Devices

Charge coupled devices (CCDs) are a means to collect photons and convert them to a digital signal. Their invention revolutionised astronomy and many other fields as their operation is linear with the incoming flux. Before CCDs astronomy was performed with photographic plates which are non-linear and have a much lower sensitivity, requiring longer exposure times. Each incoming photon received causes the excitation of N_{QE} electrons, a wavelength-dependent quantum efficiency factor (QE) of the CCD. Typically modern CCDs can achieve QE values of $> 90\%$ [Howell, 2006].

A CCD is made up of pixels each consisting of a potential well to hold the




Figure 1.19: Analogy of a CCD. From Janesick & Blouke [1987].

charge caused by each photon. Figure 1.19 shows a representation of the charge collected as drops of water. Each pixel is shown as a bucket. During the readout phase of the exposure each row is shifted with a parallel shift register, and the last row of the CCD is shifted serially to the output amplifier. The output amplifier converts the analogue voltage recorded by each pixel to a digital representation, often with 16-bit conversion giving $2^{16} - 1 = 65535$ distinct flux levels available.

1.10.2 Calibration procedure

Each exposure contains electrons from the target of interest, but also sources which are not important to the measurement. The CCD cannot record a flux outside of the range 0 - 65535, so to record higher count rates the number of electrons is scaled by an inverse gain factor G commonly referred to as the gain

$$f_{ADU} = G f_{e^-} \quad (1.30)$$

where f_{e^-} is the received flux in electrons, and f_{ADU} is the flux stored in the units of Analogue-to-Digital Units (ADU). The gain has the units of $e^- \text{ ADU}^{-1}$. To prevent a negative value being recorded a bias voltage is applied in the CCD, corresponding to a bias offset. This bias level is removed before the analysis of every frame by reading out the CCD with an exposure time of 0 seconds and subtracting the recorded frame level from each image.

The quantum nature of electrons in the CCD cause spontaneous excitations to occur even in complete darkness, caused by heat in the silicon and is therefore temperature dependent. This dark current from thermal excitations is characterised by taking dark frames, exposing without opening the shutter of the camera. The

dark current of modern CCDs is typically very low, on the order of $1 \text{ e}^- \text{ s}^{-1} \text{ K}^{-1}$ so long exposures are used and the number of electrons produced is scaled to the exposure time used for the science images.

Flat frames are used to correct for the different sensitivities between pixels, and features blocking the optical path such as vignetting from the telescope or secondary mirror, or dust on the lens. A flat frame is taken by illuminating the CCD with uniform light either at twilight, or using reflected light from a lamp inside the dome the telescope is housed in. This frame is used to scale the per-pixel sensitivity for each science image.

1.10.3 Estimating the brightness of an object

To measure the flux from an object the number of photons collected which originated from that object are counted. This is often achieved by placing a circular aperture around an object and recording the pixel values within the aperture, in a technique known as aperture photometry. This aperture is often circular due to the radial symmetry of the point source profile, providing that photometry of stars is required. Photometry of galaxies for example can be performed with aperture photometry, but usually an elliptical aperture matching the shape of the half-light radius of the object in question. Point sources such as stars are usually measured with circular apertures.

Aperture weights are often calculated with a functional form such as a Fermi-Dirac like function [Eq. 1.31, Pollacco et al., 2006, in the WASP project for example], or other weighting functions which provide a similar profile. The goal is for a flat topped function with a smooth transition from inside the aperture to outside, as the star's flux is contained within an area defined by the instrument PSF

$$w(r) = \frac{1}{e^{r/A} + 1} \quad (1.31)$$

where r is the pixel distance radially from the centre of the aperture, A is a scaling factor, tuned to the project and $w(r)$ is the pixel weight at distance r . The weights are normalised so that $\sum_i w_i = \pi R^2$ where R is the desired aperture radius.

Contained in the aperture is not only the source flux from the object in question, but the sky background level. This sky background must be subtracted from the total flux measured to estimate the flux solely from the target object. A common method of estimating the sky background is to place an annulus around the aperture (see Fig. 1.20 for a schematic). The pixel values are summed in the sky annulus, usually involving some outlier rejection especially if stars are present




Figure 1.20: Geometry of a photometric aperture with its annulus.

in the annulus to calculate the sky flux level. This value is scaled by the area in the sky annulus $A_{sky} = \pi R (S_{out}^2 - S_{in}^2)$ where $S_{in,out}$ are scaling factors used to define the annulus size. By scaling by the annulus size we have a robust estimate of the local sky background around the target star. As the uncertainty of the average sky background per pixel scales as σ/\sqrt{N} with N the number of pixels, the larger the annulus the less noise due to the sky level uncertainty is introduced. The disadvantage to increasing the sky annulus is that other objects in the field will be more common in the sky annulus effectively reducing the number of pixels used for the sky background estimation, as pixels with contaminant stellar flux would be rejected as outliers. Another consideration is the inner scaling value S_{in} , which needs to be close to the target star to sample the local background, but at enough distance to include negligible flux from the source itself.

Choosing the optimal aperture size for the target object in question is important, ideally we want to collect every photon, but due the Gaussian-like nature of a typical PSF we would require a very large aperture indeed to incorporate the wings of the distribution. A common value for the size of the aperture is $1.5 \times \text{FWHM}$ as this includes almost all of the flux (missing $6 \times 10^{-5}\%$).

1.10.4 Sources of uncertainty

A photometric pipeline has a series of uncertainties from a range of sources. Observing astrophysical objects and collecting photons obeys Poisson statistics. The Poisson distribution has some key features: given the expectation value of $E(\text{Poisson}) = \lambda$, the variance is also λ so the standard deviation $\sigma = \sqrt{\lambda}$. The precision of a photometric measurement is given by the signal to noise ratio (S/N), the higher the S/N

the more significant the signal is compared to the background noise. For a CCD the S/N is given by [Howell, 2006]:

$$\frac{S}{N} = \frac{N_*}{\sqrt{N_* + n_{pix} (N_S + N_D + N_R^2)}} \quad (1.32)$$

unofficially named the CCD Equation, where n_{pix} is the number of pixels in the source aperture, and the sources of noise listed above are explained in the following subsections. The number of photons received from a particular source is denoted N_i where i is $*$ for the signal, S for the sky background, D for the dark current and R for the read noise.

Signal N_* Photons from the object of interest (signal) follow Poisson statistics, so correspondingly have a \sqrt{N} noise characteristic. For standard aperture photometry this is the total flux originating from the object inside the aperture, after sky background subtraction.

Sky background N_S Photons that originate from the background flux, generally the sky for ground based observations, contribute unwanted signal into the source aperture. The sky background is estimated through the sky annulus and is removed from the total source counts N_* , but the uncertainty on this level is $\sqrt{N_S}$ which remains in the image as a noise level. In Eq. 1.32 N_S is the number of photons received per pixel in the sky aperture. This is then scaled up by the number of pixels in the source aperture n_{pix} to calculate the number of photons in total from the sky.

Dark current N_D Even when the CCD is not collecting light electrons are spontaneously generated, and contribute a background level to each image [Martinez & Klotz, 1998]. The dark current level is completely reproducible and depends on the CCD temperature and exposure time. It affects each pixel independently and so must be scaled by the number of pixels n_{pix} . This dark current level N_D is removed through dark frames, but the uncertainty on the dark current estimation remains as a source of noise for each exposure $\sqrt{N_D}$.

Readout N_R Each readout of the CCD the number of electrons contained within each pixel is estimated with the electronics in the camera. Due to losses during charge transfer, and noise in the output electronics this estimation has an associated uncertainty and introduces another source of noise. The readout noise is measured

from bias frames, by exposing for zero seconds the readout stages of the CCD are characterised. The bias level added by the camera is subtracted from science images, and the uncertainty in the bias level is the readout noise. As this is characterised as a direct measurement of the noise and therefore non-Poissonian, it is squared when added in quadrature with the other terms.

Combining noise sources For the signal to noise calculation, the noise sources are added in quadrature so that

$$\sigma_T = \sqrt{\sum_i \sigma_i^2} \quad (1.33)$$

where σ_i are the individual sources of noise. As discussed above, given the number of electrons N_i , the noise sources are in turn: signal noise $\sqrt{N_*}$, background noise $\sqrt{n_{pix}N_S}$, dark current noise $\sqrt{n_{pix}N_D}$ and read noise $n_{pix}N_R$. The total noise is therefore given as

$$\begin{aligned} \sigma_T &= \sqrt{\sigma_*^2 + \sigma_S^2 + \sigma_D^2 + \sigma_R^2} \\ &= \sqrt{N_* + n_{pix}(N_S + N_D + N_R^2)} \end{aligned}$$

1.11 Thesis structure

In this thesis I study the selection biases in the WASP survey in order to determine the underlying population of hot Jupiters and constrain the evolutionary process acting on them (Chapters 2 and 3). I also discuss the Next Generation Transiting Survey (NGTS) which is a new transit survey designed to find smaller planets around bright stars. I discuss my role during the development of the project and analysis of two prototype instruments along with optimisation of the observing strategy and modelling of the noise processes (Chapters 4 and 5).

Chapters 2 and 3 use tools and data from the WASP project [Pollacco et al., 2006; Collier Cameron et al., 2006, 2007a], with extra tools developed, and analysis performed by myself. Chapters 4 and 5 involve a new transiting project for which the design was complete prior to my analysis. Analysis of the instrument capabilities was performed through standard techniques but applied to a new instrument.

Chapter 2

Quantifying the WASP selection effects

A major goal of exoplanet surveys is to determine the underlying population of planets in the Galaxy. With a significant statistical sample of planets, correlations can be analysed, trends found and physical understanding of the formation processes can be uncovered. The absolute occurrence rate of planets provides information on the efficiency of planet formation, and the underlying distribution of certain parameters contributes to the understanding of how and where planets are most likely to form. The occurrence rate requires a detailed knowledge of statistical and observational biases, and selection effects in a survey or groups of surveys.


2.1 Motivation

Hot Jupiter type planets are on order the size of Jupiter, orbiting their parent star within 0.1 AU, well inside the orbit of Mercury. The first hot Jupiter caused a flurry of activity from planet formation theorists to explain how such a large planet could get to such a close proximity to its host star. These planets were among the first to be detected due to the relatively large signal they create, but since more sensitive technology has become available they were found to be rare relative to smaller or more distant objects [Howard et al., 2012]. Hot Jupiters must have formed beyond the snow line where the necessary mass accretion rates can exist (see Section 1.5), but yet they appear much closer to their parent stars than any planets in the Solar System. Type-II migration is one proposed migration mechanism, where a giant planet forms beyond the snow line and migrates inwards due to accretion from the protoplanetary gas disc. This mechanism accounts for the observed pile-up of close



planets at around 3 days, and planetary pairs in mean motion resonances but cannot account for large eccentricities or observed Rossiter-McLaughlin results [Nagasawa et al., 2008]. The Rossiter-McLaughlin effect can however be explained with planet-planet scattering or the Kozai mechanism in multiple planet systems, but ejection is possible leading to single planet systems or differing mutual inclinations [Kozai, 1962]. The pile up of planets may alternatively be explained by truncation of the inner edge of the protoplanetary disk, under the influence of the stellar magnetosphere [Rice et al., 2012].

The WASP survey [Pollacco et al., 2006] is the most successful ground based search for transiting exoplanets to date with over 70 confirmed planets discovered since 2004. Over the course of the project WASP has observed over 37 million stars from two sites at Roque de los Muchachos Observatory on La Palma, and the South African Astronomical Observatory (SAAO). The two sites cover a large range of declinations, between $-90 \leq \delta \leq 60$ increasing the number of these rare hot Jupiters. Kepler is another successful project observing from space but with a single 130000 star field. As of 30/9/2013, 123 confirmed planets have been detected of which 12 are gas giants, with many Kepler object of interest (KOIs) remaining unconfirmed using established confirmation methods as they are too faint (see Section 1.4.2). The Kepler team have calculated that nearly 90 % of the KOIs have a false positive probability (FPP) $< 10\%$ and over half have FPP $< 5\%$ [Morton & Johnson, 2011]. More recently Fressin et al. [2013] find the FPP is globally 9.4 % but higher for giant planets ($6 - 22R_{\oplus}$) at 17.7 %. The false positive rate for both Morton & Johnson [2011] and Fressin et al. [2013] increase towards larger transit depths and by extension larger planets, due to the assumed underlying planet radius function decreasing while the false positive rate of hierarchical triples remains relatively flat.

Figure 2.1a indicates that WASP is proficient at detecting large close in objects whereas other projects may be able to probe the smaller or wider separation systems. Figure 2.1b shows the distribution of orbital period and planetary radius for the detected WASP planets. The large number of hot Jupiters detected by WASP gives the chance to draw more robust conclusions about the underlying distribution of hot Jupiters than with any other project. This is due to the limiting noise of the instrument preventing the detection of shallow transits, and the limited phase coverage biasing the sample to short periods and therefore small separations. The level of noise for Kepler is orders of magnitude lower than WASP, and the phase coverage is much higher which improves on both selection effects mentioned, however the sample size of hot Jupiters is lower. The number of detected planets and their distribution of parameters is heavily weighted by such selection effects, and without



(a)



(b)

Figure 2.1: Figure 2.1a: the current distribution of known exoplanets. Objects detected by WASP are marked in red and other objects marked with open circles. Figure 2.1b: distributions for the orbital period (top) and planetary radius (bottom) for the confirmed WASP planets. Data source: <http://exoplanets.org>.

Table 2.1: Literature hot Jupiter occurrence rates. The upper values are from transiting surveys, the lower values are from radial velocity surveys. Reproduced from Wright et al. [2012].

Work	Rate (per thousand)	Sample
Gould et al. [2006]	$3.1^{+4.3}_{-1.8}$	OGLE-III Transits (90 % confidence limits)
Howard et al. [2012]	5 ± 1	Kepler transits
Marcy et al. [2005]	12 ± 1	Keck, Lick, and AAT RVs
Cumming et al. [2008]	15 ± 6	Keck RVs (entire target list)
Mayor et al. [2011]	8.9 ± 3.6	HARPS and CORALIE RVs
Wright et al. [2012]	12.0 ± 3.8	Keck and Lick RVs

understanding and properly characterising them, we cannot make any meaningful claims as to the underlying distribution of planets in our galaxy. For example Penev et al. [2012] correct for *a priori* sensitivity assumptions to statistically constrain the tidal dissipation in planet hosting stars, which relies on the underlying distribution of planets.

Table 2.1 shows the occurrence rates of giant planets as calculated by other surveys (see [Wright et al., 2012] for limits on the parameter spaces searched.) The discrepancies between the radial velocity surveys and transiting surveys have a proposed explanation by Wolfgang & Laughlin [2012] as being due to different stellar samples for the two detection techniques, and either of two proposed mass-radius relationships. The stellar samples for HARPS and Kepler differ as HARPS target stars are selected for their slow rotation and low magnetic activity, and rejection of spectroscopic binaries. The proposed mass radius relationships treated terrestrial and low-density gaseous planets independently, and it was found that by simulating planets assuming an underlying distribution in mass, the number of planets detectable by Kepler was consistent with the number actually detected.

Understanding the selection effects of a survey is critical to estimating the completion and therefore predicting the true number and distribution of extrasolar planets. Usually analytic estimates are made [e.g. Hellier et al., 2012; Penev et al., 2012], but this work aims to outline specific areas of sensitivity or lack of sensitivity, and map proposed underlying distributions through these sensitivity maps to the observed distribution of exoplanets. Understanding the selection effects requires characterising the WASP detection process. Each real transit signal has a chance to be detected, which encompasses the observational biases and analytical techniques used during the transit detection process. To calculate this detection probability

p , synthetic transiting systems were generated and the transit signals they would produce were inserted into WASP data. With knowledge of the input parameters of these synthetic systems the detection probability was calculated across a range of orbital period and planet radius to generate a sensitivity *map*. The detection probability was then used to assess underlying occurrence rate predictions to assess their validity. This chapter outlines the transit synthesis process: section 2.2 summarises the implementation of the WASP hardware and analysis techniques, section 2.3 describes the process to generate a sample of synthetic transiting systems and the alteration of the WASP lightcurves accordingly.

The project proceeded in two phases: firstly the lightcurves of WASP planet-hosting stars were analysed to explore the selection effects with well constrained stellar parameters from exoplanet validation follow up measurements, providing confidence in the method. This was followed by a generalisation of the method to target stars without transiting planets, with which the underlying distribution of hot Jupiters was determined. This is described in Chapter 3.

The work described in this chapter and the following chapter involves studying the WASP planet search pipeline in detail, to understand the selection biases involved. The planet search pipeline was created by the WASP project, and the lightcurves used for this analysis were extracted by the WASP photometry pipeline. The detrending and transit search tools were also developed by the WASP project, with minor modifications performed by me to accommodate our analysis. All subsequent analysis of the results from applying the planet search pipeline to my synthetic data was implemented through my custom tools.

2.2 WASP project description

In the following section the WASP hardware and analysis techniques used are described.

2.2.1 Hardware

Each WASP site houses one equatorial mount with eight cameras (Fig. 1.10). The mounts are fully robotic allowing for autonomous control without requiring human intervention. The enclosures are also fully automatic opening the roof every night unless some predetermined conditions requiring the roof to shut are met, for example bad weather or high levels of dust. Each mount has a pointing accuracy of 30" rms over the entire sky, with a tracking accuracy of 0.01 arcseconds per second [Pollacco

et al., 2006]. This gives WASP an excellent ability to track stars for the duration of a night while they change position in the sky.

Each camera was sourced from Andor, with an e2v CCD. The CCDs have $2048k \times 2048k$ pixels, each is $13.5\mu m$ in size covering a sky area of 13.7 square arcseconds. The CCDs are back illuminated allowing for a $> 90\%$ peak QE at the operating temperature of $-50^\circ C$. The readout time for each exposure is ~ 5 s with 8-10 e^- readout noise from the 16-bit digitisation. The cameras are operated with a gain of $2e^- ADU^{-1}$.

Finally the lenses are Canon 200mm f/1.8 telephoto lenses allowing for a total field of view of 64 square degrees per camera. Each site observes multiple fields per night, by taking either one or two exposures per field then slewing to the next field, where each image is exposed for 30 seconds, with 5 second readout time. Each field is visited on average every 7 minutes allowing for up to 8 fields to be observed simultaneously per night [Smith et al., 2006].

2.2.2 Data reduction pipeline

The data reduction pipeline has been in place since the first data collection in 2006 excluding some improvements, and it remains largely unaltered for its main capacity of detecting planets. For more details, see Kane et al. [2004] for a description of the prototype instrument, Pollacco et al. [2006] for an overview of the project, Collier Cameron et al. [2006] for a description of the planet search algorithm and Collier Cameron et al. [2007a] for discussion of the parameter estimation.

Nightly bias, dark and flat frames are taken at dawn and dusk for every night the enclosure is open. Statistical validation is carried out for each frame with outlier frames rejected, and nightly master calibration frames are created from the remaining images. Bias and dark masters are computed from iteratively sigma-clipped means of the individual bias and dark frames. The master bias frame is subtracted from all individual dark, flat and science images, the dark frame is scaled to the exposure time of each science image before subtraction from flat and science images. Sky flats are taken in an automatic exposure sequence resulting in a constant maximum flux of ~ 28000 ADU. The images are dithered to move the star positions, and gradients are removed from each flat image by rotating each through 180° around the centre of the vignetting pattern, subtracting the rotated image and performing a planar least squares fit to the residuals. The gradient is then divided out of each flat frame, and a shutter correction is applied to the short exposure flat frames. Nightly flats are median combined with previous flats with an exponential weighting in time to create a nightly master flat frame, which each science image is

divided by.

The mount coordinates give the field centre to a precision of a few arcminutes. A subset of the Tycho-2 catalogue is extracted for this field centre; it is slightly larger than the coverage of a single camera to allow for stars near the edge drifting off the CCD. Star positions are extracted from each image using the Starlink `EXTRACTOR` package with a 4σ detection threshold above the background. The astrometric solution is calculated from matching the 100 brightest stars in the Tycho-2 subset with the extracted positions on the CCD allowing for translation, rotation, scaling and a barrel distortion term. The rms scatter on the difference between the resulting computed star positions and extracted source locations is always close to 0.2 pixels. A catalogue of targets for each frame is built from the USNO-B1.0 catalogue [Monet et al., 2003] of objects brighter than a magnitude of $R = 15$. All photometric measurements are then tied to these objects, and their positions are accurately known at the time of flux extraction.

For each frame an exclusion mask is created, by flagging any pixels which are within a magnitude-dependent radius of the catalogue star positions. This is used for calculating the sky background, by iteratively fitting a quadratic surface to the unflagged pixels. Each object in the input catalogue has three apertures placed at 2.5, 3.5 and 4.5 pixel radius along with a sky annulus from 13 to 17 pixels. The sky annulus rejects any pixels flagged in the exclusion mask to prevent stars affecting the sky background estimation. By observing the ratio of flux in the three apertures the PSF of the instrument and blending with neighbouring objects are monitored.

For each frame the heliocentric Julian date (HJD) is calculated for each object, as this varies over the field of view. The post processing pipeline PPWASP is run over the field which primarily removes known trends from the data. Primary and secondary extinction is corrected for with an iterative least squares fit made to the instrumental magnitudes calculated by the photometry pipeline and the airmass at which the observation was made. Stars showing excessive variance are down-weighted to reduce their impact to the global extinction correction. Following this around 100 bright non-variable stars are chosen and calibrated to their Tycho-2 magnitudes simultaneously calculating the instrumental colour response and system zero point. By using this each object is converted to the WASP V magnitude system and stored in large binary field files. Target stars are grouped into fields and typically observed by a single camera, though each field may be observed in multiple seasons. The lightcurves in each field form a set over which the planet hunting analysis is performed. A field file holds all information about a single WASP field per season and usually contain around 100k lightcurves, each with up to 40000 time points,

and information describing the observations or targets in question.

The lightcurves are detrended using the SYSREM [Tamuz et al., 2005] algorithm customised to run on the WASP data format. The algorithm removes systematic effects such as atmospheric extinction without prior knowledge of the effects. A sub-sample of bright non-variable stars are chosen from the field as standard stars with which per-image coefficients are iteratively fit to determine global trends in the data. Each object in the frame is given a weight value from iteratively minimising the residuals based on the per-image coefficient. It was determined that running the algorithm four times is optimal for the highest quality lightcurve without reducing any real transit signal. Detrended lightcurves are then further detrended with the TFA algorithm [Kovács et al., 2005] removing periodic trends in the data. The initial stage of TFA is similar in implementation to SYSREM: a template set of lightcurves is used to generate per-image basis functions for the dataset, and each star is allowed a coefficient which minimises the square of the residuals. A periodic basis function is iteratively computed from each lightcurve’s residuals after phase-folding and binning on a trial period and epoch, and is subsequently least-squares minimised to minimise the residuals further. This basis function is created during the transit search stage, so TFA is tightly coupled to the transit detection stage. Both TFA and SYSREM require all objects in the field to be present as the other lightcurves are used in the detrending process as basis functions.

The transit search is then run on each lightcurve using an adaptation of the box least squares (BLS) algorithm [Kovács et al., 2002; Collier Cameron et al., 2006], with a soft-edged transit model from Protopapas et al. [2005]. This implementation is known to the WASP consortium as ORION. This function is computed over a grid of period and transit width. The transit depth and goodness-of-fit statistic χ^2 are computed. The transit significance is given by the difference between the fitted χ^2 and null hypothesis of no transit ($\delta = 0$) and is denoted $\Delta\chi^2$. At every period the $\Delta\chi^2$ value for the best fitting transit depth and width is stored giving a transit significance for every period analysed. The five candidate detections with the highest signal detection efficiency (SDE, the significance of the periodogram peak compared to the fitted background level, see Section 2.4.2 for further explanation) are stored for further analysis by later stages of the transit search process. Candidate transiting objects found by the BLS are then analysed by a Markov-Chain Monte Carlo (MCMC) algorithm for transit refinement [Collier Cameron et al., 2007a]. During this process many parameters characterising the detection and most likely system configuration are calculated: both physically, describing the planet, star and their interactions; and parameters describing the detection process, for example the

ratio of the transit signal to the level of noise in the lightcurve.

Measurements made by the processes discussed in this section are uploaded to a MySQL¹ database, and statistics describing the WASP dataset are generated from this database to retrieve groups of objects which meet certain criteria. The next step in the WASP detection process is to select groups of candidates, which match criteria set to reduce the number of candidates to a manageable number to be inspected by the consortium (see Section 2.4.2). These parameters are designed to include all of the true planets whilst minimising the amount of time spent inspecting the WASP data. During this eyeballing stage candidates are classified based on a discussion by the consortium members. Candidates are broadly classed into two categories: probable false positives, or objects which will require further study either more data collected by the WASP instrument, or from larger telescopes for photometry and spectroscopically obtained radial velocity information. The promising objects will be discussed until the consortium is convinced that the object is a true planet.

2.3 Calculating the selection effects

2.3.1 Overview

To assess the selection effects for the WASP project, synthetic lightcurves were inserted into the standard WASP pipeline. The detection probability p was calculated from the fraction of detections to proposed models across a two dimensional parameter space, in orbital period and planetary radius. This procedure was applied in two phases: analysing the known planets, followed by analysing a random sample of the WASP catalogue.

The known planets were limited to contain all detected planets up to and including WASP-40 b. This was done to have a consistent data set to work on without needing to continually analyse new data when another planet was detected. The stopping point of WASP-40 b was chosen as a complete set for which refereed publications had been made when the analysis was performed. We note that the stopping point is somewhat arbitrary, provided sufficient planets were included. This set contains all planet designations except WASP-9 b which was found to be a false positive. By analysing the known planets (up to and including WASP-40 b) we studied each sensitivity measurement per planet in detail, and investigated the differences between sensitivity maps. The analysis was then extended to incorporate the analysis of a random sample of WASP stars to assess the sensitivity of the complete WASP project.

¹<http://www.mysql.com>

This chapter describes the lightcurve synthesis process, the modifications to the WASP pipeline and data and the discussion of the sensitivity maps for the known planet hosting stars. Chapter 3 describes the process of studying the sensitivity for the random sample of WASP stars, the comparison of sensitivity maps produced and the estimation of the occurrence rate of giant planets.

2.3.2 Simulation parameters

To place constraints on the underlying population of hot Jupiters, a method of applying the sensitivity to models of underlying populations was required. The transiting systems are characterised by parameters describing the bodies involved, and the orbital dynamics. Hot Jupiters exhibit a pile up in period at around 3 days, and to be able to study this feature we needed to be able to determine the underlying distribution in terms of period. The radius distribution of hot Jupiters was also an objective, as some planets are inflated beyond their predicted radius. We determined the selection effects in terms of these two parameters. The orbital period of the system is directly measured to high accuracy in transit lightcurves, so we chose this parameter over the orbital separation, which is inferred. Other parameters such as transit depth do not provide the same level of understanding as these direct properties of the system despite their high measurement precision. Previous studies of transiting planets [e.g. Howard et al., 2012] have expressed their sensitivity in this parameter space and it allows for a natural comparison. Studies of radial velocity planets on the other hand have expressed their sensitivity in terms of orbital period and planet *mass*. Comparisons are difficult between the occurrence rates measured for transit surveys and RV surveys, especially for hot Jupiters due to the radius degeneracy with mass, but consistent results have been found [Wolfgang & Laughlin, 2012, at least for $M_{pl} \leq M_{Nep}$]. Additionally expressing the results in this parameter space allows for a simple form for the Roche limit (Section 2.3.6).

The ORION implementation used by the WASP project searches by default from 0.35 days to 10 days, which we used as our period range for study, and is sufficient to allow the sensitivity to hot Jupiters to be estimated. At shorter periods than this, medium sized planets would reach their Roche limit, at longer periods the decreasing probability of transit and assumed lack of long period sensitivity of WASP reduces the probability of detection to almost zero. No WASP planets had been detected with periods > 8.5 days (WASP-84 b, Anderson et al. [2013]) so comparisons with the hot Jupiters could not be made outside this range. Radius values were limited to the range of $0.3 \leq R_p \leq 6 R_J$ which covers Neptune class planets to late type main sequence stars, where $R_J = 71\,492$ km. We did not want

to limit this range completely to that of the current population of known planets, this would inhibit possible conclusions about the existence of extremely inflated objects, and to ensure that the entire parameter space was explored whilst studying the limits of sensitivity at low radii, planet parameters were drawn from the logarithm of the ranges. To ensure uniformity across the two dimensional parameter space, an even grid was created with N_i bins for each i dimension (P or R_p). A uniform random offset δ_i was added to the bin centre point i_0 so that the period and radius values drawn for a synthetic system are

$$P = P_0 + \delta_P$$

$$R_p = R_{p,0} + \delta_{R_p}.$$

The random offset δ_i was chosen to be uniform within the bin; with low numbers of random points this guarantees an even coverage. The number of bins per dimension N_i was set to 50 providing a log spacing in period of $\delta \log_{10}(P) = 0.029119$, and $\delta \log_{10}(R_p) = 0.035563$. This allows for small scale structures to be visible whilst retaining low Poisson errors per bin. It has the effect of performing a Monte Carlo analysis whilst ensuring even coverage across the parameter space.

2.3.3 Other parameters

A transiting system cannot be fully described by period and radius alone. To close the system, the other parameters were synthesised or randomised whilst ensuring a self-consistent model.

The inclination i has a large affect on the transit recovery, the majority of possible inclination values show no transit signal at all. For a given system a transit is only visible if $\cos i < (R_p + R_\star)/a$. Since $R_p \ll R_\star$ the inclination i is often close to 90° . The orientation of a transiting system is expected to be distributed uniformly but the projection of this orientation is such that $\cos i$ is distributed uniformly.

Eccentricity e should not affect the detectability of a transit-like signal, only vary the transit duration, and alter the conclusions made about the object. The signal will be detected within a large range of eccentricities. Hot Jupiters generally have low eccentricities, often consistent with $e = 0$. A few WASP planets have $e > 0.1$ indicating that a non-zero eccentricity does not prevent a planet detection. Only circular systems were generated, but a possible extension is to simulate eccentric systems and their effect on the sensitivity of WASP.

The WASP data quality is variable, with large systematics at 1-day periods and certain harmonics. As the project observes from the ground the day/night sam-

pling pattern is also important. Any assessment of the detectability of a proposed synthetic planet must not be dominated by these regions of missing or low quality data. By randomising over the time of mid transit t_0 the data quality was averaged over a whole lightcurve. Multiple synthetic transits were inserted into the same lightcurve, so the detectability was thoroughly explored.

A transiting survey is only sensitive to planet mass in that more massive planets tend to be larger, at least below approximately the mass of Saturn. The assessment of the selection biases is therefore not sensitive to mass, as planet radius is a parameter of interest. The radii and masses of hot Jupiters are typically independent, as the mass-radius relation is degenerate at around $1M_J$ (see Fig. 1.18).

The planet-hosting stars have been well studied during the planet confirmation stage, and so have reliable stellar parameters from spectroscopic estimation of the effective temperature, and mass and radius determination from stellar isochrones. The stellar radius is used to estimate the transit width and depth, and the effective temperature is used for the limb darkening coefficient estimation. No point of this study was concerned with synthesising alternative stellar parameters. We required the stellar information for an accurate transit lightcurve but we did not alter the nature of the star. For this reason we do not synthesise alternative stellar parameters.

2.3.4 Lightcurve synthesis

Figure 2.2 shows the geometry used in the transit synthesis. To alter the lightcurves and generate a realistic transit signature, the analytic model from Mandel & Agol [2002] was used. The model incorporates non-linear limb darkening from Claret [2000] of the form

$$\frac{I(\mu)}{I(1)} = 1 - \sum_{n=1}^4 c_n \left(1 - \mu^{\frac{n}{2}}\right) \quad (2.1)$$

where $I(1)$ is the specific intensity at the centre of the disc, and $\mu = \cos(\theta)$ θ is the angle between line of sight and the emergent intensity (γ in Claret [2000]). We note that through geometry

$$\mu = \cos(\theta) = (1 - r^2)^{1/2}. \quad (2.2)$$

where $0 \leq r \leq 1$ is the normalized radial coordinate of the planet on the disk of the star.

Claret [2000] tabulate the coefficients from ATLAS simulations [Kurucz, 1979] producing the coefficients c_n for a temperature range of $3500 \leq T_{eff} \leq 50000\text{K}$, for turbulent stellar velocities of 0, 1, 2, 4 and 8 km s^{-1} , for surface gravities ranging

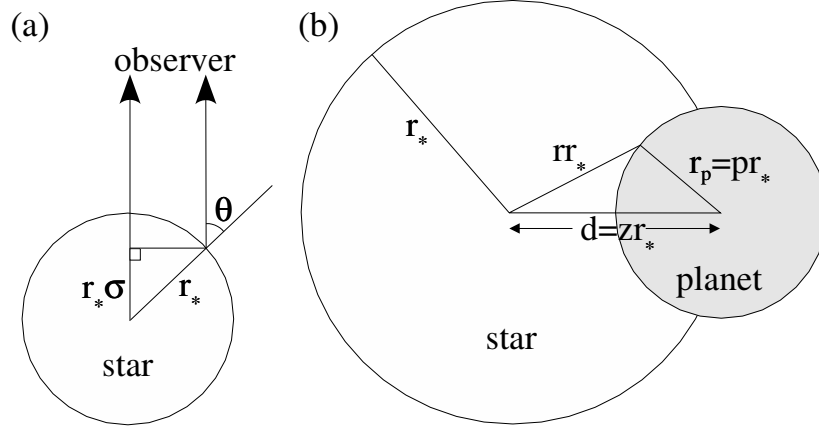


Figure 2.2: (a) Geometry of limb darkening. The star is seen edge-on, with the observer off the top of the page. The star has radius r_* , and θ is defined as the angle between the observer and the normal to the stellar surface, while $\mu = \cos \theta$. (b) Transit geometry from the perspective of the observer. Reproduced from Mandel & Agol [2002].

from $0 \leq \log g \leq 5$ and for metallicities ranging from $-5 \leq \log[M/H] \leq 1$ in multiple optical and near-IR filters. Following Collier Cameron et al. [2007a] the stars micro-turbulent velocity was assumed to be 2 km s^{-1} , the surface gravity was assumed to be $\log g = 4.5$ and the metallicity assumed to be $\log[M/H] = 0.1$ allowing for planet occurrence trends with metallicity. Limb darkening coefficients were then taken in the R band from interpolating the four coefficients based on the stellar effective temperature. Figure 2.3 shows the flux profile of three synthetic stars with varying temperatures by substituting Eq. 2.2 for r in Eq. 2.1. Table 2.2 shows the calculated limb darkening coefficients c_n for the three temperatures shown in Fig. 2.3.

To generate the synthetic transit signals the small planet approximation from Mandel & Agol [2002] as used, for simplicity and speed of evaluation. The approximation holds for the majority of the planet range. We followed the method of Mandel & Agol [2002], which is summarised here. For a given lightcurve the JD values t_j were converted to normalised separation of body centres $z(t_j)$ by the following relation [Collier Cameron et al., 2007a]:

$$z(t_j) = \frac{\sin^2 \phi_j + (bR_*/a)^2 \cos^2 \phi_j}{R_*/a} \quad (2.3)$$

where ϕ_j is the orbital phase for data point j : $\phi_j = 2\pi(t_j - T_0)/P$, b is the impact parameter $b = a \cos i / R_*$ and a is the orbital separation. We generated normalised




Figure 2.3: Stellar flux profile $I(r)$ for three synthetic stars with temperatures indicated in the legend in Kelvin. Limb darkening parameters used are given in Table 2.2.

Table 2.2: Limb darkening coefficient values calculated for the three temperatures shown in Fig. 2.3.

T_{eff} / K	c_1	c_2	c_3	c_4
4000	0.586	0.190	0.701	0.222
5600	0.610	0.340	1.09	0.534
10000	0.595	0.193	0.151	0.0298

model lightcurves so we assumed $I(1) = 1$ in Eq. 2.1 therefore yielding the normalised stellar disk intensity as a function of normalised radial coordinate on the disc of the star $I(r)$. For convenience we define $\Omega = \sum_{n=0}^4 c_n(n+4)^{-1}$, $k = R_p/R_\star$ and $a \equiv (z - k)^2$. The normalised flux received at distance z is

$$F = 1 - \frac{I^\star(z)}{4\Omega} \left[k^2 \cos^{-1} \left(\frac{z-1}{k} \right) - (z-1) \sqrt{k^2 - (z-1)^2} \right] \quad (2.4)$$

[Mandel & Agol, 2002] where $I^\star(z) = (1-a)^{-1} \int_{z-k}^1 I(r) 2r dr$.

To assess the accuracy of the transit synthesis process, a real WASP lightcurve was altered. The full lightcurve for WASP-12 b was extracted and normalised by dividing by the median of the out-of-transit flux. Points with low precision measurements (the fractional error of the flux measurement $> 2\%$) were rejected, and the lightcurve of the remaining points sigma clipped to emphasise the transit signal. The phase-folded lightcurve is shown in Figure 2.4 in the top left panel, and the true transit signal is clearly visible. The model flux for each point was generated using the method described above and is shown in the bottom left panel of Fig. 2.4, taking the required parameters from the exoplanet database [Wright et al., 2011]. The flux points were corrected through $f_{corr,i} = f_i - m_i + 1$ where f_i is the measured flux at point i , m_i is the model flux and $f_{corr,i}$ is the corrected flux at point i , where the resulting lightcurve is shown in the top right panel of Fig. 2.4. The model lightcurve has not been fitted in any way, just the literature parameters have been used, and the resulting lightcurve shows no signs of a transit. The test was completed by inserting the transit of a different planet WASP-6 b into the lightcurve, where the altered flux measurement is given as $f_{altered,i} = f_i + m_i - 1$. This is shown by the bottom right panel of Fig. 2.4, and displays a transit with much shorter duration due to the longer period of WASP-6 b.

2.3.5 Testing the transit synthesis method

To determine the accuracy of the transit alteration method, an initial test was performed. The parameter space of interest was restricted to only altering the planet radius, whilst keeping the other parameters the same. A series of transit models were generated for WASP-12 b, the planetary radius was increased in increments of $0.1R_J$ from $1R_J$ to $2.4R_J$ enclosing the true value of $1.736R_J$ [Chan et al., 2011]. For each model generated, the lightcurve for the field file containing WASP-12 b was altered subtracting the real transit signal and inserting the new model in its place. To determine whether the object was detected or not, the BLS search from ORION (described in Section 2.2.2) was run using its single object mode. This




Figure 2.4: The transit removal and synthesis process. Top left: the original WASP-12 lightcurve after SYSREM, zoomed around the transit region. Bottom left: the transit model used to remove the WASP-12 b transit signal. Top right: lightcurve of WASP-12 after transit removal. Bottom right: the WASP-12 lightcurve with a different transit inserted, based on WASP-6 b and re-phased based on WASP-6 ephemeris.




Figure 2.5: Initial test results for only varying the planetary radius of WASP-12 b. Black circles indicate tested values, the dashed vertical line represents the literature radius value of $1.736 \pm 0.092 R_J$ [Chan et al., 2011], the grey region representing the uncertainty.

test was preliminary as no detrending was performed. Even without detrending the significance of the transit signal for WASP-12 b is enough for a rudimentary detection. A clear point where the signal transitions from not detected to detected is seen, suggesting that the selection effects can be measured successfully, though the true planetary radius is close to the detection limit emphasising the importance of accurately characterising the selection effects in detail.

The determination of the critical radius point for WASP-12 b suggested that the sensitivity of WASP was calculable. To calculate the sensitivity of the project itself, the ability to apply this test for detection was required for an arbitrary collection of objects. The planet search analysis used in WASP implements the detrending described in Section 2.2.2, so for an accurate characterisation of this pipeline the same steps needed to be repeated. Both SYSREM and TFA required the full set of lightcurves contained in a field, and so for the analysis of an arbitrary object the full field files were required.

2.3.6 Model rejection

Proposed models are rejected if they meet certain criteria. Planets cannot exist if they have reached their Roche limit, which defines the minimum orbital separation the system can have before the planet is disintegrated by tidal forces from the star.

It is given as

$$R_p = 0.462 a_R \mu^{1/3} \quad (2.5)$$

where $\mu = m_p/M_\star$ is the planet-star mass ratio [Ford & Rasio, 2006]. By substituting for a_R in Eq. 2.5 using Keplerian dynamics

$$P^2 = \frac{4\pi^2}{GM_\star} a^3 \quad (2.6)$$

the Roche limit can be expressed independently of the host star, only requiring the orbital period, planetary radius and planetary mass. Recasting the Roche limit separation into a Roche limit planetary radius, the radius at which the planet is tidally disrupted:

$$R_{RL} = 0.462 \left(\frac{Gm_p}{4\pi^2} P^2 \right)^{1/3} \quad (2.7)$$

For this study we set m_p to $13M_J$; the point at which core burning starts and the object is classed as a brown dwarf, consequently we overestimated the Roche limit for most synthetic planets. We were not concerned with the planetary mass and the radius of a hot Jupiter is independent of the planetary mass (see Fig. 1.18), so we do not calculate the true Roche limit for the synthetic planetary systems. Systems which definitely cannot exist were rejected but no systems which can exist were rejected, ensuring that we were exploring the full parameter space.

Non-transiting systems were not synthesised as there would be no transit signal to detect. A system is transiting if $|b| \leq 1 + R_p/R_\star$, which includes grazing transits, where the disk of the planet is only partially obscuring the stellar disc at the mid-transit point. We generated non-transiting planet models in order to include the probability of transit in our calculations; they are noted for inclusion into the non-detection statistics.

2.3.7 Implementation of the data modification process

The analysis tools were written modularly in five parts, two of which were custom written and the remaining three comprise of the standard WASP analysis pipeline with some alterations. These alterations had to be made to accommodate the extra lightcurves properly, between 15500 and 16500 lightcurves were inserted for each data file. The tasks operate on a WASP field file (see Section 2.2.2) as the unit of data over which the analysis is run.

The detrending programs SYSREM and TFA heavily rely on defining a set of basis functions from the bulk photometric behaviour of a set of lightcurves. Inserting

thousands of almost identical lightcurves (at least their out of transit data would be identical) would heavily bias the detrending towards removing just the real signal. The algorithms used by these programs were altered to exclude any objects from the ensemble basis function calculation, that had been inserted as part of the selection effects study. The original WASP lightcurve was left in the file so the basis functions used by the detrending should be identical to those calculated during the original analysis. For each run of SYSREM each image has a weight and each object has a weight, the synthetic objects were simply not allowed to affect the per-image weights. A similar exclusion was incorporated for TFA, not allowing any synthetic objects to affect the detrending behaviour. This was tested by adding extra lightcurves to a field file and comparing the per-image basis functions to the original file. The points were identical to within floating point errors indicating that the extra lightcurves did not adversely affect the detrending to make this alteration unrealistic. A block diagram of the data flow is shown in Fig. 2.6, and the process involved:

- 1 The first step involved the model parameter generation and storage of these parameters for further subsequent use. These parameters are the inputs to the transit lightcurve generation stage. Each model system was modelled from altering the true underlying system, and categorised into transiting or non-transiting. The model parameters were stored in an SQLITE3 database for easy querying.
- 2 Next and in a separate task, the model lightcurves were generated for the synthetic systems which were transiting. The source lightcurve was duplicated in the field file, the existing transit removed and the transit signal from the new synthetic system inserted. The two steps were kept separate for ease of development, and so verification tests could be performed on them individually.
- 3-4 SYSREM and ORION are capable of running on multiple cpus simultaneously and so require a multi-core computer to maximise their efficiency. These tasks are mostly unmodified from their behaviour in the standard WASP pipeline, but as noted in a previous section they account for a large number of very similar lightcurves when generating the per-image detrending coefficients. ORION was further customised to only analyse synthetic objects inserted by step 2.
- 5 The MCMC task is single-core only but some efficiency modifications have been made. The program itself analyses a single lightcurve with command line arguments for the filename and lightcurve id, so a wrapper script was written in PYTHON to run this program for every synthetic object detected by ORION in the file.

To maximise the efficiency of manipulating the data, and to automate the analysis, the tasks above were run on the computing cluster at Leicester University. A processing queue was allocated for this project, providing six six-core machines for SYSREM and ORION which could run in parallel, and six other machines with multiple task slots each for the single-core tasks, all managed by the SUN GRID ENGINE² system. The parallel tasks had to be run separately on the parallel queue meaning that the whole process could not be run as a single job, some inter-job communication was required.

A job submission architecture was created so the first job comprising of the model parameter generation and file alteration when finished on the single-core node would submit the next job on the parallel queue containing the parallel tasks SYSREM and ORION, which in turn would submit the final single core task MCMC. Each task script managed the required input and output files and handled moving the files to each node’s local disk for speed.

A full analysis of each field file depends on the number of objects inserted, but would take around two days to complete. As only six parallel tasks were able to run simultaneously this created a bottleneck. Each file had to pass through the ORION stage and with only six running simultaneously this was the limiting factor to the speed of the analysis, even though each MCMC task was slower to run individually. In total three months of computing time were spent to analyse the WASP selection effects.

2.3.8 Testing and verification

Some debugging techniques were used to ensure the models generated and inserted are valid. The shape of the transit was plotted for various orbital separations (Fig. 2.7) to ensure expected behaviour. As the planet moves further from the star, providing it is still transiting then the transit duration decreases as it’s orbital velocity is higher to maintain the same period. From Winn [2010] the transit duration is given by

$$T_{tot} \equiv \frac{P}{\pi} \sin^{-1} \left[\frac{R_{\star}}{a} \frac{\sqrt{(1+k)^2 - b^2}}{\sin i} \right] \quad (2.8)$$

so the phase-normalised transit duration T_{tot}/P decreases as the separation increases.

²<http://www.oracle.com/us/products/tools/oracle-grid-engine-075549.html>, now managed by Oracle

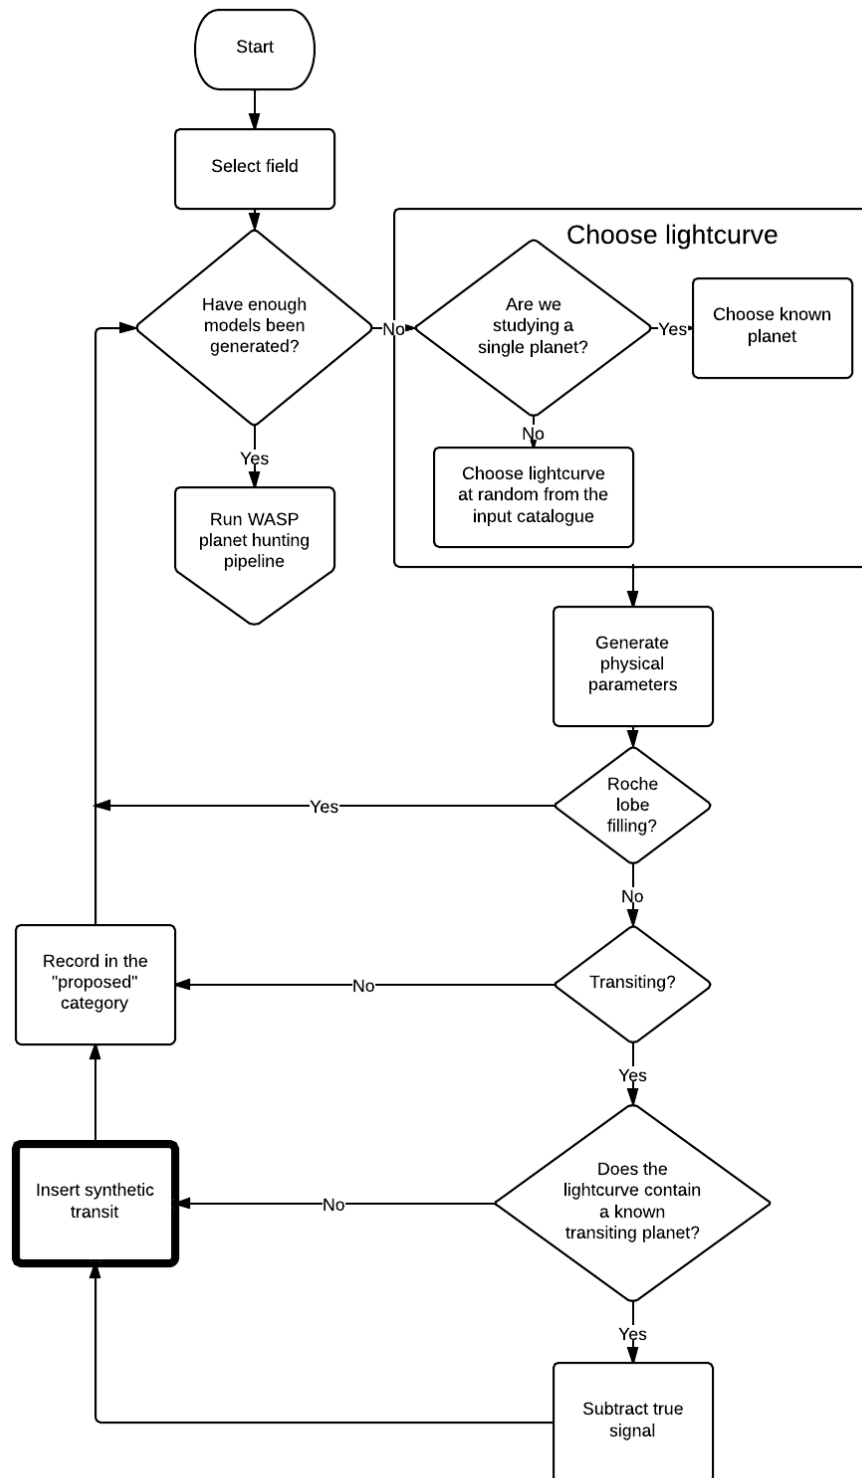


Figure 2.6: Schematic diagram of the lightcurve alteration process.




Figure 2.7: Transit shapes for a range of semi-major axes indicated by the legend, for the same model.

Figure 2.8 shows a series of verification images for the model generation stage, representing the proposed models before their insertion into WASP data. This was created to ensure the parameters generated were consistent with the input assumptions. Of particular interest are panels regarding rejecting non-transiting models: the top-right panel shows how the range of inclinations of the accepted objects decreases as the orbital period increases and the probability of transit decreases. The top-left panel shows that models were proposed evenly in log-period and log-radius parameter space, whereas the top-middle panel shows the decreasing number density with long orbital periods as the probability of transit decreases. Interestingly the bottom-middle panel shows a steep decrease in the number of accepted models with an impact parameter $b > 1$ suggesting that grazing transits are uncommon. The pie chart shows that roughly 2/3 of the proposed models are rejected either as non-transiting or Roche limit reaching, which matches expectations as the transit probability is low even at modest separations.

Figure 2.4 shows the subtraction and re-insertion process for a single WASP-12 dataset. The original transit has clearly been removed from the original data source with only the literature values for the system parameters.

2.4 Computing the sensitivity map

The synthetic lightcurves were inserted into WASP field files and analysed by the WASP pipeline, slightly altered to accommodate the large number of synthetic objects inserted. The sensitivity map is a tool to study the sensitivity of WASP in a two dimensional period-radius plane. We wanted to know how likely a planet

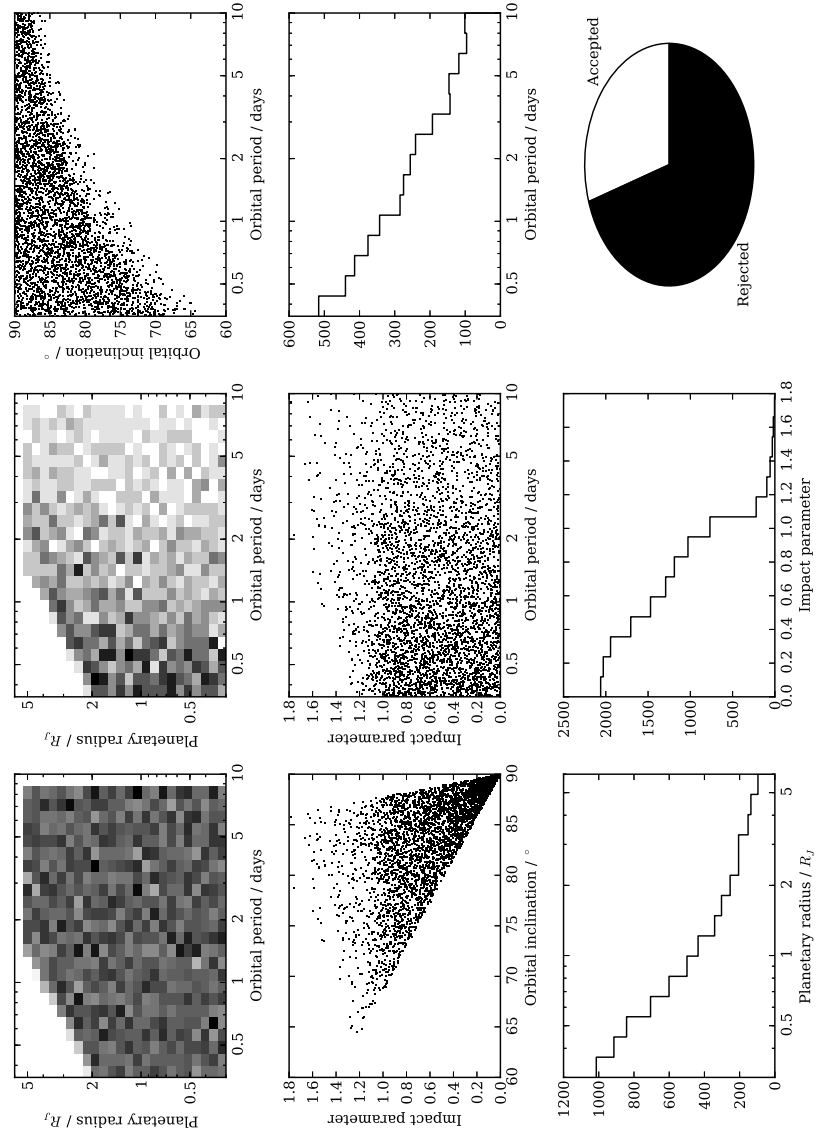


Figure 2.8: Statistics run on input models to ensure correct behaviour. Models are based on parameters from WASP-12 b. Important parameters are plotted against others and histograms made to verify the accuracy and correctness of the input model generation. Panels show the parameters for models accepted and were generated and analysed, apart from the pie chart which reflects the ratio of accepted to rejected, and the top-left panel which represents all models proposed.

with period P and radius R_p was to be recovered after being studied by the WASP analysis methods. To do this the fraction of synthetic models which were *accepted* by the WASP pipeline was calculated by binning the parameter space into the same 50 bins as the underlying uniform grid on which the systems were inserted on. The definition of *accepted* now follows as some selection criteria must be met before the object is classed as detected.

2.4.1 Acceptance

A synthetic model is defined as accepted if it meets all the following criteria:

- the object is detected by ORION,
- the detected period is the same or a harmonic of the input period,
- the object passes the standard WASP consortium selection cuts

We assume that a WASP consortium member would identify periods which are a factor of two or a half of each other, so we allow periods which match to within 1% of an integer multiple of the period, where the multiples are 1/2, 1 and 2. This ensures both that we are detecting the input signal, and that mismatched periods are accounted for. It is common during the eyeballing process to investigate period harmonics when assessing the validity of a target.

The WASP consortium defines a set of selection criteria for a candidate, reducing the number of potential lightcurves that require human analysis by removing likely false positives and lightcurves with questionable data quality. These cuts are performed against parameters calculated by the analysis pipelines, and they range from observational (e.g. the number of points in a lightcurve) to fitted parameters of the physical system (e.g. estimated planetary radius). The cuts are made both to reduce the amount of bad data that make it to human inspection, and to reduce the number of false positive objects which imitate transit-like events. Since visually inspecting each lightcurve is time consuming these cuts help focus the analysis onto only possible candidates.

The WASP planet search process continues for one more stage, the eyeballing stage which involves each consortium member who is participating to make a judgement on these lightcurves which pass the selection cuts. The lightcurves being eyeballed are also subject to secondary checks from a different consortium member. Finally a discussion between all participating consortium members takes place in which each candidate is discussed. The purpose of the eyeballing stage is to reject spurious detections made by ORION. A sample of the detected synthetic

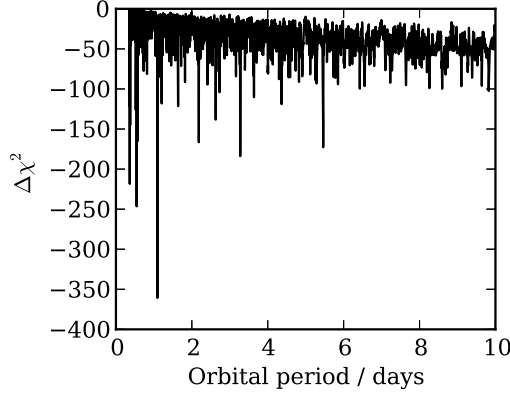


Figure 2.9: $\Delta\chi^2$ periodogram for planet WASP-12 b. The clear negative peak at 1.09 days is the signal picked up by the BLS search.

lightcurves which passed the selection cuts imposed were inspected by four members of the consortium. Proposed transiting systems that fall in regions of low sensitivity were examined to determine if the smaller radius, or long period signals would be accepted as candidates. The lightcurves were found to be convincing and would pass the standard eyeballing selection.

2.4.2 Selection cuts

Selection cuts are made during the planet search stage to reduce the list of targets to only likely candidates. False positive signals for which a measurable statistic of the lightcurve can clearly determine its non-planetary nature are rejected so the only remaining targets are true planets, and false positives which require human assessment. These selection cuts are defined below.

$\Delta\chi^2$ The BLS search calculates a $\Delta\chi^2$ value, the difference between the χ^2 of a lightcurve without the transit model and the χ^2 with the transit model for each period. An example periodogram showing the $\Delta\chi^2$ is shown in Fig. 2.9 and is used by the WASP project for visual inspection of the strength of this signal, and any harmonics that may exist. It is the largest difference for every combination of box width and epoch for a given period. This value is negative for strong signals, and the presence of a large negative peak indicates a good match to a transit-like object. The cut for this parameter is designed to pick out a high absolute difference in χ^2 which represents a significantly better match for a transit-like signal than without.

Accept: $\Delta\chi^2 < -40$.

SDE This parameter is the significance of a peak compared to the background level over the entire period range searched. A linear fit is made to the background level of the $\Delta\chi^2$ periodogram. A better match to a transit will produce a higher $\Delta\chi^2$ value compared to other test periods, and will stand out from the background more. The parameter was suggested by Kovács et al. [2002] and can be expressed as

$$SDE = \frac{S_{peak} - \bar{S}}{\sigma_s}$$

where S_{peak} is the height of the peak, and \bar{S} and σ_s are the mean level and scatter in the noise continuum of the periodogram. This parameter represents the significance of any peak detected by the BLS search and so we make a cut on this to reduce the number of marginal detections that are likely false positives.

Accept: $SDE > 6$.

$P(R_p)$ The probability that the fitted planet radius $P(R_p)$ is $\leq 1.5 R_J$ is calculated from the MCMC analysis and represents the likelihood of the radius being of planetary origin. 3000 proposals are made to explore the parameter space and if 10% of the proposals have $R_p \leq 1.5 R_J$ the object is accepted. This cut was studied by Collier Cameron et al. [2007a] and was found to be a good way to exclude false positive systems. Note that inflated WASP planets with larger radii than $1.5 R_J$ have been found, indicating that perhaps this cut is too restrictive.

Accept: $P(R_p) > 0.1$.

$(S/N)_{red}$ The signal to red noise ratio. The significance of each transit is calculated by comparing the transit depth to the noise observed in the lightcurve. Each significance is combined, but assuming the noise is red and so does not combine simply with the square root of the number of transits, but instead with a fitted parameter β (where for white noise $\beta = 0.5$).

Accept: $(S/N)_{red} < -6$.

Q This parameter is calculated during the MCMC analysis, and is the acceptance criterion for a set of proposal parameters. Each step in the chain i has an associated χ^2 value calculated from a lightcurve model and the data. Q (following the formalism of Collier Cameron et al. [2007a]) value incorporates the main sequence prior in the Bayesian probability estimation:

$$Q_i \equiv \chi_i^2 + \frac{(M_{\star,i} - M_0)^2}{\sigma_M^2} + \frac{(R_{\star,i} - R_0)^2}{\sigma_R^2} \quad (2.9)$$

where χ_i^2 is the goodness of fit criterion for the transit model fitting, M_0 and R_0 are the prior stellar mass and radius based on an estimate from J-H colour (see Section 3.1.2 for an explanation). The Metropolis-Hastings rule defines if $Q_i < Q_{i-1}$ the new set is accepted, or if $Q_i > Q_{i-1}$, the new set is accepted with a probability $\exp(-\Delta Q/2)$, where $\Delta Q \equiv Q_i - Q_{i-1}$.

Accept: $Q < 7$.

(S/N)_{ellipse} The ellipsoidal modulation of the lightcurve is fitted from the out of transit data. Ellipsoidal modulation is an indicator of a binary star where the secondary object gravitationally distorts the star and causes a flux modulation twice per orbit. This feature is not unique to binary stars: HAT-P-7 b was found to exhibit ellipsoidal modulation but this was only detected using the high precision Kepler satellite [Welsh et al., 2010]. The amplitude of the noise is calculated and compared to the signal which gives the signal to noise of the ellipsoidal noise. This cut prevents obvious variable stars with clear sinusoidal modulation.

Accept: $(S/N)_{\text{ellipse}} < 6$.

clump_idx The ORION implementation records the five most significant peaks, and searches are run multiple times for each lightcurve with different detrending techniques, or if the object is measured in multiple fields. The CLUMP_IDX parameter describes how close these candidates' detected periods and epochs are together. A high value suggests that the same period was detected multiple times with either different data or analysis techniques, so the periodic signal detected is likely to not be a systematic error.

Accept: $\text{CLUMP_IDX} > 0.25$.

dilution_v The dilution for each catalogue object is calculated from the V magnitudes from the NOMAD catalogue, where the flux from each object j within 48'' of the target object i is calculated and summed. This flux is contaminant flux, and so the dilution is

$$\text{dilution} = 1 - \sqrt{\frac{f_i - \sum_j f_j}{f_i}} \quad (2.10)$$

where $f = 10^{-0.4V} \times 10^6$, scaled up by 10^6 to account for the zero point of the instrument in micro-Vegas.

Accept: $0 \leq \text{DILUTION_V} \leq 25$ (%).

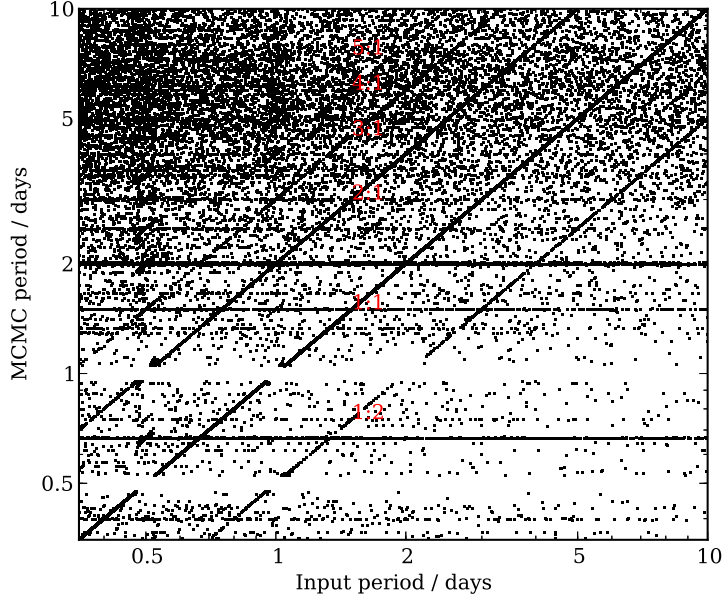


Figure 2.10: Input period against recovered period for every accepted synthetic object. Period harmonics are noted in red.

$N_{pts,good}$ This cut rejects short lightcurves which will not have enough flux measurement density around the transit to quantify the transit well enough.

Accept: $N_{pts,good} > 1000$.

N_{trans} This simply requires more than three transits to have been detected, giving an accurate period.

Accept: $N_{trans} > 3$.

2.4.3 Period matching

The synthetic models were accepted if the period detected was the same, twice or half the input period, under the assumption that during the eyeballing process this would be noticed and alternative periods tested. Figure 2.10 shows the input periods and recovered MCMC periods for every synthetic object which passes all of the other tests described above. Clear gaps are visible horizontally at 0.5 and 1 days reflecting ORION's hard coded rejection of these objects due to likely contamination from periodic signals reflecting nightly systematics. Related to this rejection are the sharp features at 0.75 and 2 days as these are common one day harmonics, and are not searched at one day periods. Strong lines diagonally show matches at 1:1, 1:2 and 2:1 with fainter lines at 3:1, 4:1 and 5:1. These show that a non-negligible

number of input synthetic objects were detected at harmonics of the true period. For input periods of 0.5 and 1 days on the edges of the 5% rejection regions we see deviations from the constant period ratio behaviour observed elsewhere in the plots. This is due to uncertainties in ORION's or the MCMC estimation of the period where input periods within the rejection region are fitted as having periods just outside the rejection region. This shows that the rejection regions for ORION are not completely inhibiting the possibility of detection of objects with near-day periods.

2.5 Sensitivity maps

To generate a sensitivity map, detected models and proposed models were binned separately into 50 period and radius bins across the entire range. The ratio of the resulting 2D histogram represents the fraction f of detections n_{detect} to proposed transiting systems n_{prop}

$$f_{P,R_p} = \frac{n_{detect}}{n_{prop}} \quad (2.11)$$

and therefore the sensitivity of the WASP project to planets with the particular period and radius P and R_p .

Figure 2.11 shows two sensitivity maps generated for WASP-12 b and WASP-7 b. Sharp drops in sensitivity are visible around 0.5 days and 1 day. ORION rejects any periodic signals within 5% of 1 day or its half day harmonic as the source of such signals is likely to be from nightly effects as fields rise and set, or similar systematics related to observing nightly from the ground. The upper left region has no sensitivity because planetary systems cannot exist here as they would exceed the Roche limit for $13M_J$ planets. The diagonal dashed lines represent regions for different assumed planet masses: $0.5M_J$, $2M_J$ and $10M_J$ where a planet cannot exist as it would be tidally disrupted. The two maps shown in Fig. 2.11 show differences in the shape. The sensitivity map for WASP-7 shows an increased sensitivity to smaller longer period planets, whereas the map for WASP-12 shows an increased sensitivity to larger ($R_p \geq 2 R_J$) planets. It is interesting to note that WASP-7 b would likely not be detected in the lightcurve for WASP-12 based on the sensitivity measurements shown here. The differences in the sensitivity for each star emphasises the need to study each object individually in this detail.

2.5.1 Planet trends

As discussed in Section 2.3.1 the first phase of the project involved studying the known transiting planets individually. With a sensitivity map for each planet, po-

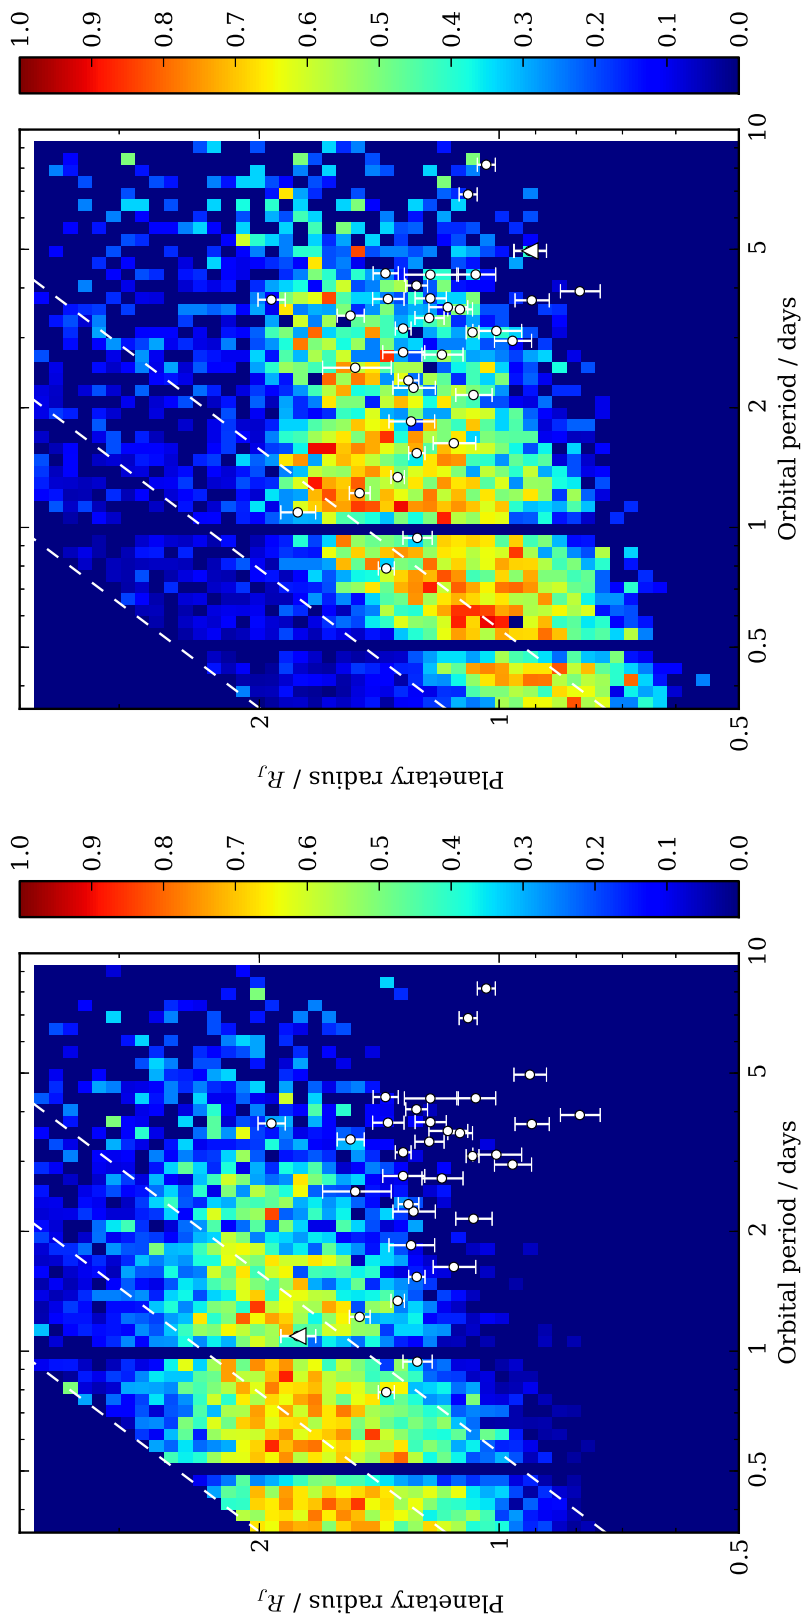


Figure 2.11: Sensitivity maps for WASP-12 b (left, $P = 1.09$ days, $R_p = 1.736 R_J$) and WASP-7 b (right, $P = 4.95$ days, $R_p = 0.915 R_J$). The colour scale represents the fraction of detections f_{P,R_p} . The parameters for the true planet are marked with a white triangle, and all other WASP planets are marked with white circles. The white dashed lines represent the Roche limits for planets with masses $0.5M_J$, $2M_J$ and $10M_J$ with higher mass lines being closer to the top left of the plot.

tential trends between the map shapes and measured parameters were assessed. A trend in the ancillary parameters such as the stellar brightness would require the mapping out of the overall sensitivity as a function of these parameters. The sensitivity maps made from each individual planet’s data were ordered by a selection of stellar parameters: the V magnitude, $[\text{Fe}/\text{H}]$, T_{eff} , R_{\star} and M_{\star} . The V magnitude maps well to the photon noise of a lightcurve assuming no other systematics dominate. Given a photon noise limited observation, a transit signal in a brighter object is more significant compared to the background noise compared to a fainter object. We choose T_{eff} as a proxy for spectral type, as later spectral types are smaller and so transit signals for similar sized planets might appear more significant. The stellar mass and radius are tested also for a direct comparison in addition to the temperature proxy. The two maps shown in Fig. 2.11 show a possible trend with V magnitude as WASP-7 ($V = 9.5$) is 2.2 magnitudes brighter than WASP-12 ($V = 11.7$), and shows an increased sensitivity towards smaller planets which would be expected if only white noise were present in the WASP system.

Sensitivity maps for the targeted stars were generated and ordered for the parameters listed above. An example showing the trend in V magnitude is shown in Fig. 2.12. The array of sensitivity maps for each parameter listed above were visually inspected for trends. See Appendix A for the remaining trend panels for T_{eff} , $[\text{Fe}/\text{H}]$, R_{\star} and M_{\star} . The V magnitude array for example shows no clear trend in either absolute value of sensitivity, or the general shape of the sensitivity with some bright stars being quite insensitive (e.g. the panel labelled (a) with $V = 9.75$), and fainter stars showing a wide range of sensitivity (e.g. the panel labelled (b) with $V = 11.89$).

Similarly the other parameters chosen for analysis show no clear trend either. The lightcurves of the known planets used for the sensitivity maps were inspected and no trends were seen. The lack of variation and trends suggests that the data quality is the dominant factor for studying the sensitivity of a lightcurve to transiting planets. Each sensitivity map shows significant differences between objects, which suggests the need to perform the transit synthesis process for each target star. In the following chapter we drew stars from a complete sample of stars around which we could have detected a transit. The lack of trends observed with the known planet hosting sample motivates the need for the transit synthesis process for each star.

2.5.2 Shaping the sensitivity map

The acceptance requires three steps (c.f. Section 2.4.1): the signal needs to be detected by ORION, the detected signal must have the same period or is a nearby

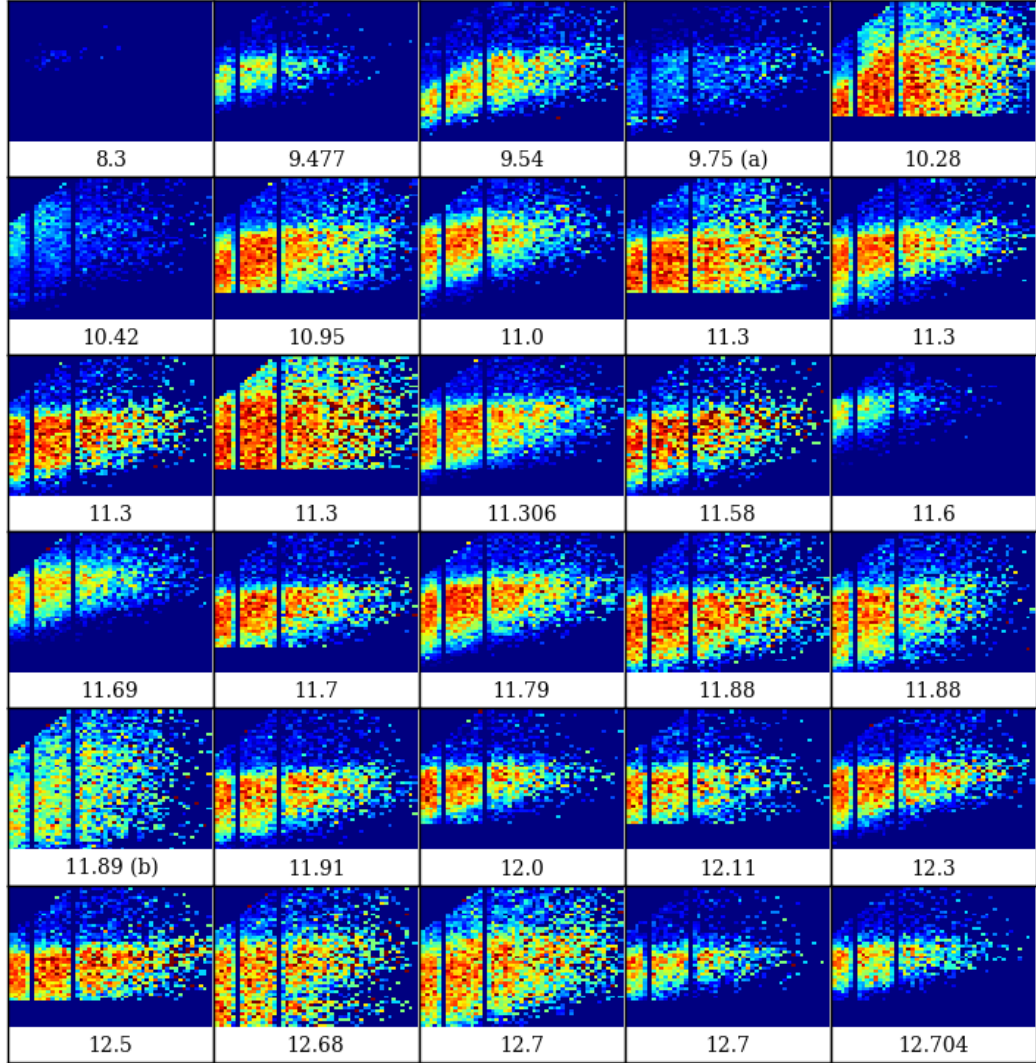


Figure 2.12: Sensitivity maps generated for the WASP planet host stars, sorted by V magnitude which is labelled in the panels. All maps share the same P and R_p axis ranges, and colour scale ranges, and do not include non-transiting systems in the normalisation. Panels labelled (a) and (b) are referred to in the text.




Figure 2.13: Orion detection map. The colour scale represents the fraction of models recovered by ORION.

harmonic of the input period, and the object must pass all of the selection criteria. To diagnose regions where synthetic systems were recovered or not, some tests were performed. The fraction of ORION detections across the parameter space is shown in Fig. 2.13. The sharp line is an artefact of initially incorrectly setting the radius range at input, and requiring extra data to be analysed at smaller radii. In general ORION detects $> 60\%$ of the targets across the majority of the parameter space, being most successful at short periods and large radii as expected. We note the fraction of detections does not decrease to zero for radii $< 0.5 R_J$, which is due to particularly favourable chance alignments in the synthesised lightcurves.

To determine which areas of the sensitivity maps were shaped by which selection cut parameters, and which selection cuts have the most effect on the shape of the sensitivity maps, a collection of 2D maps were created showing only the fraction of synthetic systems that were accepted based on a single selection cut alone. Figure 2.14 shows these maps for the four which have the most affect on the shape of the sensitivity map. Each map shows the fraction of objects in each period and radius bin which pass the cut specified.

By inspection of Fig. 2.14 the only selection cut affecting the detection of transit signals comes from the $P(R_p)$ cut, which restricts the maximum detectable planet radius. The cut is designed to reject false positive systems whilst including

real planetary systems [Collier Cameron et al., 2007a], but was determined before the presence of inflated planets with radii $R_p > 1.5 R_J$ was known. This reduced sensitivity to inflated planets may have prevented true planets from being detected by WASP, though some do exist. Typically the posterior planet radius from the MCMC analysis of these inflated WASP planets indicate a smaller radius than follow up observations have determined, and the selection cut is based on the planet probability rather than a cut in radius itself, but as shown it does have a detrimental effect on the sensitivity of WASP to inflated planets. The other selection cuts do not affect the detectability of real planets, and are mostly involved in rejecting false positive systems.

Period match The region in which the periods are matched is quite restrictive: only the longer periods and larger radius planets have been detected at the correct period. This requirement on the detected models decreases the sensitivity to planets with $P > 3$ days and $R_p < 1 R_J$, but otherwise does not affect the construction of the sensitivity map.

Each selection cut is shown independently of the other cuts making up the criteria for eyeballing the planet. The probability of acceptance is therefore the product of the maps shown in Fig. 2.14 combined with the 70% chance of ORION detecting the periodic signal. The overlapping regions of the parameter maps shown in Fig. 2.14 is small and mostly defined by the $P(R_p)$ cut. A surprising feature is that it is limiting the upper range of radius sensitivity due to its role in rejecting likely false positives.

2.5.3 Combining individual sensitivity measurements

To understand the sensitivity of the WASP project rather than the sensitivity of each individual lightcurve, the sensitivity maps were combined. By generating the map from the total number of accepted and proposed synthetic systems across each planet hosting star, the complete sensitivity is mapped. Figure 2.15 shows the complete sensitivity map for the planet hosting stars. Note the decreased pixel to pixel variation compared to the maps shown earlier, as the uncertainty with the counting statistics is lower.

At radii $< 0.5 R_J$ the sensitivity drops to zero indicating that despite ORION detecting periodic signals for these synthetic systems the subsequent cuts cause the signals to be rejected. Inflated planets have a non-zero chance of being detected which is encouraging, but it is significantly reduced compared to planets with $R_p \sim 1.5 R_J$. The $P(R_p)$ cut rejects false positive systems at the potential expense of




Figure 2.14: Maps of the acceptance fraction when only including the cut listed at the top of each panel. The bottom right panel shows the fraction of synthetic models which have the period matched. The colour ranges from white at 100% acceptance to black for 0% acceptance. Note the darker region at smaller radii is caused by initially incorrectly setting the period range (see Section 2.5.2).




Figure 2.15: Combined sensitivity map for the planet hosting stars. White circles mark the positions of the confirmed WASP planets. The white dashed lines represent the Roche limits for planets with masses $0.5M_J$, $2M_J$ and $10M_J$ with higher mass lines being closer to the top left of the plot.

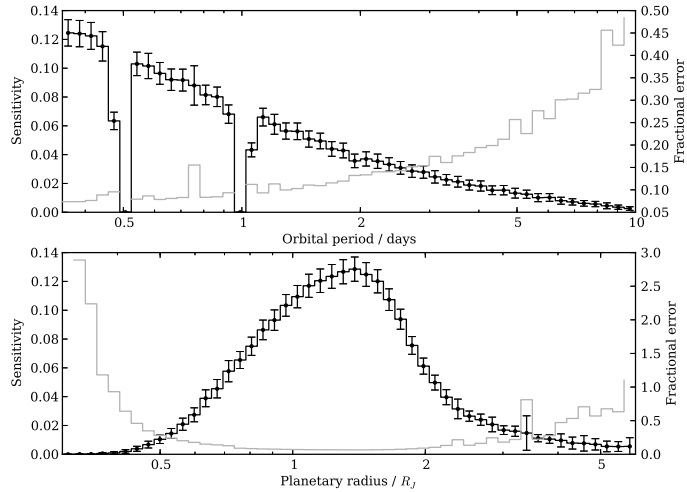


Figure 2.16: One dimensional sensitivity profiles marginalised over the second axis. Panels show the sensitivity as a function of orbital period (upper panel) and planetary radius (lower panel). The black line shows the mean sensitivity in each bin, with \sqrt{N} errorbars. The grey line shows the fractional error in each bin, and the values are given on the right axis.

detecting inflated Jupiters. The detected planets (white circles in Fig. 2.15) are mostly not located at regions of maximum sensitivity, but are often located in regions of moderate to low sensitivity. This suggests that planets are rare in regions of high WASP sensitivity (e.g. $0.5 \leq P \leq 1$ days and $1 \leq R_p \leq 2 R_J$), and common in areas of low sensitivity emphasising the need to correct for selection effects.

Figure 2.16 shows the sensitivity map collapsed separately onto the period and radius axes, and average sensitivity computed. When considering the period axis (top panel) we see a decrease in sensitivity towards longer periods, as expected from *a priori* assumptions as planets on long period orbits exhibit fewer transits for a given length of observation. The ORION period search algorithm relies on combining multiple observed transits in orbital phase to increase the detection probability. With fewer transits the sensitivity towards transiting planets decreases which is reflected accordingly in Fig. 2.16. The period cuts performed by the ORION period search are clearly visible at one day and half a day, with the average sensitivity dropping to zero. The lower panel of Fig. 2.16 shows the average sensitivity as a function of planetary radius and shows a clear peak at around $1.5 R_J$ implying that WASP is most sensitive to moderately inflated planets. The sensitivity decreases towards extremely inflated planets due to the $P(R_p)$ sensitivity cut, chosen to reject

false positives. The lower radius range also shows decreased sensitivity reflecting the lower significance of shallow transits. Both sensitivity profiles indicate a peak detection probability of around 0.13, but collapsing the sensitivity in this way does not reflect the two dimensional nature of the sensitivity which has a peak detection sensitivity of around 0.78 (Fig. 2.15) showing the value of studying the sensitivity in two dimensions.

2.6 Summary

By inserting synthetic transiting systems into WASP data and analysing the recovered systems, we have mapped out the WASP selection effects to a range of hot Jupiter planets. We find that the individual maps for each planet are similar, showing typically a decreasing sensitivity with planetary radius and orbital period, with a sharp drop in sensitivity above $2 R_J$ from the $P(R_p)$ selection cut. Each sensitivity map has maximum sensitivity at around 0.75 days and at $1.5 R_J$. The per-planet sensitivity maps do however show differences, in particular with the range of radius sensitivity. We have shown that this reflects the variable data quality of WASP rather than properties of the underlying star. We find that the known planets are not located in the peak sensitivity regions, suggesting that the population of planets missed by WASP may be large. In chapter 3 the analysis technique described here is extended to encompass an arbitrary stellar sample, drawn from the WASP catalogue. We build a catalogue of the complete set of stars around which WASP could have detected planets, sample them with a similar method, and determine the underlying population of hot Jupiters.

Chapter 3

Determining the underlying population of hot Jupiters using the WASP survey

In the previous chapter the sensitivity of WASP was calculated for the planet hosting stars. The method involved analysing individual lightcurves in detail, synthesising alternative planet configurations, analysing the synthetic lightcurves with the standard WASP pipeline and assessing their detectability. I described the process of measuring the detection probability across a range of orbital period and planetary radius values, and the reasons why the regions of high or low sensitivity existed.

In this chapter I discuss expanding the list of target stars to contain any star around which a planet could have been detected through its transit. I describe the process of synthesising the required stellar parameters for these target stars, as they have not had the detailed spectroscopic characterisation that the candidate planet hosting stars have. I define how we determined this sample, and the modifications needed to study the selection effects for these stars in Section 3.1. Investigation of the shape of the sensitivity map and how it is shaped is described in Section 3.2. The inclusion of the probability of transit is also discussed, which allows the underlying distribution of hot Jupiters to be determined. I consider aspects of the analysis process and how it affects the detectable signal in Section 3.4. We used the results from the Kepler project to constrain the population of planets around solar type stars, and propose new models for the giant planet population in Section 3.5.

3.1 Modifications for analysing the full WASP stellar sample

To understand the selection effects of the WASP project rather than just the planet hosting stars, analysis of the full sample of 37 million lightcurves was required. It is impractical to study the whole dataset, so a representative sample was taken which is believed to represent the WASP data in which a planetary signal could conceivably be detected.

3.1.1 Composing a sample of stars

We built a complete sample of stars for which we could have detected a planetary transit. To reduce the number of targets, some assumptions were made. We cannot find planets around stars fainter than a critical V magnitude. The Poisson uncertainty in each brightness measurement would mask even the largest transit depths, and fainter stars are more difficult to follow up with radial velocity. We excluded the bright stars to avoid including saturated objects for which the lightcurve characteristics are atypical. The spectral type was restricted to only include FGK stars as these are the bulk of targets in the WASP archive. Likely giants were not studied as they would be rejected at the eyeballing stage.

V magnitude The faint V magnitude limit was set by considering the known transiting planet sample. The target planets were only found around stars with V magnitudes brighter than 12.8 so we restricted our input catalogue to this range. This is advantageous also since there are exponentially more faint stars than bright stars, and we would potentially have wasted a lot of effort analysing stars for which the photon noise would be considerably higher. The saturation point for WASP is variable and sensitive to sky background and focus levels. We used the known planets sample again to restrict the magnitude range to stars fainter than 9.3.

Spectral type We also made a restriction on the J-H colour of the stars which is a proxy for spectral type. The WASP project is sensitive to predominantly FGK stars. Later type stars tend to be too faint in the WASP bandpass. Earlier type stars are observed in large numbers but planets around these stars have shallower transits, and they are not well suited to radial velocity follow up as they are fast rotators. We used the J-H colour to infer spectral type, and used the known planets to restrict the range. No WASP planets have been found around stars with J-H colours outside the range of $0.17 \leq (J-H) \leq 0.49$ which maps to the effective temperature range of




Figure 3.1: Dwarf probability used in the classification of stars for the WASP reduced stellar input catalogue.

$5000 \leq T_{eff} \leq 6500$, covering late K stars to early F stars (see Section 3.1.2 for an explanation of how the effective temperature was derived from J-H colour.) We excluded WASP-33 b from calculating this range as it orbits an A5 star [Collier Cameron et al., 2010].

Dwarf probability Another cut was made on the dwarf probability, calculated empirically from proper motion information. Giant stars are a common contaminant for exoplanet surveys as they have similar photometric colours and brightnesses than main sequence stars, but are larger leading to shallower transits. They can be identified by their lack of proper motion as they are significantly more distant than dwarfs of the same magnitude so exhibit substantially smaller proper motions. We followed the method of Collier Cameron et al. [2007a] which is performed in the WASP analysis pipeline.

The classification was modified from Gould & Morgan [2003] but using the J-H colour index, and proper motions from the USNO-B1.0 catalogue [Monet et al., 2003]. A subsample of 2000 FGK stars with high quality high resolution spectra for which the surface gravity had been determined to better than ± 0.1 dex were used to calibrate the classification system, with dwarfs exhibiting $\log g \geq 4.0$ and giants $\log g \leq 3.0$. The stars were cross-matched in the USNO-B1.0 catalogue and their J-H colour. The number density of dwarfs and giants were independently placed in reduced proper motion H_J versus J-H space. The fraction of dwarfs to total objects in each bin was calculated to give the dwarf probability, shown in Fig. 3.1. The known WASP exoplanets all have a dwarf probability of $> 98\%$ suggesting that

Table 3.1: Parameter cuts used to restrict the WASP stellar sample.

Parameter	Range
V magnitude	12.8 - 9.3
J-H	0.16 - 0.49
Dwarf probability	> 98%
Number of data points	> 1000

known planets are only detectable around very likely dwarfs. We excluded stars with a dwarf probability lower than this value.

Additionally to the cuts outlined above a cut was performed on the number of data points in the lightcurve. Objects with fewer than 1000 points are rejected by the selection cuts, and so we do not simulate these poorly sampled systems. These cuts are summarised in Table 3.1 and reduce the input stellar catalogue from around 37 million to 326620, a much smaller sample. Figure 3.2 shows the distributions of the parameters of the first 40 known planets used to reduce the stellar sample, along with the full stellar sample in the WASP catalogue. The J-H distribution peaks at around 0.25 indicating that the typical target stars are similar in effective temperature, and therefore spectral type. The underlying sample distribution shows that we are selecting the bluer stars preferentially and are not including some of the redder stars but these are likely to be much fainter in the WASP bandpass so will be rejected most likely through the V magnitude cut. Alternatively the red stars are distant giants and are rejected by the dwarf probability cut. The V magnitude distributions match well down to $V = 12$ but the full sample continues to around $V = 13$ when it starts to decrease, suggesting that the known planets are preferentially found around the brighter stars, as expected due to increased difficulty in detecting the transits, and radial velocity follow up constraints. The dwarf probability of the selected stars has a strong peak near 1.0. The full sample shows a strong peak with a low but flat distribution between 0 and 1, with a slight increase at 0 showing that the number of objects for which the probability is neither exactly 0 or exactly 1 is low. We are deliberately rejecting giant stars for our sample, and the lower cut of 98% rejects uncertain objects.

3.1.2 Calculating unknown stellar parameters

In Chapter 2 the target set of lightcurves were measured from planet hosting stars. As part of the planet candidate confirmation procedure these stars were spectro-

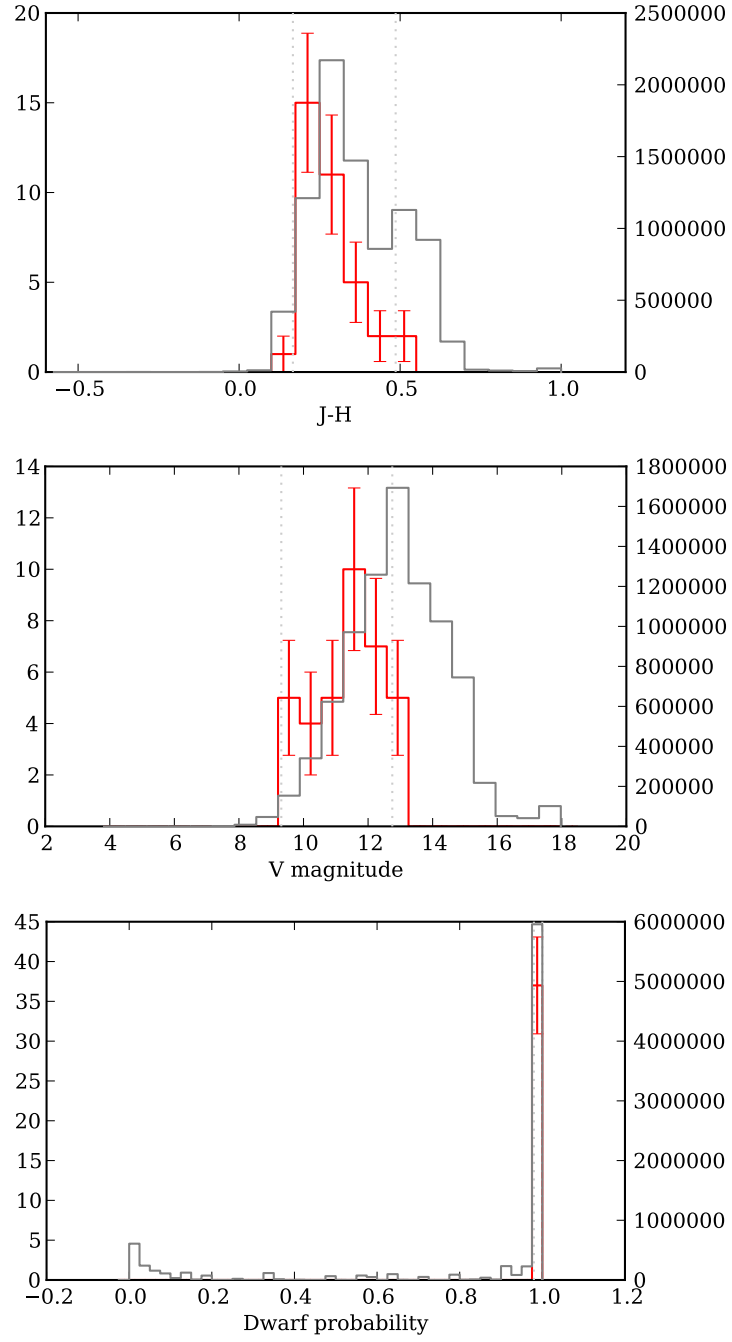


Figure 3.2: Distributions of the stellar parameters used to reduce the input sample. Red lines show the known planets' values, grey lines show the distribution for the full WASP sample. The left axis represents the known planets sample, the right axis represents the full sample. Grey dashed vertical lines represent the cuts made.

scopically characterised to rule out false positive systems.

We followed the method of Collier Cameron et al. [2007a], who used J-H colours to estimate the stellar parameters, using a sample of 65000 Tycho-2 FGK dwarf stars from Ammons et al. [2006]. Temperatures were spline-fitted to broadband Tycho-2 [Høg et al., 2000] and 2MASS [Skrutskie et al., 2006] photometry using a training set of stars observed with the Keck High-Resolution Echelle Spectrometer. The relation between the effective temperature and J-H colour was fitted as

$$T_{eff} = -4369.5(J-H) + 7188.2 \quad (3.1)$$

with an rms uncertainty of 114K. The stellar radius was then calculated from a polynomial fit to the main sequence temperature/radius relation tabulated in appendix B1 of Gray [1992]:

$$\begin{aligned} \frac{R_{\star}}{R_{\odot}} = & -3.925 \times 10^{-14}(T_{eff})^4 + 8.3909 \times 10^{-10}(T_{eff})^3 \\ & - 6.555 \times 10^{-6}(T_{eff})^2 + 0.02245(T_{eff}) - 27.9788. \end{aligned} \quad (3.2)$$

The stellar mass was then estimated via the main-sequence mass-radius relationship

$$\frac{M_{\star}}{M_{\odot}} \sim \left(\frac{R_{\star}}{R_{\odot}} \right)^{1.25}. \quad (3.3)$$

The WASP catalogue was mostly cross matched with the 2MASS catalogue [Skrutskie et al., 2006] for the J-H colour, which allowed the stellar parameters to be estimated. We note that the quality of the sensitivity estimates do not depend on these values being accurate for individual systems. We did this to ensure an overall consistency of assumed stellar parameters with the population of stars in the study.

3.1.3 Analysis pipeline modifications

To allow target stars to be chosen at random from the WASP reduced stellar input catalogue, the implementation discussed in Section 2.3.7 was altered. Only the additions custom written for this analysis were modified. The modifications discussed in Section 2.3.7 remained to prevent this analysis from biasing the normal operation of the WASP planet hunting tasks, but otherwise the analysis pipeline remained unchanged. The WASP data are structured in large files containing the photometry and associated metadata for a field. In the initial phase of this study a single lightcurve (containing a true transit-like signal) was duplicated thousands of times and studied in great detail. As we wanted to sample a much larger collection

of objects it was impractical to give each single object the same scrutiny. Instead a broad search was performed consisting of a lot of targets studied in less detail. As the target objects have been filtered to a representative sample of WASP stars, the techniques are equivalent.

The WASP lightcurves are stored as fields, all of the data for one pointing for one camera for one season. Each field had the same number of synthetic transits inserted. More synthetic objects were proposed per field than there were valid target lightcurves, so each lightcurve was duplicated multiple times and different synthetic transits inserted into each. We considered the possibility that a trend in the sensitivity may exist with brightness, which was not apparent with the planet hosting stellar sample, and that this may require accounting for in the final analysis. When the target stars were chosen per field we ensured that an equal number of synthetic lightcurves were inserted to the bright stars as were the faint stars, which created consistent uncertainty across the brightness range. This was achieved by grouping the target stars into three brightness bins of equal size in V magnitude, and drew an equal number of target stars from each bin.

Once a field file was altered, the analysis procedure followed the same path as described in the previous chapter (steps 3 - 5 in Section 2.3.7.) The detrending tools were applied and the ORION search was performed to search for periodic transit-like signals. Our MCMC parameter fitting code was then run over the detected lightcurves to refine the system parameters. The sensitivity maps were constructed with the same selection cuts as previously. It no longer makes sense to analyse the selection map on a per-object basis as the signal to noise is likely very low. Instead the ensemble sensitivity map, and the components that define it will be discussed.

A total of 16698754 objects were proposed, and 3418201 synthetic objects were inserted, into 213 separate field files giving 16038 objects per field file. This is the same density of objects per field file, but split over a total of 125304 objects instead of 40 covering a much wider range. Figure 3.3 shows the distribution of the number of synthetic objects inserted per lightcurve. The average number of synthetic systems per lightcurve is 27, but some objects were studied in much more detail. These are the brightest objects given that the magnitude range was split and an equal number of objects inserted for each magnitude group.

3.2 Shaping the sensitivity map

To determine how regions of sensitivity are shaped in the sensitivity map, two measurements were made: the regions where ORION detects objects, and the areas




Figure 3.3: Distribution of the number of synthetic transiting systems inserted per lightcurve.

of acceptance based on the selection cuts. The ORION detection map is shown in Fig. 3.4 (for comparison, see Fig. 2.13). The map shape is similar to that of Fig. 2.13, ORION is proficient at finding larger objects and has a decreasing sensitivity towards smaller planets, which is slightly period dependent.

For the planet hosing star sample (Chapter 2) the acceptance fraction across the map was calculated for each selection cut in turn. This gave regions where the selection cuts provided regions of rejection for input systems. This analysis was repeated for the randomly drawn stellar sample and is shown in Fig. 3.5. Similar to the previous analysis, the main restriction for the detection is $P(R_p)$, which restricts the large radii. Each map has a decreasing acceptance fraction towards smaller radius which is the detection sensitivity of ORION (Fig. 3.4) coupled with the strength of the transit signal in the data.

Both $(S/N)_{red}$ and $(S/N)_{ellipse}$ show a slight decrease in acceptance with $P \approx 0.5$ days and $R_p \approx 2 R_J$ as these lightcurves will have frequent large transits which would appear similar to ellipsoidal modulation or red noise and so the S/N of these lightcurves would be above the threshold.

These maps suggest that only the lack of detection from ORION, and the $P(R_p)$ significantly change the shape of the sensitivity. Conversely each selection cut does not let every synthetic model through, as the mean level of each acceptance map is < 1 therefore each selection cut has the effect of reducing the sensitivity across the parameter space. The selection cuts were imposed by the WASP planet hunting pipeline to reduce the number of false positives, and it is encouraging that the selection cuts mostly do not restrict true planets, with the exception of some inflated planets.




Figure 3.4: ORION detection map. The colour scale represents the fraction of models recovered by ORION.

3.3 Sensitivity maps

We can now compare the sensitivity maps generated from the two different stellar samples to assess how representative the planet hosting stars are. Figure 3.6a shows the sensitivity map for the randomly drawn stellar sample. For ease of comparison, the sensitivity map generated from planet hosting stars is shown in Fig. 3.6b. An initial assessment shows the two maps look similar, which is encouraging as it suggests that the lightcurves of stars found to host planets are not significantly higher quality than the bulk of the WASP lightcurves. Instead these data were representative of the data collected across many similar stars.

Figure 3.7 shows the ratio for each bin between the map made from transiting planet lightcurves, and the map made from a random sample from the WASP reduced stellar input catalogue. A histogram of values is shown in Fig. 3.8 with a median value of 1.002 ± 0.005 , consistent with 1 suggesting that there is no difference between the data drawn from the known planet sample, and the data drawn randomly from the stellar WASP reduced stellar input catalogue. The distribution width is $\sigma = 0.42$ suggesting that despite a near unity median value the distribution of values is dominated by regions of the sensitivity maps where the fewest simulations were run, at long periods and large radii.






Figure 3.5: Maps of the acceptance fraction for the randomly drawn stellar sample when only including the cut listed at the top of each panel. The bottom right panel shows the fraction of synthetic models which have the period matched. The colour ranges from white at 100% acceptance to black for 0% acceptance.



(a)



(b)

Figure 3.6: Sensitivity maps created from the two different WASP stellar samples. Top: the sensitivity map for the representative sample of WASP stars. Bottom: the sensitivity map for the planet hosting stars only.




Figure 3.7: Ratio of the sensitivity map generated from using exclusively the data containing real planets to the sensitivity map generated from the random field choice.




Figure 3.8: Histogram of values in the map ratio > 0 . The median value for the dataset is 1.002 ± 0.005 .

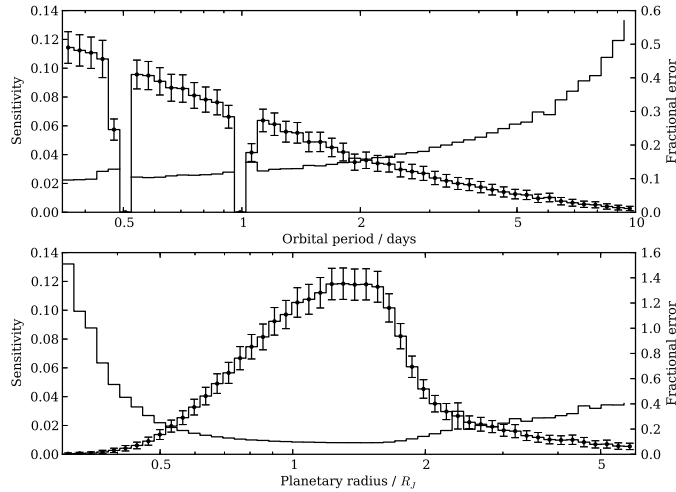


Figure 3.9: Flattened sensitivity profile.

Figure 3.9 shows the sensitivity profiles for the period and radius axes collapsed in each dimension in turn, for comparison with Fig. 2.16. The shape appears similar which is expected given how similar the 2D maps are.

Brightness comparison

We compared the three different magnitude ranges sampled. Figure 3.10 shows three sensitivity maps, one created for each magnitude range. On initial inspection the maps are similar in shape, suggesting broadly similar behaviour across the entire brightness range. Despite the overall similarity, some differences are apparent. The map of brightest objects shows the widest range of radii with sensitivity > 0.3 , and also a lack of very large ($\sim 2 R_J$) objects at very short periods. This is likely caused by the $(S/N)_{\text{ellipse}}$ cut as shown in Fig. 3.5. It is rejecting the large short period objects around bright stars. The sensitivity map for the faintest objects exhibits a narrower radius range with decreased sensitivity to Saturn sized planets at periods longer than two days. The peak sensitivity is also higher than the brighter objects. The sensitivity map for the middle magnitude range has the highest normalisation of sensitivity.

All maps have similar peak sensitivity values, to within 0.07 of each other suggesting consistency in the WASP dataset. We determine that the differences between the maps are not significant enough to require separate treatment.




Figure 3.10: Sensitivity maps built up by splitting the magnitude range of study into three. The magnitude range is quoted at the top of each figure. See Fig. 2.15 for a description of the features.

3.3.1 Incorporating probability of transit

During the model system generation stage the proposed non-transiting systems were recorded along with the transiting systems which were subsequently simulated by altering the WASP lightcurves. To include the probability of transit in the sensitivity maps the number of non-transiting systems were included in normalisation. The sensitivity then becomes

$$f_{P,R_p} = \frac{n_{detect}}{n_{ins} + n_{nt}} \quad (3.4)$$

where n_{detect} are the systems recovered by the WASP detection and analysis pipeline, n_{ins} are the inserted transiting systems and n_{nt} are the proposed but non-transiting systems (see Eq. 2.11). This method is preferable to calculating the probability of transit across the sensitivity map


$$p_{trans} = \frac{R_p + R_\star}{a} \equiv (R_p + R_\star) \left(\frac{GM_\star}{4\pi^2} P^2 \right)^{-1/3} \quad (3.5)$$

as this will be different for each object due to the differing stellar parameters per star. By counting the non-transiting systems the probability of transit is appropriate for the population taken as a whole.


The sensitivity map incorporating the probability of transit is shown in Fig. 3.11a for the target sample drawn from the WASP reduced stellar input catalogue. Compared to the sensitivity map without the probability of transit (Fig. 3.11b) the peak absolute sensitivity level drops by half, and the long period sensitivity is decreased significantly. This naturally matches the probability of transit decreasing with increasing period as $P^{-2/3}$. This new measurement allows proposed underlying distributions of hot Jupiters to be assessed and the occurrence rate to be calculated.

3.4 Comparing input systems to detections

To assess the effect of detrending on the transit shapes, the input radius for each model was compared to the posterior radius computed by our MCMC implementation. Figure 3.12 shows a contour map of the input radius against the fitted radius. The MCMC analysis recovers a systematically smaller radius than the value input, possibly meaning the detrending itself is distorting the transit shape. Towards larger input radii after around $2 R_J$, the MCMC algorithm tends to fit a smaller radius biasing the detection away from larger planets. TFA is a more aggressive detrending process than SYSREM, and has a detrimental effect on the transit shape itself, generally decreasing the transit depth as indicated by Fig. 3.12. Commonly during the



(a)



(b)

Figure 3.11: Upper: sensitivity map incorporating the probability of transit for the target sample drawn from the WASP reduced stellar input catalogue. Lower: sensitivity map without incorporating the probability of transit, for comparison.




Figure 3.12: Ratio of the planetary radius fitted by the MCMC algorithm against the input planetary radius inserted into the lightcurve. The dashed red line marks the unity point.

eyeballing stage the non-TFA lightcurve is used to validate any transit shape.

The MCMC analysis does not fit impact parameters greater than 1, with the explanation that only solutions yielding a transit of observable depth are interesting from a planet hunting perspective [Collier Cameron et al., 2007a]. Figure 3.13 shows a comparison between the impact parameter of the recovered synthetic systems against the best fit impact parameter from the MCMC analysis. Due to a large number of data points the number density is used to represent the values, with the density displayed on the colour scale. The impact parameter fitted by the MCMC analysis, b_{mcmc} is limited to the range $0 \leq b_{mcmc} \leq 1$ but this does not prevent input models with higher impact parameters b_{in} to be recovered. Across the entire range of b_{in} the recovered impact parameter is overestimated with a higher value detected than inserted. The MCMC routine assumes a prior uniform distribution of b starting at 0.5. Input impact parameters $b_{in} < 0.5$ are overestimated, sometimes considerably, but above this the match is better. This could be an approach to improving the MCMC fitting routines.

3.5 Constraining the underlying hot Jupiter population

To investigate the underlying hot Jupiter population the sensitivity of WASP was used to invert the population of observed planets to account for real planets missed either because they are non-transiting or because of the lack of sensitivity of WASP. This population is the underlying distribution of hot Jupiters, and we use this to




Figure 3.13: Comparison of the input and detected impact parameters from the analysis. The red line indicates the same value for input and output.

calculate the occurrence rate. As we have used the observed distribution of WASP planets to determine the underlying distribution, this calculation is only valid within the region of parameter space in which a WASP planet has been detected. We therefore define the result as the underlying distribution of hot Jupiters up to periods of 10 days and planet radii between 0.7 and 2 R_J , covering the hot Jupiter parameter space only.

We use the list of confirmed planets with references in the literature as our planet sample. Each observed planet has $1/f_{P,R_p}$ unseen companions where f_{P,R_p} is the sensitivity to a planet with orbital period P and radius R_p . To calculate the underlying distribution a new population of planets was created from the number of observed planets and the number of planets not detected. The distribution of the number of missed planets per star is shown in Fig. 3.14 and implies that each observed planet represents a sometimes large population of undetected objects. The median number of extra planets is 20 and the number of extra planets based on transit probability alone for a typical hot Jupiter around a solar type star ($a = 0.05$, $R_\star = R_\odot$) is 13 reflecting the incomplete sensitivity of the project itself.

Figure 3.15 shows the distributions of orbital period and planetary radius before and after correcting for the selection effects of WASP. The underlying orbital period distribution suggests that longer period planets are more numerous than




Figure 3.14: Correction factors calculated by inverting the WASP planets through the sensitivity map.

is evident from the observed distribution. The radius distribution shows a large correction performed for the smaller radii reflecting the limits of sensitivity for such planets.

The Kepler candidates have been used for a study of the planet occurrence by Howard et al. [2012] for planets orbiting within 0.25 AU of solar-type stars. WASP has detected twice the number of Kepler candidate hot Jupiters and so is in the best position to understand their population, but has no sensitivity to longer period planets so we combine datasets to calculate a joint constraint.

3.5.1 Joint constraint with Kepler

The Kepler project has been extremely beneficial for the exoplanet community, providing candidate planets ranging from Earths to hot Jupiters. Howard et al. [2012] calculate a planetary abundance rate using the planet candidates and some stringent restrictions on the lightcurve signal to noise estimates and stellar input catalogue. The stellar catalogue is limited to solar type stars only, with the restrictions used listed in Table 3.2. The number of stars around which each planet candidate could be found is calculated by detection significance calculations incorporating the Kepler combined differential photometric precision (CDPP), a measure or prediction of the photometric precision for a set of timescales [Gilliland et al., 2011]. With a measure of the number of candidates detected and the number of stars around which that planet could have been detected, the occurrence rate of planets around solar type stars within 0.25 AU was determined. The parameter distributions for the period

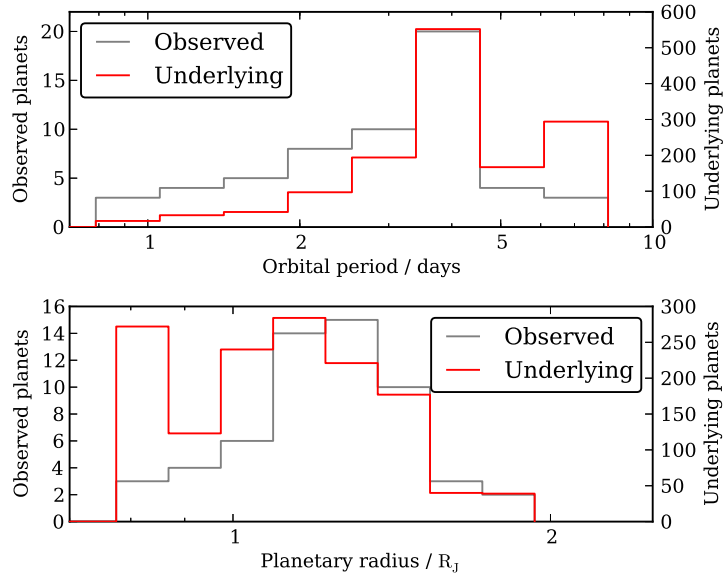


Figure 3.15: Comparison of the distributions of (top) orbital period and (bottom) planetary radius for the observed distribution (grey) and corrected distribution (red) of WASP planets.

Table 3.2: Properties of the Kepler stellar and planetary sample. From Howard et al. [2012].

Parameter	Value
Stellar effective temperature	4100-6100K
Surface gravity	4.0-4.9
Kepler magnitude	< 15
Number of stars	58041
Orbital period	< 50 days
Planet radius	2-32 R_{\oplus}
Detection threshold, SNR (90 days)	> 10
Number of planet candidates	438

Table 3.3: Best fit parameters of the period model. Taken from Howard et al. [2012].

$R_p(R_\oplus)$	k_P	β	P_0 (days)	γ
2-4	0.064 ± 0.040	0.27 ± 0.27	7.0 ± 1.9	2.6 ± 0.3
4-8	0.0020 ± 0.0012	0.79 ± 0.50	2.2 ± 1.0	4.0 ± 1.2
8-32	0.0025 ± 0.0015	0.37 ± 0.35	1.7 ± 0.7	4.1 ± 2.5
2-32	0.035 ± 0.023	0.52 ± 0.25	4.8 ± 1.6	2.4 ± 0.3

and radius dimensions separately are given as

$$\frac{df(R)}{d \log R} = k_R R^\alpha \quad (3.6)$$

$$\frac{df(P)}{d \log P} = k_P P^\beta (1 - e^{-(P/P_0)^\gamma}) \quad (3.7)$$

for $k_R = 2.9^{+0.5}_{-0.4}$, $\alpha = -1.92 \pm 0.11$ and the coefficients in Eq. 3.7 given in Table 3.3. As Howard et al. [2012] split the period relation into radius groups we can directly compare the giant planet category ($8 - 32 R_\oplus$) with the confirmed WASP planets.

Figure 3.16 shows the observed WASP planets compared to the observed Kepler candidates. Shown are 59 WASP planets¹ and 24 Kepler candidates. The period distribution of the WASP planets has an increasing occurrence towards longer periods until a period of 6 days after which the number of planets drops significantly. This naturally matches the calculated drop of sensitivity as calculated by the simulations, and has in the past been attributed to selection effects [e.g. Hellier et al., 2012]. The radius distribution shows a strong peak between 11.3 and 16, but the bins below and above show significantly fewer objects reflecting the lack of sensitivity and lack of inflated objects respectively.

For each WASP planet we have a sensitivity f_i which is calculated from the sensitivity map shown in Fig. 3.11a. We used the information contained in Fig. 4 of Howard et al. [2012] reproduced here as Fig. 3.17 to ascertain the occurrence rate of different classes of planets, in particular the hot Jupiters by selecting the bins in the upper left corner representing $0.68 \leq P \leq 10$ and $8 \leq R_p \leq 32 R_\oplus$.

We plot the Kepler data represented in Fig. 3.17 in period and radius, marginalising over the other dimension in Fig. 3.18, along with the WASP occurrence rate, which is taken from inverting the observed population through the sensitivity map and normalising by the total number of stars in the WASP reduced stellar input catalogue, $N_\star = 326620$. The data shown in the Kepler period distribution

¹taken from <http://exoplanets.org>

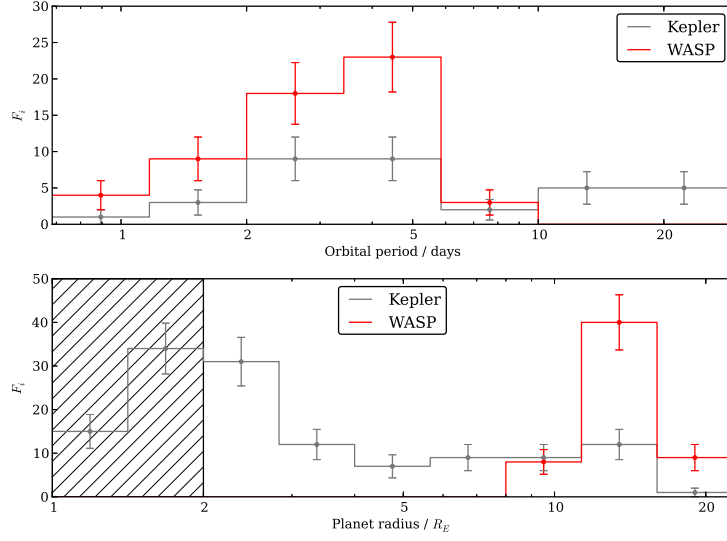


Figure 3.16: Confirmed WASP planet, and Howard et al. [2012] Kepler candidate distributions. The hatched region represents where Howard et al. [2012] believe to be incomplete. The upper panel shows the period distribution for planets $8 \leq R_p \leq 32 R_{\oplus}$. The lower panel shows the radius distribution for planets $P \leq 10$. Bin sizes are chosen to correspond with Howard et al. [2012].

(top panel) are from the giant planet ($R_p \geq 8 R_{\oplus}$) category. The data shown in the Kepler radius distribution (bottom panel) have orbital periods ≤ 10 days. As WASP has detected more giant planets than Kepler the uncertainties in each bin are lower allowing tighter constraints to be placed on the planet occurrence.

Both WASP distributions are consistent with the Kepler results from their candidate planets. We note that the consistency of the normalisations is likely a coincidence as the false positive rate for Kepler especially for giant planets is around 17.7% [Fressin et al., 2013], and WASP continues to find new planets whilst the stellar sample does not increase at the same rate which increases the WASP normalisation. We continue to treat the two datasets as consistent for this analysis, but this is considered further in Section 3.6. The Kepler period distribution shows a dip between 5.84 and 10 days which is not significant but is consistent with a more strongly detected dip in the WASP data. The consistency again suggests that the selection effects are well corrected for in this region. When we consider the radius distribution (lower panel of Fig. 3.18) we see consistent results for the two datasets. The absence of planets in the $16 \leq R_p \leq 22.6 R_{\oplus}$ bin likely reflects the rarity of inflated planets.

With smaller uncertainties per bin in the WASP data due to higher total objects, the number of bins can be increased to provide more detail about the




Figure 3.17: Kepler candidate planet occurrence as a function of planet radius and orbital period for $P < 50$ days. White numbers indicate the values represented in each cell: the upper left number gives the number of candidates detected in each bin, with the number corrected for geometric transit probability in parentheses. The lower left number gives the number of stars around which a typical planet in the bin could be detected. The lower right number gives the ratio of these values and represents the planet occurrence in the bin. The upper right number gives the planet occurrence per logarithmic area unit. Reproduced from Howard et al. [2012].

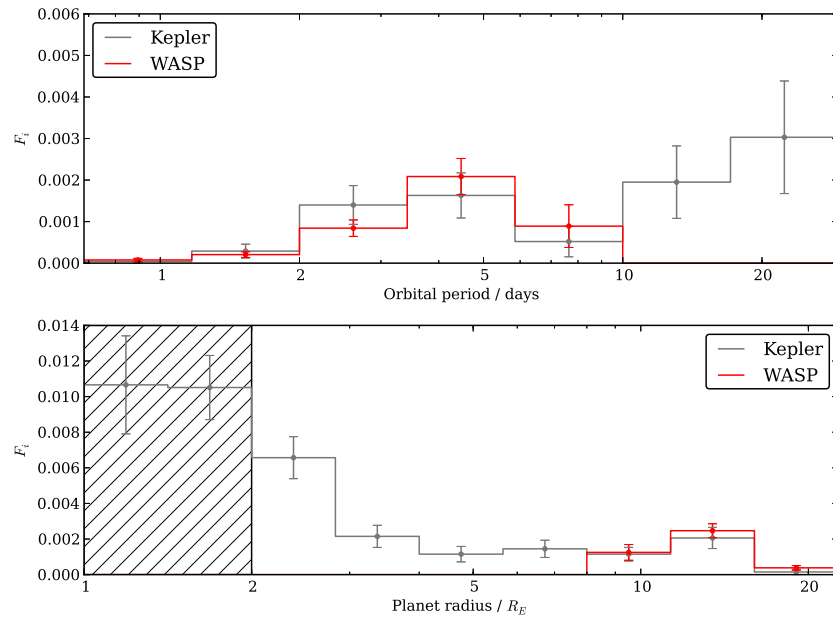


Figure 3.18: Occurrence rate in planets per star showing the Kepler and WASP results. The Kepler period distribution in the top panel is from the giant planet ($R_p \geq 8 R_\oplus$) category, the Kepler radius distribution is planets with orbital period $P \leq 10$ days. The grey line is the measured distribution from Howard et al. [2012], the red line is the corrected distribution from WASP. The data have been binned to the same resolution for comparison.

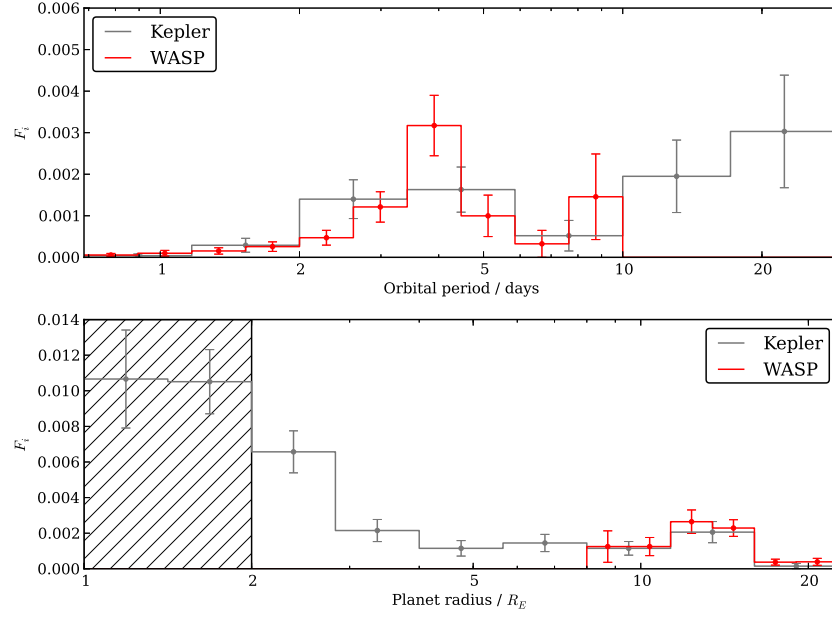


Figure 3.19: The occurrence rate in planets per star. The lines have the same meaning as in Fig. 3.18. The WASP data have been binned to twice the resolution, and the occurrence rates multiplied by two for comparison.

structure of the underlying distribution, without sacrificing precision. Figure 3.19 shows the same data with twice the number of WASP bins. Note the uncertainties are still comparable with the Kepler bin uncertainties. We observe a much higher peak of the number of hot Jupiters orbiting with periods $3.42 \leq P \leq 4.47$ days than the Kepler measured distribution, with a deficit at longer periods, but the radius distribution does not show extra structure after increasing the bin resolution. We now demonstrate that the functional form proposed by Howard et al. [2012] is not sufficient to model the combined giant planet distributions.

3.5.2 A new model for the underlying period distribution

The WASP project is only sensitive to hot Jupiters and so cannot place constraints on the longer period planets. We therefore use the combination of Kepler results and WASP results to estimate the functional form of the underlying distribution of giant planets.

To ensure consistency between our technique and that of Howard et al. [2012], the Kepler data alone were re-fitted with a Levenberg-Marquardt algorithm [Press, 2007]. To ensure the same local minimum was explored the best fit parameters from Howard et al. [2012] were used as the starting point of the fit. These parameters




Figure 3.20: Planetary occurrence marginalised over radius for three planet classes: super Earths (orange), Neptunes (green) and giant planets (blue). The lines show the functional form shown in Eq. 3.7 with best fitting parameters listed in Table 3.3. From Howard et al. [2012].

are given in Table 3.3. The upper panel of Fig. 3.21 shows the Howard et al. [2012] fit along with the repeat fitting. The maximum likelihood parameters were found to be consistent as is apparent in the upper panel.

To constrain the fit further, the Kepler data were then combined with the WASP data and re-fitted using the same technique. The lower panel of Fig. 3.21 shows the result of adding the WASP data. We find a χ^2 of 33.2 with 14 degrees of freedom and reject the null hypothesis that the original Howard model is valid at the 1% level. The same functional form is clearly not a good approximation to the binned occurrence rate. The effect of the pile up at four days constrains the fit around that region and it is no longer well-suited at longer periods. When allowing the maximum likelihood solution to be computed (the red line in Fig. 3.21) we find a χ^2 of 21.9 with 14 degrees of freedom suggesting that we cannot formally reject it based on statistical arguments alone, but the model is not physical as it would imply decreasing planet numbers at longer periods which is not observed, for example the radial velocity and direct imaging planets.

We propose a new functional form where the occurrence rate is composed of two terms: a general increase in planets towards longer periods as observed by the Kepler team, together with a Gaussian excess of hot Jupiters between 3-5 days.




Figure 3.21: Top: the occurrence rate in each bin measured by Howard et al. [2012] shown with their best fit model (green) and the best fit model found by reanalysis (red). Bottom: the occurrence rate with the addition of the WASP data, with the previous best fit model from Howard et al. [2012], and the best fit model found by reanalysis (red).

The form used is

$$\frac{df(P)}{d \log P} = k_p P^\beta + \frac{\gamma}{2\sigma_p \sqrt{2\pi}} e^{-((P-P_0)/\sigma_p)^2}. \quad (3.8)$$

where k_p and β are named as such to be consistent with Eq. 3.7 as this term remains the same. This functional form was fitted to the combined WASP and Kepler dataset and is shown in Fig. 3.22, providing the best fit parameters shown in Table 3.4, giving a $\chi^2 = 6.91$ for 13 degrees of freedom, and a p -chance of 0.907. This is a better fit than the Howard et al. [2012] proposal ($\chi^2 = 21.9$, 14 degrees of freedom, p -chance of 0.08) and fits the hot Jupiter pile up well.

Table 3.4: Best fit parameters for the power law model with Gaussian excess calculated from minimising the least squares.

Parameter	Value
k_p	0.00029 ± 0.00009
β	1.15 ± 0.12
γ	0.0185 ± 0.0026
P_0	3.81 ± 0.13
σ_p	0.83 ± 0.11




Figure 3.22: Occurrence rates in planets per star for giant planets incorporating Kepler (grey) and WASP data shown on a logarithmic axis. The best fit model for Eq. 3.8 is shown in blue. The residuals normalised to the uncertainty are shown in the bottom panel.

The methods described above assume the count in each bin is enough to be able to approximate the uncertainty through the normal distribution, but as suggested by Fig. 3.16 this is not entirely the case. We therefore use the Cash statistic and aim to minimise

$$C = 2 \sum_{i=1}^N (e_i - n_i \ln e_i) \quad (3.9)$$

[Cash, 1979] where e_i is the expected value in bin i , and n_i the observed value, derived from the log-likelihood ratio of the Poisson distribution. When applying this statistic and minimising through application of the Nelder-Mead simplex minimisation, the best fit parameters are shown in Table 3.5 with 1σ confidence limits calculated through the bootstrap method, and resulting fit and residuals in Fig. 3.23.

A free normalisation scaling parameter was added between the WASP and Kepler datasets to account for affecting the normalisation of the WASP fit, and was calculated to be 0.94 suggesting that the WASP data indicate a 6% higher overall occurrence rate.

Table 3.5: Best fit parameters for the model (Eq. 3.8). The previous values from Table 3.4 have been repeated for comparison.

Parameter	Previous value	Cash value
k_p	0.00029 ± 0.00009	0.00011 ± 0.00007
β	1.15 ± 0.12	1.01 ± 0.25
γ	0.0185 ± 0.0026	0.018 ± 0.003
P_0	3.81 ± 0.13	3.73 ± 0.12
σ_p	0.83 ± 0.11	0.72 ± 0.19

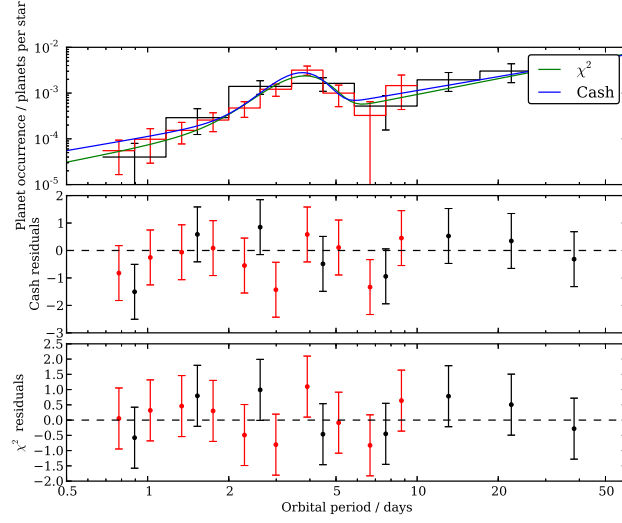


Figure 3.23: Fits to the occurrence rate of giant planets with simultaneous fit to Kepler and WASP data. The top panel shows the occurrence rate along with the fit using the least squares method (green) and minimising the Cash statistic (blue). Grey points are the Kepler data, and red points the WASP data. The middle panel shows the χ^2 residuals for the Cash statistic. The bottom panel shows the χ^2 residuals for the least squares fit.

Table 3.6: Best fit parameters from applying the Cash statistic for the underlying radius distribution, with 1σ errors calculated through the bootstrap method.

Parameter	Value
k_R	0.62 ± 1.05
α	-2.03 ± 0.36
μ	13.29 ± 0.62
σ	1.91 ± 0.56
γ	0.066 ± 0.020

3.5.3 Investigating the radius distribution

We followed the same method to investigate the radius distribution of giant planets, each step is shown in Fig. 3.24. We used least squares minimisation when comparing models, as the non-Gaussian nature of the uncertainties has the effect of reducing the χ^2 , but an improved model can still be determined. We applied minimisation of the Cash statistic however to estimate the best fit parameters.

First we fitted the Howard et al. [2012] relation described in Eq. 3.6 but we selected only planets with $P \leq 10$ days to be consistent with the WASP parameter space. We used a Levenberg-Marquardt algorithm with the initial parameters $k_R = 2.9$, $\alpha = -1.92$ found by Howard et al. [2012] to find the best fit parameters and to ensure the method for analysing the period distribution was consistent. As we were comparing the planets with orbital periods $P \leq 10$ days rather than $P \leq 50$ days we expect some difference. We found the best fit parameters of $k_R = 0.61 \pm 0.31$ and $\alpha = -2.06 \pm 0.39$, and a clearly unacceptable fit statistic of $\chi^2 = 17.97$ with a p -chance of 0.003.

We then combined the Kepler data with the WASP data with half bin sizes, and the normalisation scaled up by two to perform a simultaneous fit, and found $\chi^2 = 42.26$ with a p -chance of 1.46×10^{-5} implying an even worse fit.

By inspection, the distribution may be consistent with a power law with excess at around $R_p = 15R_\oplus$. We fitted the functional form of Eq. 3.8, i.e. a power-law plus Gaussian, and found the result shown in the third panel of Fig. 3.24, with a statistically acceptable fit with $\chi^2 = 9.67$ and a p -chance of 0.289. The excess was modelled with a central value of $\mu = 13.29 \pm 0.62 R_\oplus$ and a width of $1.91 \pm 0.56 R_\oplus$ whilst the power law coefficient becomes -2.03 ± 0.36 . The full list of parameters is given in Table 3.6 where we denote the power law normalisation k_R instead of k_P and β as α as the power law term is the same as used by Howard et al. [2012].




Figure 3.24: Result of characterising the radius distribution. The blank lines represent the Kepler occurrence rate of planets with $P \leq 10$ days. The red lines represent the WASP occurrence rate scaled by two due to the difference in bin sizes. The green lines represent the best fit for the three scenarios. Top: fitting the Howard et al. [2012] proposal function to the selected Kepler occurrence rates. Second: fitting the Howard et al. [2012] proposal function to the selected Kepler and WASP occurrence rates. Third: fitting a power-law with Gaussian excess to the selected Kepler and WASP occurrence rates. Bottom: residuals to the power law with Gaussian excess.

3.6 Discussion

We have provided significant evidence that the pile up in hot Jupiters is real and not an observational artefact. The central value of the peak has been constrained to 3.7 days and the behaviour outside the peak is a rising power law. This implies that an excess of giant planets exist at a characteristic period, which must be a marker for their formation.

Disk migration is driven by torques exerted by material at resonant locations in the protoplanetary disk. The inward migration torque is predominantly provided by the most distant outer Lindblad resonant point where $m = 1$ in Eq. 1.24. The material orbiting interior to the corotation radius of the disk (the radius at which the material in the disk orbits with the same period as the stellar rotation period) is evacuated due to the experienced drag, so opens up a gap in the centre of the disk. When the outer Lindblad resonance point is interior to the inner edge of this gap, the planet no longer experiences inward torque from the disk and the migration halts. This point is when the inner disk edge is in 2:1 resonance with the planetary orbit [Lin et al., 1996; Kuchner & Lecar, 2002]. We therefore expect to see an over abundance of planets at a period half that of the rotation period at the time of migration, which is likely much faster than that seen today due to magnetic braking. The subsequent evolution of the host star has decreased the spin period so we must consider T Tauri stars. To explain the 3.7 day pile up observed in this study the rotation periods must be near 8 days, which has been found when studying the rotation rates of T Tauri stars in the Orion Nebula Cluster (ONC) [Herbst & Mundt, 2005]. The distribution of rotation rates in the ONC is shown in Fig. 3.25, though other clusters studied do not show such a clear feature indicating that this may be an interesting coincidence.

Another proposed argument for smooth disk migration leading to a pile up is where the inward migration is halted by mass loss due to Roche lobe overflow, widening the orbit [Trilling et al., 1998]. This transfer of mass provides a balancing torque to the inward torque from the disk. Once the disk dissipates the planet remains at the distance required to be Roche lobe filling, with depleted mass. This would lead to a mass deficit for planets which orbit within the pile up. Figure 3.26 shows the planetary mass against orbital period. A mass deficit does seem to be apparent around the peak pile up value calculated, and is visible especially when taking the running median of the dataset. This perhaps could be explained instead by low mass giants being evaporated close to the star, and low mass giants being undetectable to radial velocity measurements further from the star.




Figure 3.25: Rotation rates for T Tauri stars in the Orion Nebula Cluster. The vertical line represents twice the central peak value found when fitting the WASP and Kepler data. Adapted from Herbst & Mundt [2005].




Figure 3.26: Planetary mass against orbital period for planets $R_p \geq 8 R_\oplus$. The red line represents the average of the dataset when grouped into period bins. The blue line represents the median of the dataset when grouped into period bins. The vertical dashed line marks the best fit pile up centre.




Figure 3.27: Roche limit distribution for the WASP planets before (grey) and after (red) correcting for selection effects. Distributions are normalised to unit area contained in the histogram. The vertical dashed line represents the inner edge value proposed by Ford & Rasio [2006].

Ford & Rasio [2006] predict an inner edge of pile up to occur at twice the Roche limit, by equating the initial orbital angular momentum with high eccentricity $e \approx 1$ and final orbital angular momentum. We correct the WASP sample for selection effects and plot the distribution of $x = a/a_{RL}$ the ratio of observed separation to Roche separation in Fig. 3.27. By correcting for selection effects the central peak of the distribution moves further out to around $x = 3.3$ and arguably an inner edge does occur at $x = 2$. We observe a pile up rather than an abrupt cutoff, as perhaps tidal circularisation has not been strong enough to bring planets with larger Roche fractions into their expected value of $x = 2$. If the dominant explanation for the pile up is from tidal recircularisation due to conversion of orbital energy to heat in the planet, we might expect to see inflated planets in the pile up region, depending on how long it is since migration. Each successive pass through the periastron causes the planetary orbit to circularise, and causes heating of the planet, which increases the radius, and as the cooling timescale the planet is longer than the orbital period, the planet remains inflated during the remainder of its orbit. Figure 3.28 shows the radius of hot Jupiters as distributed by period. It is uncertain at this point if this is the case due to the low number of inflated hot Jupiters, but a small peak in the radius is apparent at close to the pile up periods. If this peak becomes more significant it might be evidence to point to this migration method.




Figure 3.28: Radius for hot Jupiter planets with $R_p \geq 8R_{\oplus}$. Horizontal dashed lines mark the edges of the radius bins from Howard et al. [2012]. The vertical dashed line marks the peak of the pile up.

Our model of the period distribution includes a power-law term $k_P P^\beta$ which defines the longer period planets. The calculated power law slope is consistent with the power-law slope found by Howard et al. [2012] for the middle planet class where $4 \leq R_p < 8 R_{\oplus}$. This slope was found to be (Table 3.3) $\beta = 0.79 \pm 0.50$ and we find $\beta = 1.01 \pm 0.25$, and may suggest that the period distribution of both Jupiter class and Neptune class planets is similar, apart from clear excess at 3.7 days for the Jupiters.

The excess modelled for the radius distribution is likely due to the degenerate mass radius relation for giant planets. Mass radius relations for giant planets [e.g. Fortney et al., 2007] result in a 1 R_J planet for a wide range of masses and so an excess of planets of this size is expected.

We note that the consistency between the WASP and Kepler occurrence rates is very high, and is likely a coincidence. WASP is constantly observing and detecting new planets, but the number of stars observed is not increasing at the same rate, which would have the effect of increasing the measured occurrence rate from WASP. Conversely repeat observations of the same stars provide a longer timebase of observations which allows longer periods to be investigated. Under a white noise dominated system adding more observations decreases the noise, which increases the significance of shallower transits and the detectability of smaller planets. This would alter the WASP sensitivity which would not increase the occurrence rates. In reality the true effect is likely to be between the two. The Kepler false positive rate was found to be $17.7 \pm 2.9\%$ [Fressin et al., 2013], or $18.2 \pm 6.7\%$ [Santerne et al., 2012] for giant planets with orbital periods < 10 days, which would decrease the

Howard et al. [2012] normalisation factor and cause a further discrepancy. Howard et al. [2012] note that the occurrence of hot Jupiters in the Kepler field is only 40% of that in the solar neighbourhood, in comparison with Marcy et al. [2005]. WASP typically targets stars in the solar neighbourhood which typically have solar metallicity. Conversely Kepler observe typically low metallicity stars [Howard et al., 2012]. Correlations between giant planet occurrence and host star metallicity have been found [e.g. Santos et al., 2003; Fischer & Valenti, 2005] so discrepancy in the absolute occurrence rate is expected.

Chapter 4

NGTS: Design and prototype

The Next Generation Transit Survey (NGTS) is a new ground-based wide-field survey for transiting Neptunes and super Earths. The primary goal is to find small planets around stars bright enough for radial velocity confirmation, including bright super Earths for atmospheric studies. With mass determinations from radial velocity measurements the bulk density of these objects will be calculated. We aim to find and characterise enough of these planets with density measurements to be able to perform analysis of the population. Detected NGTS planets will be brighter than similar Kepler planets and so easier to characterise using current ESO instruments (e.g. HARPS and ESPRESSO) and next generation technologies (e.g. JWST and E-ELT).

4.1 Introduction

Current ground based transiting exoplanet surveys have provided many hot Jupiters, even observing some inflated objects much larger than Jupiter (e.g. HAT-P-32 b [Hartman et al., 2011] or WASP-17 b [Anderson et al., 2010]). Space based surveys are successful at finding the smaller planets due to the increased photometric stability and precision [Borucki et al., 2010]. Space-based missions such as Kepler typically observe fainter stars than would be targeted from the ground, so the confirmation of the detected planets is more difficult. Figure 4.1 shows the distribution of the planet radius of known exoplanets and shows a distinct lack of planets between 4 and 11 R_{\oplus} . They are difficult to detect in significant numbers because they are relatively rare so space-based surveys do not contain many in their target sample, and the transit signal is low so ground based surveys do not detect them. The small planets that have been detected typically have not undergone the same




Figure 4.1: Planet radius histogram for the known exoplanets. The black line indicates planets detected from the ground, the red line indicates planets detected from space.

levels of atmospheric characterisation that has been performed on the hot Jupiters, as the atmospheric scale height of planetary atmospheres typically increases with the planet radius, so these small planets typically have thin atmospheres. Future space missions (e.g. EChO [Tinetti et al., 2012] or CHEOPS [Broeg et al., 2013]) will be able to achieve this but require targets to study.

NGTS is an international collaboration between institutions in the UK, Geneva, Berlin along with a partnership with ESO. The team bring expertise and experience from past planet search projects including WASP, HARPS and CoRoT, and a combination of wide field robotic telescopes and state of the art spectrographs. The project will provide targets for the next generation of follow up instruments with which the atmospheric studies will be performed, providing both increased number and higher precision measurements required to study these small planets. NGTS has four times the sky coverage of Kepler and so will have a larger sample of bright and small exoplanets.

Smaller planets will be detected with NGTS by observing smaller and therefore redder stars, and by making high precision observations of the stellar flux to be able to detect shallower transits. Figure 4.2 shows the known transiting planets detected from the ground in terms of the radii of the planet and host star. There is a distinct lack of confirmed planets around solar type stars with a precision between $0.1\% \leq \sigma_{frms} \leq \sim 0.5\%$. The shaded region represents the parameter space which NGTS is targeting. Compared to the other areas of the plot the number of planets is very low and so we aim to increase our understanding of the smaller planets, and




Figure 4.2: Planet radii and host star radii for planets detected from the ground. Planet parameters are taken from online catalogues exoplanets.org and exoplanet.eu. Dashed lines show the transit depth of the system. WASP planets are shown in red, other planets are shown in grey. The shaded region represents the parameter space where NGTS is targeting, representing K and G type stars with precisions down to 0.5 mmag.

the population. Dressing & Charbonneau [2013] find an occurrence rate of $0.9^{+0.04}_{-0.03}$ planets per star for stars $T_{eff} < 4000\text{K}$ suggesting a large number of potential detections for NGTS. With increased precision and smaller target stars NGTS will push the limits of planet detection towards smaller planets around smaller stars. With NGTS we aim to measure stellar lightcurves with higher precision than has been achieved before from a small wide-field survey. A target of sub-1 mmag ($< 0.1\%$) around an $I = 13$ star in an hour was set as it is required for a detection of a Neptune around a solar type star, or an Earth around an M star [Pollacco et al., 2011].

The Transiting Exoplanet Survey Satellite *TESS*, is a project led by MIT in conjunction with NASA, and has similar goals to NGTS: high precision observations of late type stars. Observing from space, the TESS satellite will observe almost the entire sky over the course of 2.5 years, targeting G and K type stars of magnitude 5 to 12, looking for nearby transiting exoplanets. The satellite’s orbit allows an unobstructed view of the sky, with 27 days of continuous observation on each field before shifting the field of view, providing high period phase coverage up to around 10 days. This increased period phase coverage, and the higher precision possible from space allow TESS to find smaller planets than NGTS, but at much higher cost (around 100 times the budget of NGTS), and is not due to begin observations

until NGTS has been operating for 3 years. The selection of TESS did not occur until shortly before the submission of this thesis.

The work described in this chapter and the following chapter was performed through application of techniques implemented by myself, to data collected from instruments the design of which I was largely not involved with. The aperture photometry pipeline was constructed from the standard software packages AUTOPHOTOM and SEXTRACTOR, and was used to extract lightcurves from images reduced by João Bento, and the detrending tool SYSREM was implemented by Richard West. The planet catch simulations (Section 5.5) included analysis performed by my supervisor and a colleague Tom Loudon, with field selection from João Bento.

4.2 Achieving the targets

The goals of the NGTS will be achieved through understanding of the sources of noise in the experiment, the location of the instrument and high precision hardware design and construction.

4.2.1 Design

Instrument

The cameras were sourced from Andor in the UK, and contain CCDs manufactured by e2v also in the UK (Camera brochure number: iKon-L 936BR-DD, CCD catalogue number CCD42-40¹). To maximise the sky area and therefore number of stars observed, the chips are large with 2048x2048 13 μm pixels, coupled with the telescope design allows each pixel to cover 24.01 square arcseconds of sky each providing a field of view of 7.77 square degrees per telescope and 93.24 square degrees for the full instrument. To observe smaller stars which tend to have a spectral energy distribution which peaks towards the red, the CCDs are deep depleted and back illuminated to maximise the longer wavelength flux. Figure 4.3 shows the sensitivity response of the full instrument, incorporating the atmospheric transmission and throughput of the telescope and filter systems. A filter has been added to include the wavelength range 5500 Å to 9000 Å.

The telescopes were manufactured by ASA in Austria, and consist of a 200 mm diameter Newtonian mirror system, with a focal length of 560 mm providing a f-number of 2.8. The low f-number allows more sky area to be observed,

¹<http://www.e2v.com/e2v/assets/File/documents/imaging-space-and-scientific-sensors/08-42-40-fi-aimo-cpak.pdf>




Figure 4.3: Throughput of the NGTS instrument (black line). Johnson-Cousins V , R , and I filters are added for reference. The grey line is the transmission of the atmosphere at airmass 1.5; the cyan line is the QE of the CCD, taken from the CCD brochure from Andor; the magenta line is the throughput of the filter and the black line the combination of these components.

and subsequently more stars, though it limits the faint magnitude limit for target stars. The scattered light entering the optical path is reduced by adding a 400 mm baffle to the front of the telescope tube. Each telescope tracks its own field by guiding from the science images, and provides the ability to track targets with sub-pixel precision [McCormac et al., 2013]. The PSF size is $12\ \mu\text{m}$ across the field of view which matches the pixel size well giving sharp in focus images, which coupled with the objective to keep stars on the same pixels during observations reduces the errors in estimating the flat field, and reduces the noise from the sky background.

The NGTS facility is an array of 12 robotic telescopes, able to operate without human intervention for extended periods of time. Each mount hosts one telescope and one camera allowing individual units to follow fields independently of the others, and allow for increased tracking precision. The mounts are fork based in design allowing for a large range of movement and pointing locations on the sky. Figure 4.4 shows a single NGTS unit, installed at the Geneva Observatory.

Facility

Through a partnership with ESO, NGTS is located on Cerro Paranal in Chile, providing the best observing conditions available due to excellent atmospheric con-




Figure 4.4: A complete NGTS unit, assembled at the Geneva Observatory.

ditions, high altitude and low humidity. For example, Fig. 4.5a shows the percentage of photometric nights is 11.74% higher than another Chilean site La Silla, and Fig. 4.5b shows the level of precipitable water vapour at both sites showing that on average Paranal is drier. It was shown that the probability of planet detection for Paranal is higher at longer periods especially during the winter months [Bento, 2012].

The facility has been designed to allow each instrument independence from all others providing freedom to observe any field without being restricted by the other instruments. Each unit contains a control computer along with a master computer to organise tasks for the telescopes to perform. These tasks range from taking calibration frames, nightly quality frames or science frames and the system is entirely automated once the desired tasks have been inserted into the master computer. The enclosure was built by GRPro Limited in the UK, and contains automatic controls for the opening and closing of the roof based on weather conditions and other night quality qualifications. Figure 4.6 shows computer renders of the facility.





Figure 4.5: Weather quality for Paranal. Left: monthly average percentage of photometric nights for La Silla and Paranal. Data taken from <http://www.eso.org/gen-fac/pubs/astclim/paranal/clouds/> Right: monthly average precipitable water vapour from 1999-2005. Data taken from <http://www.eso.org/gen-fac/pubs/astclim/paranal/h2o/>

4.2.2 Prototype testing


The design described above required proof that the targets could be achieved. Once the design of the instrument had been decided, two test units were assembled. The first was a prototype instrument consisting of similar hardware but available for purchase immediately to the consortium, whereas the final instrument is built from customised hardware. The second test unit was a complete telescope system consisting of the mount, telescope, final optics and camera and was used for full unit tests. The first unit was installed on Roque de los Muchachos in La Palma, the second unit was installed at Geneva Observatory.

4.3 La Palma prototype

Between October 2009 and May 2010 a prototype unit was installed on Roque de los Muchachos in La Palma (Fig. 4.7 shows images of the prototype instrument) with the aim to prove that the requirements for NGTS could be achieved. The hardware used for the prototype was slightly different to that of the final instrument (Table 4.1 lists the hardware differences); all components were already available for purchase rather than requiring custom specifications, but was similar enough to test the main principles. Note two telescopes and cameras are shown in Fig. 4.7 but only one was available for testing. The prototype instrument was constructed with a separate



(a)



(b)

Figure 4.6: Computer generated renders of the NGTS facility. Top: the facility with roof open to see the instruments. Bottom: close up of the instruments. Courtesy of R. West.



Figure 4.7: Images of the NGTS prototype installed in La Palma. Left: view of the enclosure and telescope setup. Right: close up of the prototype.

Table 4.1: Comparison of detector features. WASP values are included for reference.

Feature	Prototype	Final Instrument	WASP
Camera	iKon-M 934 BR-DD	iKon-L 936 BR-DD	iKon DW436
Number of pixels	1k x 1k	2k x 2k	2k x 2k
Pixel scale	5.3"	4.9"	13.7"
Total FOV	2.27 sq. deg.	7.77 sq. deg.	64 sq. deg.
Readout time	~ 1.1 s	~ 1.5 s	~ 5 s
Filter	Unfiltered	550 - 900 nm	400 - 700 nm
Mount type	German equatorial	Fork	Fork

autoguiding telescope, but it was not possible to test the autoguiding system as the prototype suffered from mechanical flexure between the science and autoguiding telescopes. This and imperfect polar alignment caused the fields to drift during the course of a night.

The photometric reduction of the images was performed by a colleague João Bento and consisted of bias, dark and flat correction. A total of 909 flat frames were collected during the course of the prototype with an average of 14571 counts per pixel. 1608 bias frames and 267 dark frames were taken to characterise the system. Average bias frames were created each night to subtract the bias level added by the camera electronics. The camera did not read out an overscan region yielding a measurement of the bias level per image, so these bias frames were crucial to estimate the bias level. This bias level was subtracted from each calibration image in turn before continued correction. Average dark frames were created and normalised by the exposure time to remove the dark current from the flat and science images; the typical dark current value is $0.06e^- \text{ pix}^{-1}$ at the operating temperature of -79°C . To correct for vignetting and other optical features (e.g. dust), flat frames were taken each night during twilight to ensure even illumination of the CCD. These flat frames have an average of 14571 counts per pixel and are median combined




Figure 4.8: Graphic representing when data was taken by the NGTS prototype at La Palma, each month from October 2009. Red boxes represent the nights where data were collected, grey boxes mark days which do not exist in that month.

after each frame is normalised by its own median. Each science image is divided by the flat frame to reduce the effect of vignetting and inter-pixel dependencies. To estimate the photometric quality of the instrument, the calibrated frames from the prototype (reduced by João Bento) were analysed with an aperture photometry pipeline.

4.3.1 Aperture photometry implementation

To assess the quality of the NGTS data without biases from aperture photometry implementations, a standard aperture photometry package was used. To characterise the NGTS system the photometry package of choice needed to be able to extract the light originating from a point source, as the transit signals being searched for are contained in the star's flux. The flux needed to be corrected for the sky background flux not originating from the target star which contaminates the flux measurement. The sky background may not be constant across the whole image due to clouds so a sky background estimation close to each target star was required for the best measurement of the data quality. To create contiguous time series of the stellar flux, the positions in each frame needed to be extracted in an automated way, and a method to track the position of each object from frame to frame was required.

The aperture photometry pipeline was created using AUTOPHOTOM a freely available implementation used in the STARLINK² suite of astronomy tools. The photometry is configurable and allows choices of sky background estimation and flux calculation. The required input to AUTOPHOTOM is a list of target positions, which were estimated by extracting their positions from the initial frame by using

²<http://starlink.jach.hawaii.edu/starlink>

SEXTRACTOR³, a program designed to calculate the positions of any astronomical sources in an image. To track targets from frame to frame the centroiding available in AUTOPHOTOM was used. It was assumed that the stellar positions would not move significantly between images, so the position in the next frame would still allow centroiding to find the new position. The initial target catalogue created by SEXTRACTOR was used by AUTOPHOTOM as the input target list, and each subsequent frame took the calculated stellar positions from the previous frame as the input catalogue and stars were tracked. This tracking was even more important as the telescope drifted during the course of a night: typically the stars would drift by 20 pixels, but the displacement between each pair of frames was small.

Stars were included in the initial catalogue of stars if at least three pixels were 3σ above the background level. Included in the SEXTRACTOR output was a quality flag per object per frame which records problems that occurred with the source extraction such as the object was saturated or too close to the edges of the image. To prevent saturated stars or dubious detections entering the input catalogue this flag was used to filter the target list.

Figure 4.9 shows an example NGTS prototype image. The colour scale has been chosen to show only targets which were extracted and tracked. Figure 4.10 shows the region around an isolated star and the circles represent the aperture and annuli sizes. To maximise the signal to noise of the photometric measurement for each aperture, the radius was balanced between encompassing all of the stellar flux and minimising the sky background flux within the aperture. The aperture was set to 4 pixels to maximise the collected flux from the bright stars. The sky annulus size was set by scaling factors of the science aperture radius. To balance the size of sky aperture whilst maximising the number of pixels used for the sky background estimate, the inner annulus scaling factor was set to $S_{in} = 1.3$ and the outer scaling factor set to $S_{out} = 3$. To compensate for the sky annuli overlapping other stars or source flux from the target star being erroneously contributing flux to the sky background estimation, the sky background level was estimated using a 2σ clipped mean.

Standard magnitude system

The count rate of an object is specific to an instrument, so to compare the brightness of an object with other projects a standardised system of magnitudes is used. Magnitudes are quoted for the specific filter used during the photometric extraction. The sensitivity of the prototype was never calculated, but the QE response

³<http://www.astromatic.net/software/sextractor>




Figure 4.9: An example NGTS prototype image, taken on 2009-10-07. The colour scale has been configured to show from the limit of source detection at 3σ above the sky background to the saturation level of the CCD logarithmically. The colour scale shows the counts for each colour value.




Figure 4.10: Example aperture used by the photometry pipeline. The green circle represents the aperture size of 4 pixels used, the blue circles represent the sky annulus used to estimate the sky background. The colour scale is the same as used in Fig. 4.9.




Figure 4.11: QE quoted in the Andor brochure for the prototype camera, model number iKon-M 934 BR-DD.

quoted by Andor, the camera manufacturer is shown in Figure 4.11. This response is reduced further by the optics in the instrument, so is an optimistic estimate of the throughput of the system.

The zero point of an instrument represents the magnitude at which a count rate of one photon per second would be received. Given a collection of objects with measured fluxes and catalogue magnitudes, the zero point of the instrument

$$m_0 = m + 2.5 \log_{10} f \quad (4.1)$$

can be estimated, where m the catalogue magnitude and f the count rate of the object. The zero point gives a measure of the sensitivity of the instrument, where the fainter the zero point the more sensitive the instrument.

To calculate the zero point of the instrument, the UCAC3 catalogue was queried for stellar I band magnitudes in the frames taken. The I band was chosen as it is similar to the throughput frequency response of the prototype instrument. It is not a perfect match to the throughput of the prototype, the R band was an alternative but the red stars are a higher priority for the project so the ability to achieve a more accurate estimate in the red was chosen. Due to mismatches between the filter used in the reference catalogue and the instrument frequency response, each object gives a different zero point depending on its spectral type. The




Figure 4.12: Histogram of the estimated zero points for the field of the night of 2009-10-07. The dashed line represents the median zero point 20.60.

different estimated zero points create a distribution which represent the instrument, the width and shape of this distribution are related to how close the filter response of the catalogue magnitudes is to the instrumental response. The narrower the distribution the better the match. The stellar count rates as measured by the photometry pipeline were matched to the catalogue magnitudes, and the zero points estimated by substituting for f in Eq. 4.1. By matching the stars in a group of images the zero point was estimated, and this zero point was then used to convert all instrumental count rates into magnitudes. Figure 4.12 shows the distribution of the zero points calculated for the night of 2009-10-07, the median zero point was found to be 20.60. This zero point was used to convert the observed count rates to I band magnitudes.

Expected sources of noise

A transiting survey is subject to sources of noise which reduce the quality of the observations. As described in Section 1.10.4 the process of counting photons from a source is subject to Poisson noise, along with the background estimation. The thermal current from the CCD also contributes a level of uncertainty as does the




Figure 4.13: Geometry of the light path length through the atmosphere, from Birney et al. [2006].

output amplifier, and were corrected through collecting dark frames and flat frames (Section 1.10.2). The flat fielding process is subject to uncertainty, which introduces noise. To reduce this effect, we plan to keep stars centred on pixels during the course of an observation. The read noise was given in camera specifications from the manufacturer. Scintillation is caused by turbulence in the upper atmosphere and causes intensity fluctuations which add to the noise level. The scintillation is given in magnitude units as

$$m_{scin} = 0.004D^{-2/3}X^{7/4}e^{-h/H}(2t_{exp})^{-1/2} \quad (4.2)$$

where D is the telescope aperture in m, X is the airmass, h is the height above sea level, $H = 8000$ m is the scale height and t_{exp} is the exposure time [Dravins et al., 1998]. To convert into units of flux, the noise due to scintillation is

$$\sigma_{scin} = 1 - \left(10^{-m_{scin}/2.5}\right). \quad (4.3)$$

The largest systematic effect for a ground based survey is the effect of the amount of atmosphere the flux from a star has to travel through as the Earth rotates and the star changes position. As the target moves in the sky the light path length through the atmosphere changes and causes more or less attenuation as the airmass changes. The minimum path length and therefore brightest an object appears is straight overhead at zenith. All pointings other than this will receive less flux, and the object appears dimmer.

Figure 4.13 shows the simplified geometry of observing a target star from the




Figure 4.14: Extinction behaviour for a group of stars. The top panel shows the bulk photometry for this ensemble after the mean magnitude of each star is removed. The bottom panel shows the same magnitude offset from the mean against the measured airmass during the course of the night. The example 1σ uncertainty is shown in the lower left corner. Note: different stars were used for the ensemble measurement between the rising and settings fields, see Section 4.3.3.

ground through the atmosphere. For the simple approximation we assume that the atmosphere is planar and has a constant density. The path length of light is calculated as $X = \sec z$, where X is the airmass, and z as Figure 4.13 the angle between target and zenith, also known as the zenith distance. The difference in magnitude of the object is linearly dependent on the airmass [Barbieri, 2007]. Figure 4.14 shows the airmass behaviour for a group of stars observed with the NGTS prototype.

The correction for the extinction is

$$m_{\lambda,0} = m_{\lambda} - k'_{\lambda}X - k''_C CX. \quad (4.4)$$

where $m_{0\lambda}$ is the magnitude of the target if it were observed from above the atmosphere, m_{λ} the observed magnitude, k'_{λ} an empirical extinction coefficient scaling the correction required to the airmass, and k''_C is second order colour correction applied to the first order term to correct for the different stellar colours [Birney et al., 2006]. Each object will show different extinction behaviour as the airmass changes as the magnitude difference depends on the colour of the object C . The coefficient




Figure 4.15: Extinction measurements from the Carlsberg Meridian Telescope during the time the prototype was installed. Measurements are of the k_{λ} parameter using the SDSS r' filter.

k_{λ} is tabulated by observations from the Carlsberg Meridian Telescope at La Palma⁴ (Fig. 4.15 shows this measurement during the time the prototype was installed.)

4.3.2 Removing trends

Systematic effects are removed through the detrending of the data. This is accomplished with the application of the SYSREM algorithm.

System

When observing a single target, a comparison star can be used to remove instrumental effects that decrease the photometric precision, as any systematics will affect both stars together and will be removed by dividing the target flux measurement by the comparison flux measurement. Provided the comparison star is a similar spectral type the extinction removal will be good, but if the colours are different the extinction will affect both stars differently and be corrected. SYSREM is an algorithm designed for detrending a large collection of lightcurves which share similar systematics (see Section 2.2.2 for implementation details.) Systematic effects which are shared between stars are removed without prior knowledge to the source, and is superior to using comparison stars as it harnesses the power of using all of the

⁴http://www.ast.cam.ac.uk/ia/research/cmt/camc_extinction.html




Figure 4.16: SYSREM coefficients for 2009-11-19. Coefficients are separated by 5, and are ordered in time from 0 to 3 in the order they were corrected for.

stars to remove common systematics. By default, the implementation of SYSREM used runs the algorithm four times, which was found to be optimal for WASP. We use this number of iterations as it has been tested on the WASP dataset, and so has been calibrated on a much larger dataset. It was tested running performing more iterations with SYSREM and we found no improvement after running SYSREM more than four times. Figure 4.16 shows the four per-image coefficient values for each image for a single dataset, indicating the systematic signals which were present in the data but have been detected and removed by SYSREM. Each coefficient decreases in amplitude suggesting that the corrections required become less significant with each iteration. Each star is still allowed a weight for these functions but overall they decrease.

We note that errors in the flat field coupled with the stars drifting may have induced systematic trends for which SYSREM is not suited. The behaviour of each star would be independent of the others and so would not be modelled well by constructing weighting functions from the bulk of the lightcurves.

4.3.3 Limitations

As the pipeline relies on matching objects from frame to frame by assuming that the CCD positions do not change much, and that the centroiding is sufficient to track objects, any significant change in CCD position causes errors in matching flux measurements from frame to frame. This limits the lightcurves that were produced when reducing the prototype data as the mount was a German equatorial mount

which required the tracking to stop when the mount is pointing at the meridian to prevent the mount from tracking or slewing into itself. Near the meridian the telescope is turned over, and the guiding had to be reset to reacquire the field. Early in the testing the two fields were treated as separate and so would not include the same stars. A few full nights were devoted to three different fields so long lightcurves could be generated, but the matching of fields before and after meridian flip was not sufficient to match up the objects with this naive centroiding method.

This aperture mismatch also manifests itself within single fields which are within a half night. An aperture around a faint star which is near a brighter star will have a chance that the aperture will switch to the brighter star. The lightcurve from the aperture will have a change in mean flux, potentially very early in the night. An example lightcurve with an aperture jump is shown in Figure 4.17. The spread in the flux distribution is much larger if the fainter flux values are included, and the mean is at a different level. Note also that even though the data were detrended with SYSREM the lightcurve shows a clear decreasing trend with time. The aperture jump has decreased the efficiency of the detrending stage of the pipeline further decreasing the measured precision.

Since these occurrences are infrequent in the field, they were filtered out for assessing the overall quality of a dataset. The aperture jumps will be visible in the positions of the aperture centroid, sometimes in both x and y directions, but always in one. Lightcurves were rejected from the dataset before SYSREM was run if they had a large positional difference from point to point. The rejection threshold was configured during the analysis stage.

4.4 Noise model

To simulate the precision of the NGTS prototype, we constructed a noise model. This enabled us to predict for a given exposure time, assumed sky background level and mean airmass the expected level of precision we would reach. To calculate the source noise from a given stellar I magnitude, the flux rate of the instrument is required. To compare the noise levels with the NGTS target of 1 mmag the fractional noise is required, so each noise source once calculated is scaled by the source flux. The sky background in the I band was measured for La Palma by the TNG⁵, measuring typical sky backgrounds in magnitudes per square arcsecond shown in Table 4.2. The third column shows the background sky flux after conversion from magnitudes per arcsecond, to flux in electrons per second per pixel through

⁵http://www.tng.iac.es/info/la_palma_sky.html




Figure 4.17: Example lightcurve from 20010-02-04. Flux has been detrended with SYSREM. A clear aperture jump can be seen at $t \sim 2455232.36$. Lines indicate the weighted mean and standard deviation values for the lightcurve if all data is included (dotted line), and after removing the flux at the lower level (solid lines).

the application of the following relation:

$$m_{sky} = m_0 - 2.5 \log_{10} \left(\frac{f_{sky}}{p^2} \right). \quad (4.5)$$

where p is the pixel scale of the instrument (5.3 arcseconds per pixel) and m_0 the zero point as calculated in Section 4.3.1. Sky background levels for the prototype are shown in Fig. 4.18 against Moon phase, and are consistent with the predictions given in Table 4.2.

Each readout of the CCD causes spurious electrons to be added from the output stages. This noise is characterised by estimating the noise level in the bias frames. The noise level is also available from the camera specifications from the manufacturer and was 7.83 e^- . Equations 4.2 and 4.3 describe the fractional noise from scintillation in the atmosphere. For NGTS the telescope aperture D is 0.2m, the height h is 2400m above sea level. The exposure time and airmass are free parameters to be tuned to the observing conditions. The dark current was taken as $0.06 \text{ e}^- \text{ s}^{-1}$ at the operating temperature of -79°C . This value was scaled by the exposure time to calculate the level of dark current. The total noise level was




Figure 4.18: Sky background level for different Moon phases (0 is new Moon, 1 is full Moon). The plot has been truncated at $400e^{-}\text{pix}^{-1}s^{-1}$ due to cloudy conditions causing atypical sky background levels.

Table 4.2: Sky background values for days since New Moon, for La Palma. Sky magnitudes are quoted for the I band in magnitudes per arcsecond, and f_{sky} is the flux in electrons per second per pixel.

Days	I magnitude	Flux ($e^{-}s^{-1}\text{pix}^{-1}$)
0	19.70	50.39
3	19.60	55.25
7	19.10	87.57
10	18.20	200.60
14	17.30	459.55

calculated through the modified CCD equation (see Eq. 1.32 for the original)

$$\frac{S}{N} = \frac{N_*}{\sqrt{N_* + n_{pix} (N_S + N_D + N_R^2) + \sigma_{scin}^2}}. \quad (4.6)$$

where σ_{scin} is an extra term accounting for the scintillation noise, shown in Eq. 4.3. Figure 4.19 shows a noise model for the NGTS prototype with an assumed exposure time of 30 s, dark sky and an airmass of 1.3. The comparison with the data is shown in Fig. 4.21. We see that below $I=12$ the noise is sky background limited, and above $I=10$ the noise is scintillation limited. The read noise and dark current noise never become the dominant noise source throughout the range of magnitudes.

4.5 Prototype results

33 half nights worth of data were analysed with the aperture photometry pipeline to ascertain the quality of the instrument. Figure 4.20 shows the collected quality statistics for the NGTS prototype dataset. Data used to describe the results from the prototype are taken from the data sections highlighted in grey.

4.5.1 Precision

To detect Earth like planets, we require a fractional noise level of < 1 mmag on the timescale of a transit (Section 4.1), requiring the fractional noise level given in Eq. 4.7 to be $\sigma_{frms} \leq 0.1\%$.

$$\sigma_{frms} = \frac{\sigma_f}{f} = \frac{1}{\frac{S}{N}}. \quad (4.7)$$

Precision is measured using Eq. 4.7 and can be calculated from a lightcurve by noting that the rms of the deviations is identical to the standard deviation for a normally distributed noise source, so

$$\text{rms} = \sigma = \frac{1}{N} \sqrt{\sum_i \frac{f_i - \bar{f}}{\sigma_i}} \quad (4.8)$$

for a lightcurve with N points of flux f_i and uncertainty per measurement σ_i , and mean flux \bar{f} . To compare stars with differing mean fluxes the *fractional* rms is calculated from the ratio of the mean flux to the rms value, leading to σ_{frms} from Eq. 4.7.

NGTS is a wide field instrument allowing for simultaneous flux extraction




Figure 4.19: Noise model for the NGTS prototype, simulated for an exposure time of 30 seconds and an airmass of 1.3. The lines are coloured as shown in the legend, one for each source of noise.

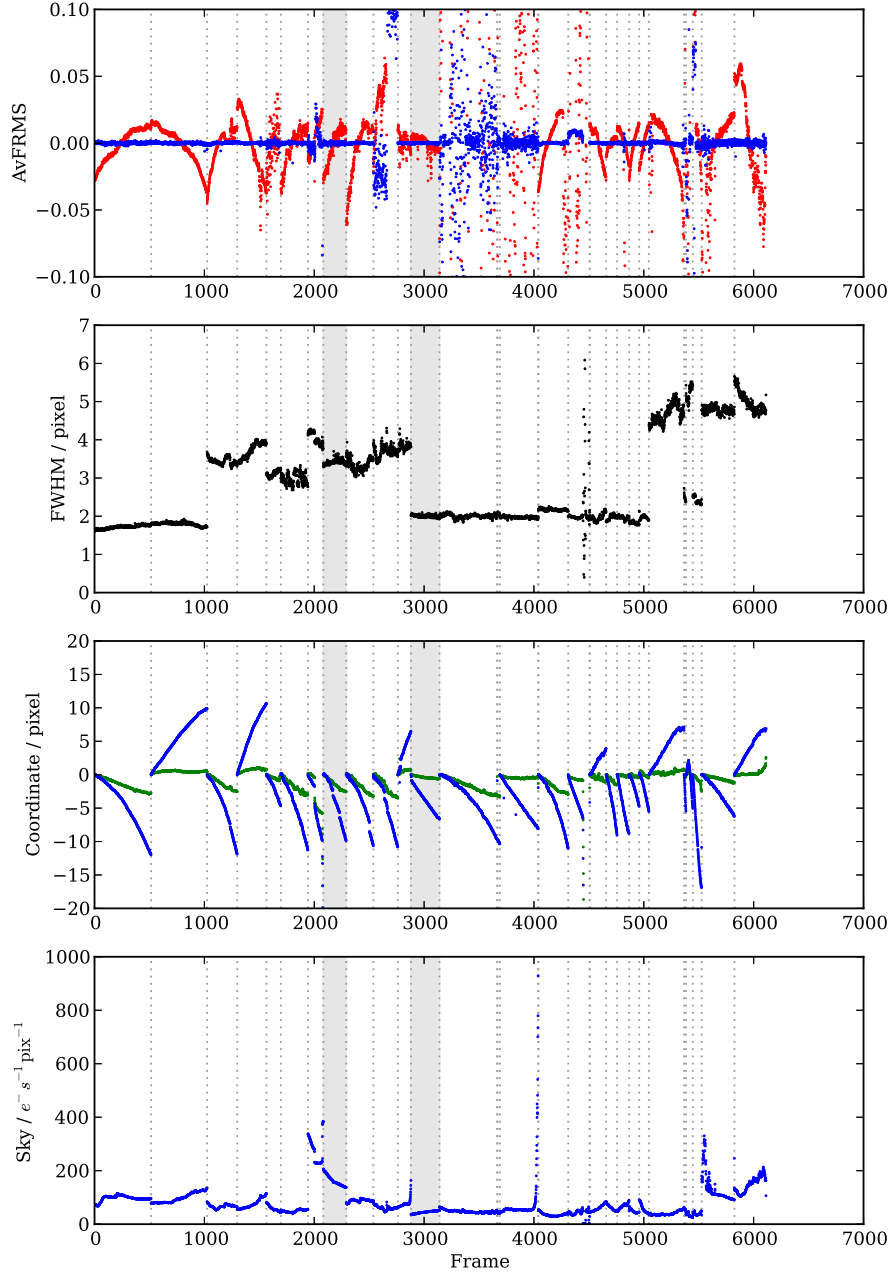


Figure 4.20: Properties of the NGTS prototype dataset. The top panel shows the median fractional residuals. Red points are calculated values before SYSREM has been used, blue points are calculated after SYSREM. The second panel shows the measured full width half maximum of the images as measured by SEXTRACTOR. The third panel shows the relative mean drift of the apertures from their starting positions, where the green points are the x coordinate and blue points the y coordinate. The bottom panel shows the sky background as measured by SEXTRACTOR in electrons per second per pixel. Each panel is plotted in chronological order against frame number, and vertical dotted lines delimit the datasets. The two grey regions represent chosen datasets to discuss the photometric quality.




Figure 4.21: Fractional rms for the field collected on 2009-11-19, pre-meridian. Each point represents the statistics for a single lightcurve. The lightcurves have been detrended using SYSREM. Overlaid are predictions from the noise model (Section 4.4) where the colours represent scintillation (cyan), read noise (green), sky noise (blue), source noise (red), dark current noise (black) and the total noise (magenta).

from multiple stars. This collection of stars is referred to as a field and defines a set of data which was collected simultaneously. This field will contain some instrumental systematics which may be common to every object, for example first order atmospheric extinction. The fractional rms for a field can be calculated to ascertain the limiting precision for any star and therefore the overall quality for that dataset. Figure 4.21 shows an example fractional rms calculation for the pre-meridian field of 2009-11-09. The data are well matched to the noise model.

The fractional residuals measured represent the brightness-independent photometric quality for an object, and the median of this value per frame assesses the overall quality of the entire dataset during the course of the night. Any trends common to the bulk of the objects will appear, and as the median is robust to outliers, the median fractional residual value represents the bulk photometric quality well, though is weighted towards fainter stars as they are more numerous. Calculating this

for the prototype dataset (shown in Fig. 4.20) clearly shows the ability of SYSREM to remove systematic trends, and immediately areas of well behaved photometric conditions (for example frames 1 – 1000) are visible from areas of bad photometric conditions (for example frames 3100 – 4000).

4.5.2 Noise colour

Noise is defined as white when there is equal power at all frequencies, and described as red when more power is present at low frequencies, and the uncertainty of a single measurement is dependent on neighbouring measurements. In contrast blue noise has more power at high frequencies. The noise level of purely white noise dependent measurements reduces as \sqrt{N} for N combined measurements. Independent noise sources, or measurements with white noise, are combined by summing the variances σ_i^2 :

$$\sigma_T^2 = \sum_i^N \sigma_i^2. \quad (4.9)$$

This is not the case when red noise is present. If we define the measurement X_n to be dependent on the previous measurement X_{n-1} plus a random component y_n such that

$$X_n = pX_{n-1} + y_n, \quad (4.10)$$

where p is the level of correlation ($p = 0$ for white noise). By squaring and averaging, the variance on point n is given as

$$\sigma_n^2 = p^2 \sigma_{n-1}^2 + \sigma^2 \quad (4.11)$$

where σ_n is the variance at point n , σ_{n-1} is the variance at point $n - 1$ and σ is the noise from component y_n [Gilman et al., 1963]. When summing noise sources as Eq. 4.9 and given that $|p| > 0$ for red noise, the combined noise level is $> \sqrt{N}$.

Pont et al. [2006] propose a method for assessing the noise colour of a set of lightcurves. By comparing the fractional rms value measured from a binned lightcurve to the predicted fractional rms value for white noise (σ/\sqrt{N}) we can assess whether the measured binned noise level achieves this value. If the values match then the lightcurve contains purely white noise, whereas if there is a discrepancy then we must conclude that red noise is present. By tuning the number of points per bin we choose the time-scale over which we wish to test the correlated noise significance.

Figures 4.22a and 4.22b show the results of binning up the lightcurves com-




Figure 4.22: a) Measured fractional noise values for the night of 2009-11-19. Black points show the unbinned values, red points are the measured fractional rms values after binning up by 60 points per bin corresponding to a total integration time of 30 minutes, blue points are the black points decreased by their theoretical white noise value of $\sqrt{N} = 60$. b) The same analysis procedure as Fig. 4.22a but for the night of 2009-10-07, the first data section in Fig. 4.20. Similarly 60 points were used.

pared to the theoretical white noise behaviour. Figure 4.22a indicates that the red noise component is low as the red points are similar to the blue values across the entire magnitude range. A small red noise component is apparent for objects brighter than $I \sim 12$. Note that there is more spread for the red values compared to the blue values, as the number of points with which a fractional rms is calculated decreases the uncertainty of the measurement increases, and are plotted on a logarithmic scale emphasising small differences. Note that the limiting precision measured for binned data reaches the 1 mmag level. Figure 4.22b shows a much higher red noise component for objects brighter than $I \sim 12.5$ as the deviation between measured and theoretical white noise behaviour is large.

As the bright stars are most affected by red noise, we used the behaviour of the binned points to measure the red noise component. We were interested in how the limiting precision of the brightest objects behaved as the flux measurements were combined. The faint stars are dominated by white noise from the sky background and so should bin up as white noise, but the brighter objects do not follow simple noise arguments. The limiting unbinned precision for the bright stars could be due to scintillation as predicted by the noise model, but it could also be due to any systematic noise or red noise in the lightcurves. Red noise would limit the precision of the project as it is much harder to detrend and remove if the source is unknown. By binning up the lightcurves the fractional noise should decrease by \sqrt{N} , and the




Figure 4.23: Binned median precision of bright stars on 2009-11-19 (dark) and 2009-11-06 (light). Dashed lines show the expected white noise behaviour.

precision will increase by the same factor. This was tested by calculating the median fractional rms of the brightest unsaturated lightcurves, generally in the magnitude range of 10 – 11 (see Fig. 4.21 for reference). For binning value N each lightcurve in the magnitude range was binned up so every N points became one measurement, and the fractional rms was recalculated. The median of this value for each lightcurve was then plotted against N to assess the noise colour. Figure 4.23 shows two such measurements for the highlighted regions in Fig. 4.20, taken during dark time and bright time. Both show the binned precision reaches sub-mmag levels and near-white noise behaviour.

Some of the nights do not exhibit white noise behaviour. For example frames 500-1000 in Fig. 4.20 show red noise behaviour. Example lightcurve data is shown in Fig. 4.24, and we find a clear 117 minute period, which corresponds to 0.08 days, which is consistent with the time taken for the star to drift across pixel boundaries. The autoguiding system for the prototype was not able to be tested, but will be implemented for the final instrument to reduce the pixel drifting systematics.

4.5.3 PSF sensitivity

By measuring the stability of the focus of the instrument, both the stability of the mount and focuser are determined. This gives important information in turn about




Figure 4.24: Example data from a single lightcurve collected during 2009-10-07. The panels show: the measured light curve (top left), the x coordinate for the duration of the measurement (middle left), the y coordinate for the duration of the measurement (bottom left), the total distance traveled by the aperture in pixels (top right), the sky background (middle right) and a Lomb-Scargle periodogram of the lightcurve (bottom right).




Figure 4.25: Examples of the focus levels shown in Fig. 4.20. Left: an in focus image from 2009-10-07 (frames 1-500). Right: an out of focus image from 2010-03-03 (frames 5600-5800). Both images are logarithmically scaled and show pixels 3σ above the background level. The images show the same number of pixels.

the stability of the hardware whilst being operated robotically. Examples of the focus levels are shown in Fig. 4.25 and show the range of focus levels utilised by the prototype.

The focus stability was measured with SEXTRACTOR by extracting the objects in every frame and calculating the median FWHM for each frame. SEXTRACTOR was used to extract the sources in every frame taken by the NGTS prototype, and the measured FWHM recorded. The median value in each frame provides a robust measurement of the *bulk* focus level in a frame as it is insensitive to statistical outliers such as blended objects. The time history of the median (Fig. 4.20, second panel) shows multiple levels of focus used during the prototype history. Extreme outliers at frame ~ 4500 are because the dome was closed during these exposures, which also causes the extreme drift in the pixel coordinates, the sky background level decreasing to zero and unusual behaviour in the average fractional frms plot at the end of the data region. Observations can be categorised into roughly three categories: in-focus (e.g. frames 1 – 1000), mildly defocussed (e.g. frames 1000 – 3000), and out of focus (e.g. frames 10000 – 11000). The levels of focus vary between relative stability (e.g. frames 3000-4000 despite bad weather conditions as evident by the top panel) and large systematic variations (e.g. frames 1000-1500). We note that focus changing affects all stars equally, as the quality of the post-sysrem data (top panel in Fig. 4.20) is consistent despite changes in focus levels. The variations in focus are most apparent where the PSF is larger, and the right panel in Fig. 4.25 shows irregular shapes during the levels of large PSF. This coupled with the pixel drifting causes uncertainties in the measurement of the FWHM, but this affects stars similarly so does not affect the bulk quality after sysrem.

4.5.4 Blending

A wide field survey such as NGTS depends on observing a large number of stars. The field of view of the instrument is set by the focal length of the telescope, and the detector. To maximise the number of stars visible with a field of view set by the hardware, we need to observe denser fields, with more stars per square arcsecond of sky. The NGTS prototype had a FOV of 2.27 square degrees, and we need to know how many stars in a field would be unblended. The PSF of the instrument causes nearby stars to merge together, and they can no longer be discerned as separate objects so both are no longer valid exoplanet target hosts. A planet transiting across a diluted star has a smaller observed transit depth, and causes the object to be more difficult to detect.

Pixel response function (PRF)

When an image is taken of a field, the PSF is sampled discretely by the pixels on the CCD. Additionally the photometry is performed by sampling the pixelated image with a circular aperture. The pixel response function is a measure of how much flux from a point source is measured from a given distance from the object during the aperture photometry. This measurement was used to predict in a field how much flux for each target is contributed by all *other* targets in the field.

To perform this analysis, an isolated star was found in an NGTS image. As the design of the final instrument involved a small psf, on the order one pixel, the data with sharpest focus was used. The aperture radius was set to $r = 1.5\sigma_{PSF}$ where σ_{PSF} the width of the PSF. A single aperture was placed over the star to measure the flux as would be measured during the normal photometric process. The aperture was moved in each orthogonal direction aligned with the pixel grid, incremented by 0.1 pixels each time and the flux at each position calculated. The orthogonal directions were used to ensure that if the PSF were asymmetric the shape would be accounted for but in practice (see Fig. 4.26) the measured response function is symmetrical. The PRF in each direction were averaged to decrease the effect of measurement error, giving a simple empirical measurement of the flux from a source at a distance r in pixels.

$$\text{PRF}(r) = \frac{1}{4} \sum_i^4 \text{PRF}_i(r) \quad (4.12)$$

where i represents the orthogonal measurement, PRF_i represents the PRF for that direction.




Figure 4.26: Overlaid response functions for the four orthogonal directions. The combined profile is plotted in black offset by 0.2.

Dilution measurement

The UCAC3 catalogue [Zacharias et al., 2010] was queried for three locations shown in Fig. 4.27, representing three typical NGTS fields, located at differing galactic latitudes allowing for an assessment of the dilution as a function of galactic latitude and star density (Table 4.3). Each field lies close to zenith and in regions of high possible yearly coverage and so represent prime target fields for NGTS. Each object in the returned catalogue has an equatorial position (α, δ) and magnitude in the I band. A double iteration over each object i with every other object j , $i \neq j$ was performed, calculating the distances between them by converting equatorial coordinates to pixel coordinates based on astrometric solutions to NGTS images. The flux of object j was converted from its I band magnitude and the prototype’s zero point as calculated in Section 4.3.1. The flux in the aperture of object i from object j is therefore the flux of object j scaled by the PRF at the distance r . For each object, the dilution is the total amount of flux in the aperture not from the target object itself, as a fraction of the total flux in the aperture:

$$D_i = \frac{1}{f_i} \sum_{j, i \neq j} \text{PRF}(r_j) f_j \quad (4.13)$$




Figure 4.27: NGTS yearly coverage. The colour scheme represents the yearly number of hours that area of sky is visible for, between 2013-01-01 and 2014-01-01. The horizontal blue line is the latitude of the NGTS site at Paranal. The blue area shows the galactic plane, with parallel lines indicating a proposed avoidance angle, which was determined by this study, and the star represents the position of the galactic centre. The sinusoidal green lines indicate the path of the Moon throughout the year. The three black boxes are fields used for crowding analysis. Image courtesy of João Bento.

where f is the flux from an object. We set a threshold of 10% to be the maximum dilution due to the high precision requirement of the project, which is a lower threshold than the WASP project where the limiting dilution value is 25%. Using the catalogue of stars, a histogram of the un-blended objects' magnitudes was created to assess where dilution would reduce the number of possible targets. We can see from Fig. 4.28 that objects brighter than $I \sim 11.5$ are completely un-blended, due to the brightness but also the rarity of these bright objects. As NGTS will observe objects down to $I \sim 17$, we will lose $\sim 30\%$ of the faintest objects, with a slight dependence on galactic latitude, but the majority of all objects are un-blended. The main goal of NGTS is to detect small planets around bright stars, so the loss of faint stars does not adversely affect the primary science goal of the project.

Table 4.3: Coordinates of the fields used for the crowding analysis. All fields are at $\delta = -45$ degrees.

Field	α (hours)	l (degrees)	b (degrees)
1	4	251.28	-48.63
2	12	293.46	16.92
3	20	354.60	-30.52




Figure 4.28: Histogram of the fraction of objects which are un-blended. The three lines indicate different galactic latitudes $|b|$ as indicated in the legend.

Observed transits

During the night of 2010-02-04 two known planets transited their host stars: GJ 436 b and WASP-11 b. The observed transit of the warm Neptune GJ 436 b is shown in the upper panel of Fig. 4.29. This planet exhibits a transit depth of 1% and is clearly significant. GJ 436 b orbits an M5 star and represents a good target for NGTS. The small stellar radius causes an increased transit depth compared to earlier type stars. WASP-11 b was also observed, and is shown in the lower panel of Fig. 4.29. The predicted signal has a depth of $1.62^{+0.03}_{-0.02}\%$ [Wright et al., 2011], and is clearly visible in the NGTS data at a significance of 3.5σ . It is a good test object as the parent star is a K type star, a prime target for NGTS, and its radius is $0.910^{+0.06}_{-0.03}R_J$ so slightly smaller than Jupiter. Both transits are clear from the unbinned data, and though the fields were chosen to contain these planets the lightcurves were extracted through the wide field photometry pipelines (in the case of GJ 436 b the photometry was not performed using the pipeline described in Section 4.3.1), giving confidence in the performance of the instrument.

4.6 Summary

The NGTS project aims to detect small planets around small stars. In this chapter I have described the application of aperture photometry techniques to the data collected using the prototype instrument, and characterised the bulk photometric quality. I have demonstrated the ability of the prototype instrument to reach sub-mmag precision around bright stars. This coupled with the increased sensitivity to smaller stars enables shallower transits to be detected.

I have calculated that only stars fainter than $I = 12$ will suffer from blending with neighbouring stars at the 10% level (see Fig. 4.28), even at galactic latitudes of 17° , which allows NGTS to observe closer to the galactic plane than e.g. WASP and observe more stars.

The clear transit signal visible in the individual transit plots (Fig. 4.29) indicates that individual transits of Jupiters and Neptunes around M-dwarfs are apparent from the unbinned data. Fewer transits will therefore be required to infer the presence of a planet, which has two advantages: the sensitivity to longer periods is increased, and the follow up instruments can be alerted faster.

A noise model for NGTS has been constructed, which predicts that stars brighter than $I = 10$ are scintillation dominated, whereas stars fainter than $I = 12$ are sky background dominated. To reduce scintillation the exposure time can be increased, and to reduce sky noise the instrument PSF can be decreased. Both noise




Figure 4.29: Top: observed transit of GJ 436 b with fitted transit (courtesy of J. McCormac). Bottom: observed transit of WASP-11 b. Red points are the measured flux values with their uncertainties in blue from the photometry pipeline. Overlaid in black is an unfitted transit model with values from the literature. Both datasets display unbinned data.

sources however are white noise sources, so the precision will increase as \sqrt{N} with binning. The main source of systematic noise for the prototype which could not be removed with SYSREM was found to be caused by the drifting of stars across pixels. The autoguiding system was unable to be tested for the prototype, but will be present for the final instrument, and we have shown that the instrument can achieve sub-mmag precision even under non-ideal circumstances.

In the next chapter, I modify the noise model developed using characteristics from the prototype instrument to include co-adding measurements to predict the precision reachable. This new model is then applied to optimise the desired exposure time for the project, and predict the number of planets NGTS will detect.

Chapter 5

NGTS final instrument and planet catch simulations

From early August 2012 a single NGTS unit was assembled and installed on the roof of the Observatory of Geneva for full system tests with a complete unit. It was the first chance to test the final system at almost its full performance. The only aspect missing from this test system was being located at Paranal, but otherwise it was realistic chance to test the design.

In this chapter I recap the goals of NGTS along with the design chosen to achieve these goals in Section 5.1. Section 5.2 describes how the noise analysis applied previously to the prototype was used to study the ability of the NGTS final instrument to achieve the goals set. Section 5.3 describes the analysis of a single NGTS camera. We used the noise models of the NGTS instrument to predict the number of detectable planets for NGTS, discussed in Section 5.4.

5.1 The NGTS project

The primary goal of NGTS was to find small planets around bright stars. These stars will be bright enough for radial velocity confirmation, therefore allowing measurements of the planetary mass and yielding the bulk densities. We aim to detect a large enough sample to determine the bulk properties of the population of these planets, which has not been done before.

The atmospheres of small planets are not well studied, as extremely precise flux measurements are required which is challenging with the current instruments. Either very precise measurements in the differences in transit depths for different wavelengths are required, or the relative flux levels before and during the very

Table 5.1: Comparison of detector features. WASP values are included for reference.

Feature	Prototype	Final Instrument	WASP
Camera	iKon-M 934 BR-DD	iKon-L 936 BR-DD	iKon DW436
Number of pixels	1k x 1k	2k x 2k	2k x 2k
Pixel scale	5.3"	4.9"	13.7"
Total FOV	2.27 sq. deg.	7.77 sq. deg.	64 sq. deg.
Readout time	~ 1.1 s	~ 1.5 s	~ 5 s
Filter	Unfiltered	550 - 900 nm	400 - 700 nm
Mount type	German equatorial	Fork	Fork

shallow secondary eclipse. With bright stars the level of photon noise decreases and so these measurements are made more possible. With NGTS we aim to produce targets for current and future instruments such as the VLT, E-ELT and JWST, as well as targets for dedicated exoplanet characterisation missions e.g. EChO [Tinetti et al., 2012] or CHEOPS [Broeg et al., 2013]. To achieve the precision required to detect the shallow transits caused by Neptunes and super Earths, we aim to achieve sub-mmag photometric precision from the ground, and the survey is optimised to observe later type and therefore smaller stars.

In the previous chapter the design specification of the final instrument was described. The key components are listed in Table 4.1 and repeated here in Table 5.1. The design of NGTS uses back illuminated, deep depleted CCDs with a high QE in the red to increase the flux collected from later type stars. Large CCDs coupled with fast optics allow a wide area of sky to be covered.

With the prototype instrument we considered the sources of noise present in the instrument and constructed a noise model. We showed that this noise model was an accurate representation of the sources of noise through comparisons with the total noise observed by the instrument. We have demonstrated that sub-mmag photometry from the ground was possible with a wide field instrument.

5.2 Reapplication of the noise analysis

In Section 5.2.1 the method for determining the noise level of the instrument and its components was described. To understand the noise components of the final instrument the noise model was updated to incorporate the change in hardware. The noise model for a 30s exposure is shown in Fig. 5.1a, and the noise model for the prototype is shown in Fig. 5.1b for comparison. We find the limiting precision for bright objects similar as it is caused by scintillation. The precision of the fainter

objects $I > 12$ increases as with a smaller PSF the level of sky noise decreases. To estimate the noise on the timescale of a transit, the noise model was updated to allow the combining of exposures to calculate the total noise present after a total elapsed time with many individual exposures.

5.2.1 Updated noise model

NGTS relies on its high photometric accuracy to achieve the goals set by the consortium. An accurate knowledge of the sources of noise in a realistic situation are crucial for this task. We need to know the noise on the time-scale of a transit to attempt to assess the significance the transit signal will have, and therefore how easily it will be detected.

Each image when exposed must be read out by the electronics which takes finite time. For a CCD with n_x columns and n_y rows this process takes t_R seconds

$$t_R = n_y \left[t_v + \frac{n_x}{\nu_h} \right] \quad (5.1)$$


where t_v is the time taken to shift a row down in parallel, and ν_h is the clock speed for the serial register. For the NGTS CCDs the values used were $t_v = 39\mu\text{s}$ and $\nu_h = 3\text{MHz}$, though other settings were available. By convention a shift in parallel is defined as the vertical direction and a shift in the serial register is the horizontal direction. Each exposure then takes t_T seconds to expose, including photon-collecting time t_{exp} :

$$\begin{aligned} t_T &= t_{exp} + t_R \\ &= t_{exp} + n_y \left[t_v + \frac{n_x}{\nu_h} \right]. \end{aligned} \quad (5.2)$$


With 2048 pixels in each dimension, the additional readout time added by the CCD electronics was 1.49s. This additional time during which the telescope is not exposing has the effect of increasing the noise levels for a given science exposure time, becoming important at short exposure times as a smaller fraction of photons are collected. The number of exposures in time interval T is therefore

$$n_{exp} = \frac{T}{t_T}. \quad (5.3)$$

We assume white noise behaviour so the levels of fractional noise decrease as $\sqrt{n_{exp}}$. The PSF for the full instrument was smaller than that of the prototype instrument so the assumed aperture size was reduced to 1.5 pixels, to match 1.5 times the PSF



(a)



(b)

Figure 5.1: Noise models for the prototype (top) and final instrument (bottom). Noise models were calculated for an exposure time of 30 s.

of 1 pixel. We observe that the level of noise for objects $I < 12$ decreases due to less noise from the sky measurement, but that scintillation is still the dominant noise source for the bright stars.

We update the noise model of Chapter 4 to incorporate the total integration time, and two examples are shown in Fig. 5.2 for two exposure times: 5 s and 30 s, integrated to one hour. Shown are the point at which the total noise model crosses the 1mmag point, and saturation limits (see Section 5.4.1) to predict the number of non-saturated objects for which the observed precision is better than 1mmag. The noise model is now dependent on two parameters of interest: the exposure time used for each observation, and the source magnitude. Figure 5.3 shows the noise models for a range of science exposure times for a 9th and 15th magnitude object for airmass 1 and 2. With longer exposures, the number of exposures decreases and so the read noise (green) decreases dramatically, but the other noise sources remain a constant fraction, apart from at exposure times $\lesssim 10$ seconds where the readout time of the CCD becomes a significant fraction of the total integration time and so fewer photons are collected.

5.2.2 Photometric quality

Most of the time the NGTS unit was installed at Geneva Observatory in 2012 was allocated to testing the hardware and guiding mechanisms. Towards the end of the installation before it was dismantled for further work and laboratory testing, one night was given to photometric testing, to assess precision and stability. One section of this testing was to take long runs on a single field to get some long lightcurves to assess the noise levels in the instrument.

We chose the Kepler field at 19:03:35 +49:11:33 for this task due to the high amount of photometric coverage that had already been undertaken on this field from the Kepler team in the construction of the Kepler Input Catalogue (KIC) [Brown et al., 2011b], and other groups observing the Kepler field, such as the Kepler INT Survey (KIS) which contributes complimentary U, g, r and i band photometry of 97% of the Kepler field [Greiss et al., 2012]. The field was chosen from within the Kepler field of view at the furthest edge away from the galactic plane. We chose this field as it was representative of a typical NGTS field.

During one hour and five minutes, 600 exposures of 4 seconds were taken to provide long lightcurves, along with various calibration frames measuring dark current and flat field effects for image reduction. Images have an overscan region where the readout electronics continue to read out pixels which do not exist, and allowed to measure the bias level added to images to prevent negative pixel values

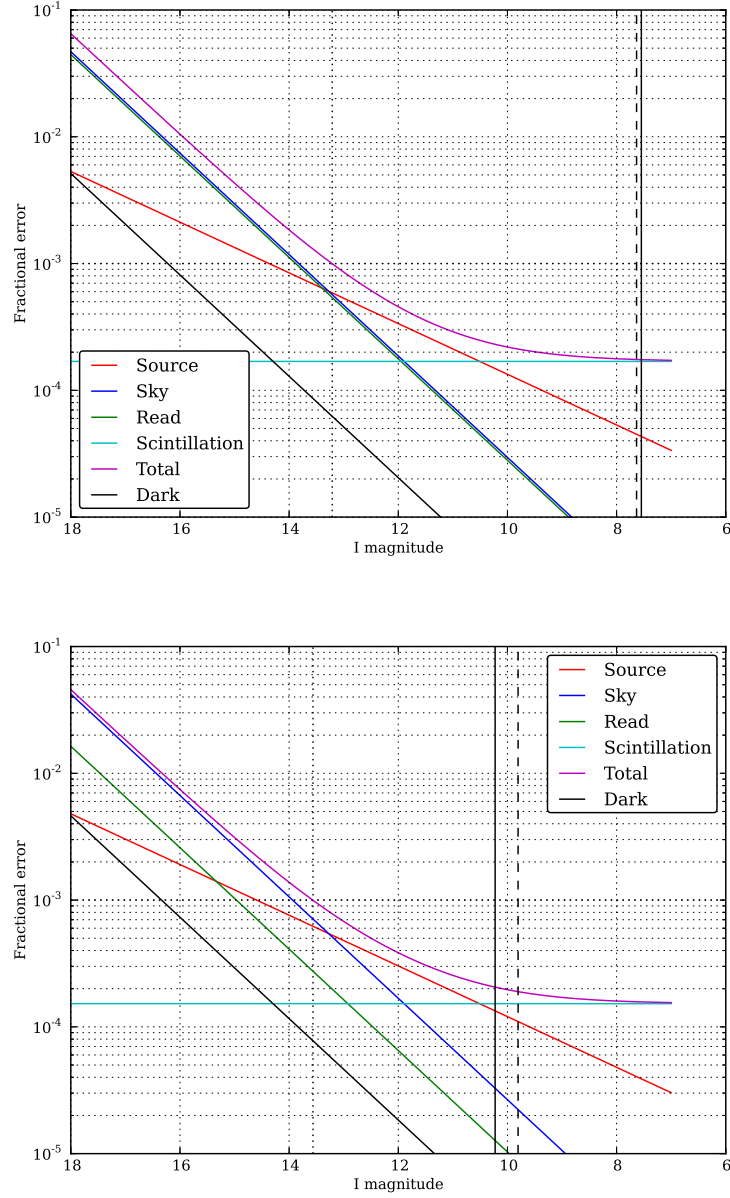


Figure 5.2: Noise models for NGTS. The dotted vertical line marks where the noise model crosses the 1 mmag point, solid and dashed vertical lines mark the saturation points for dark and bright times respectively, horizontal dotted line marks the 1 mmag point. Both panels have been integrated to a duration of one hour. Top: 5 second exposure time, bottom: 30 second exposure time.




Figure 5.3: Top: noise model for a 9th magnitude object. Solid lines denote noise contributions for airmass 1 and dashed lines denote noise contributions for airmass 2. Colours are as follows: red is noise due to the source, blue is noise due to the sky, green is noise due from electronic readout of the CCD, cyan is the scintillation noise, black is the combined noise. The total integration time is one hour. Bottom: similar to the top plot but for a 15th magnitude object.




Figure 5.4: Fractional rms for the data collected at Geneva on 2012-08-17. The noise model for the same conditions is shown with lines representing the dark current noise (black), read noise (green), sky noise (blue), scintillation noise (cyan), source noise (red) and total noise (magenta).

which cannot exist in the unsigned 16-bit electronics of the camera. Dark frames were corrected for the bias level, corrected for exposure time and combined using the median for each pixel into a master dark frame. Flat frames are bias and dark corrected, normalised to the median level and combined using the median in to a master flat frame. Science frames are all bias, dark and flat corrected using their own overscan regions and the master frames above. Photometry was performed using the pipeline discussed in Section 4.3.1. The fractional rms (Fig. 5.4) shows a limiting precision of 7.5 mmag. Due to Geneva’s low altitude of 420 m, the scintillation values are higher than the predictions for Paranal by around 30 %.

To ascertain if there was a red noise component, the limiting precision of the dataset was compared to the expected white noise behaviour (Fig. 5.5). For total exposure times ≤ 100 s show excellent white noise behaviour. Based on Fig. 5.5 we show that a NGTS unit can achieve sub-mmag precision from Geneva in 200 s, improving on the performance of the prototype, albeit on a relatively short dataset. As the behaviour follows the white noise behaviour for the entirety of the exposure time range we also determine that the red noise component during this observation must be less than 0.6 mmag. As the lightcurves only contain 600 points we cannot




Figure 5.5: Binned fractional rms of stars with magnitudes $11 \leq I \leq 9$. The dashed line indicates theoretical \sqrt{N} behaviour for white noise. Images were taken with 4 s exposures.

bin up the data any more.

5.3 Camera testing

During the on-sky tests at Geneva it was found that the bright stars exhibited unusual streaks, for example in Fig. 5.6. To determine the cause of this streaking, experiments were run on a single NGTS camera. We initially wanted to determine whether this effect was unique to the different camera used at Geneva.

In this section I discuss the process of analysing one of the NGTS cameras. The streaking behaviour seen with the on-sky data was characterised to determine the cause, and is discussed in Section 5.3.1. Section 5.3.2 describes the process of characterising the dark current. The measured behaviour with temperature is compared to the expected behaviour from the manufacturer, and the long term characteristics were measured.




Figure 5.6: Example raw image from the NGTS instrument at Geneva showing the bright stars streaking. The image has not been corrected for bias, dark or flat field effects.

5.3.1 Streak characterisation

Camera tests were also being undertaken by a team at Leicester University, where a more controllable optical environment was available. Tests other than dark current measurements were performed on the cameras and some strange artefacts were noticed during these tests. The camera being tested displayed an unexpected V-shaped feature across the centre of its dark frames, which changed shape depending on what readout settings were used for the exposure. This feature was first discovered in the camera taking the on-sky data, and was not visible with the camera being tested at Warwick. Another artefact that was apparent on both cameras being tested at Leicester and Warwick were cosmic ray signals which caused charge to spread along pixel columns. These cosmic ray hits have a large flux centred on a single pixel as they do not follow the same path as the rest of the flux through the telescope and interact with the CCD directly. They can cause the pixel they hit to saturate, so we initially attributed this charge bleed to saturation effects.

These streaks show an exponential decrease in the number of counts above the saturated pixel, and appear similar to problems with charge transfer efficiency (CTE) [see e.g. Dolphin, 2000]. This lack of efficiency causes some charge to be left behind on each vertical shift of a CCD readout, causing streaks to be left behind

objects. The streak was exponentially decreasing as the amount of charge left behind is a fraction of the original flux. The flux due to CTE effects $f(y)$ along a column from pixel 0 where the peak flux occurs is

$$f(y) = f_0 \delta^y \quad (5.4)$$

where f_0 is the flux at pixel 0, δ is the CTE value, y is the column pixel offset from pixel 0. To characterise the streaking behaviour, we used the dark frames collected for the analysis described in Section 5.3.2.

To characterise the hot or dead pixels of the CCD, which would interfere with analysing the cosmic rays, a master dark frame was constructed. This master dark frame was then subtracted from each dark frame in turn and we were left with the difference images for each exposure. Each image was then analysed for any pixels which were over a threshold value of 10000 ADU, which combined with a master dark frame value of ~ 2000 ADU gives a total threshold of 12000 ADU. This is much less than the electronic saturation point of the CCD but there should not be any pixels normally above this value without the presence of cosmic rays, so it filters frames with cosmic ray hits well. By removing any hot pixels or other similar defects we could be sure that only cosmic rays will breach this threshold. So that we could visually search for patterns in the saturation behaviour, images of the local regions around each saturated pixel were created. To diagnose the effect the regions were ranked by the brightest pixel value in the region under the assumption that this was the cosmic ray. By ranking by flux we could see that the streaking occurred only after the saturation point of the CCD.

We noticed that the CCD that we were studying had two different saturation levels for the different halves of the CCD. This was first discovered with the camera operating at Geneva when a sky flat was overexposed. The left half saturated at the electronic saturation level of 65535 ADU and the right half saturated at a lower value. Figure 5.7 shows a slice across a sky image taken at Geneva, $X = 1024$ is the mid-point. This shows that the CCD has two different saturation points, one for each half. The left half ($X < 1024$) saturates electronically whereas for the right half ($X > 1024$) the full well depth of the CCD is reached, at about $56000 \text{ADU} \times G(2.1) = 118000e^-$. The right half also shows some structure suggesting that the full well depth is not a constant across the entire frame.

To remove the effect of the CCD saturating at different points for the different halves, the cosmic ray regions were grouped into the left half ($X \leq 1024$) and right half ($X > 1024$) and considered separately. We found that some cosmic ray hits did




Figure 5.7: Horizontal slice across an overexposed flat frame taken from Geneva, averaged across 10 rows.

not exhibit streaks and some did (see Fig. 5.8 for two examples). We inferred the presence of a limiting flux value, which we attributed to the saturation point of that particular pixel.

The camera at Leicester with the V-shaped issue was also undergoing this test, taking lots of long dark exposures and waiting for a cosmic ray to hit. To ascertain whether the V-shape influenced the saturation point, the frames from this camera were analysed. It was found that saturated pixels were much more likely to streak if they were inside the V-shaped feature. This was further corroborated by shining a focussed light onto a region of the CCD that would change status from being inside the V region to outside depending on the shape of the V region (Fig. 5.9). Clearly the saturation only exhibits itself when the illuminated region is inside the V-shaped feature. It was thought that the other cameras where we observed the streaking effect but not the V-shaped region behaved as if the interior of the V-shaped region covered the whole CCD. Based on this result the cameras were returned to Andor and the streaking issue fixed by adjusting the clock voltages.

5.3.2 Dark current measurement

Specifications for the NGTS cameras were provided by the manufacturer, including a measurement of the dark current at different temperatures. To ensure the




Figure 5.8: Two regions from the dark frame series after the master dark frame has been subtracted. Left: peak flux value 18218 ADU, right: peak flux value 52836 ADU.




Figure 5.9: Illuminated region of the CCD tested at Leicester. The left panels were taken with a vertical shift speed of 76 μ s, the right panels were taken with a vertical shift speed of 38 μ s. Upper panels show the entire CCD, lower panels are centred and zoomed on the illuminated region.

performance characteristics of the camera matched the supplied specifications, a comparison of the dark current with temperature to the expected behaviour was performed.

To test the behaviour of the dark current with temperature, exposures were taken at different permutations of exposure time and temperature, with temperatures -10 , -30 and -50°C and exposure times 2, 20 and 200 seconds, logarithmically spacing the exposure times. One frame at -10°C and 200 seconds was saturated and therefore not used in the analysis. Each of the remaining 8 frames were bias-subtracted and converted to electrons per second. We compared the measured median flux values of these frames at each temperature with the specifications from Andor, shown in Fig. 5.10. A proposed relation was available from e2v¹ for the CCD used (model number CCD47-20), and is given by

$$\frac{Q_d}{Q_{d0}} = 1.14 \times 10^6 T^3 e^{-9080/T}, \quad (5.5)$$

and is also shown in Fig. 5.10. The Q_{d0} parameter represents the temperature at 293 K (20°C). In Fig. 5.10 this parameter was fitted for the black line and gives a best fit value of $Q_{d0} = 73381$ near the typical quoted value of 100000 represented by the lower of the two grey dashed lines in Fig. 5.10. Due to time constraints we did not measure the intermediate dark current values at -40 and -20°C and the lower temperatures were only achievable by a more expensive cooling.

5.4 Optimising the observing strategy

NGTS relies on precise photometry of stars to detect the shallow transits produced by the smaller planets. A study of the ideal exposure time for NGTS was performed, considering the number of stars that could be observed with mmag photometry, and the number of saturated stars. We estimate the saturation levels of NGTS in Section 5.4.1. Shorter exposure times allow brighter stars to be monitored without saturation, but they are relatively few in number. Longer exposure times allow mmag precision on a larger number of stars but they tend to be fainter, and more stars saturate. Section 5.4.2 describes finding the ideal exposure time for the project.

Contaminant flux from other nearby stars decreases the depth of any transits which may occur, and reduces the potential number of stars around which small planets can be found. Section 4.5.4 considers the number of stars which have con-

¹<http://www.e2v.com/e2v/assets/File/documents/imaging-space-and-scientific-sensors/44-47-20.pdf>




Figure 5.10: Median dark current values at different temperatures. Blue points are taken from the Andor specification sheet, red points are measured values. The dashed black line represents the behaviour given by e2v given in Eq. 5.5, fitted for the free parameter Q_{d0} . The dashed grey lines represent Eq. 5.5 but with the specified typical Q_{d0} of $100 \text{ ke}^- \text{ s}^{-1}$ (lower), and the specified maximum Q_{d0} of $200 \text{ ke}^- \text{ s}^{-1}$ (upper).

taminant flux from the neighbouring stars.

5.4.1 Saturation levels

As the total flux is calculable from the source magnitude, the saturation exposure time for an object can be calculated. When a point source is centred on a pixel, the largest amount of flux will go into that single pixel. This is often not the case, a slight misalignment may cause the peak flux to be measured in a neighbouring pixel. A realistic saturation level simulation must include this information.

We assumed a Gaussian PSF centred on an offset from the origin $(0, 0)$ up to a maximum of half a pixel away. The PSF of the instrument was used to define the width of the Gaussian, and set to 1 pixel. The 2D Gaussian PSF was integrated with a random offset (x_0, y_0) between -0.5 and 0.5 in each direction to get the flux in the central pixel (see Fig. 5.11 for a schematic diagram of the geometry).

$$f(x, y) = A \exp \left[\frac{(x - x_0)^2 + (y - y_0)^2}{2\sigma^2} \right] \quad (5.6)$$

The fraction of flux in the central pixel is the ratio of the integral within the




Figure 5.11: Schematic of a Gaussian centred at $(0.3, 0.3)$, and a FWHM of 1.5 pixels. The square represents the central pixel. We assume the pixel coordinates represent the centre of the pixel rather than the lower left corner.

central pixel to the integral of the Gaussian out to infinity. We simulated one million random offsets within the central pixel and computed the fraction of flux in the central pixel. A histogram of the results is shown in Fig. 5.12 with the most common value highlighted. From this analysis we can see that the most common flux fraction is 0.27 ± 0.03 .

We used this flux fraction to account for the observed flux from a point source filling multiple pixels. The saturation point is set by the full well depth of the CCD, and we assumed the gain is chosen such that the full well overflows before digital saturation occurs at $2^{16} - 1 = 65535$ ADU. We have a specification from Andor of 1.07×10^5 electrons. The components which increase the flux are: the sky background f_S measured in electrons per second; the dark current f_D but this is usually small; the flux from the object in question at magnitude m , scaled by the instrumental zero point m_0 and extinction coefficient k to give the flux observed at airmass 1, the brightest an object will appear; and ϵ the scaling factor calculated from the Monte Carlo simulation discussed previously. We assume the object is in the centre of the chip where the vignetting has minimal impact. Saturation occurs when the condition given in Eq. 5.7 is met.

$$\epsilon \left[f_S + f_D + 10^{(m_0 - m - k)/2.5} \right] t_{exp} \geq 107 \times 10^3 \quad (5.7)$$

The assumed sky level f_S was influenced by the prototype. We assume the sky




Figure 5.12: Distribution of pixel flux fractions from a Monte Carlo simulation. Dashed line represents the most common value.

Table 5.2: Fit coefficients for the saturation exposure time relationship.

Coefficient	Bright	Dark
α_0	7.773	6.081
α_1	-7.030	2.073
α_2	21.280	0.2552
α_3	-18.056	0
α_4	5.436	0

background at Paranal is not significantly fainter than at La Palma, and take dark sky as $50 \text{ e}^- \text{ pix}^{-1} \text{ s}^{-1}$ and “bright” sky as $450 \text{ e}^- \text{ pix}^{-1} \text{ s}^{-1}$. ϵ in Eq. 5.7 was taken as 0.27 encapsulating the fraction of flux that falls in the central pixel. The dark current was measured (Section 5.3.2) at -50°C to be $\sim 6 \text{ e}^- \text{ s}^{-1}$ which is negligible compared to the sky and source terms in Eq. 5.7 so we ignore the dark current. The magnitude at which saturation occurs is shown against exposure time in Figure 5.13, with two polynomial fits of form

$$f(\log_{10} t_{exp}) = \sum_{i=0} \alpha_i (\log_{10} t_{exp})^i, \quad (5.8)$$

with coefficients α_i are given in Table 5.2.




Figure 5.13: Saturation magnitude against exposure time for two different sky background values. Polynomial fits were made to the data of order 4 for bright sky background (red line), and 2 for a dark sky background (green line).

5.4.2 Exposure time optimisation

We use the saturation level and noise model at a specific exposure time to optimise the exposure time; we wanted to optimise for the number of high precision objects which are not saturated. Figure 5.14 shows the range of exposure times and source magnitudes which would achieve a precision of 1 mmag or higher. These targets are high priority for NGTS as they would meet our target for photometric precision, and would not be saturated as this would prevent the ability to perform photometry.

The Besançon galaxy model [Robin et al., 2003] was used for three target locations (Table. 5.3) spanning different galactic latitudes providing the predicted stellar distributions for these locations. A field with area 7.77 square degrees was used at each point representing the field of view of a single NGTS camera. The model results were queried for the I magnitudes of the objects in the frame, splitting the stars into dwarfs (class 5 using the category system defined by Robin et al. [2003]) and giants (class ≤ 5). Each object was classed as saturated, high precision, or otherwise based on the chosen exposure time and the results described previously in this section.

Figure 5.15 shows three aspects of optimising the exposure time. As the expo-




Figure 5.14: The grey region encloses the region of parameter space where any objects are not saturated in dark time, and have a precision of ≥ 1 mmag. The dotted line indicates the saturation limit during bright time.

sure time increases, the number of high precision objects (fractional noise ≤ 1 mmag) increases as the read noise decreases. The lifetime for the shutter was a concern as the quoted lifetime of one million operations is a possibly expensive time constraint as it would require returning the cameras to the manufacturer from Chile to Belfast for each breakage, which could be considerable with twelve cameras. It would reduce the ability to run robotically for any reasonable period and have a detrimental effect for the project. With longer exposure times the number of shutter operations decreases (Fig. 5.15, middle panel), reducing the overall cost of maintenance and improving the ability for the observatory to run without expensive intervention from a member of the consortium. A further benefit of exposing for longer is not plotted here, but the storage space required to hold the raw images decreases drastically with longer exposures. The main disadvantage to increasing the exposure time is that more bright objects become saturated (Fig. 5.15, bottom panel). The bright objects are more valuable as planet search targets as they have inherently less fractional noise, and are easier to follow up with spectroscopy, so can be confirmed quicker and with more confidence. Bright targets also make excellent targets for exoplanet characterisation missions.

The number of high precision objects does not change significantly over the entire range, implying that there is no real advantage to exposing for longer as the overall gain from 5 to 50 seconds is not significant. The number of saturated stars increases dramatically especially during bright time after ~ 10 seconds, where the number of saturated dwarfs is 10, rising from 5 at 5 seconds. The number of shutter operations roughly halves in this same exposure time difference, but the number of high precision objects doesn't change significantly. The optimal exposure time was left to judgement from the consortium but a value around 10 seconds seemed to be optimal.

The simulations were run again with a slower readout speed as this reduces the read noise and potentially reduces the artefacts present from the camera, but this has an adverse affect due to a longer fraction of the duty cycle devoted to reading the CCD instead of collecting photons, and decreases the number of high precision objects by around 20%.

5.5 Planet catch

Based on the predicted noise models (Section 5.2.1) and information about the observing strategy (Section 5.4) we estimated the predicted number of planets NGTS would detect. This work was undertaken as a collaboration with my supervisor who simulated the instrument throughput and applied the results to models of the galaxy, a colleague João Bento who simulated the period sensitivity function for NGTS incorporating weather conditions from Paranal, and a colleague Tom Loudon who analysed the significance of the radial velocity signals.

Simulations were performed by my supervisor, based on the Besançon galaxy model [Robin et al., 2003], which predicts the distribution of spectral types and their magnitudes based on particular galactic coordinates (Table. 5.3). The predicted throughput for the NGTS instrument was calculated from predicted response functions of the optical and electronic components taken from manufacturer specifications or assumed values (see Fig. 4.3, black line for the total throughput), along with the sky transmission at airmass 1.5. The throughput was combined with the spectral response of each spectral type taken from Pickles [1998] to give the flux that would be observed per spectral type for a given exposure time. The signal to noise for each object is then calculated based on Poisson noise and a white noise component to simulate the photometric quality reaching a maximum for the very brightest stars:

$$\frac{S}{N} = \frac{S}{\sqrt{S} + \sigma_w(S)} \quad (5.9)$$




Figure 5.15: Top: number of high precision objects (fractional noise ≤ 1 mmag) for three target fields for a single telescope. Middle: projected thousands of shutter operations per year based on the given exposure time for a typical NGTS target field with 1200 observing hours available. Bottom: number of saturated stars for the same three target fields for a given exposure time. Dashed lines represent the number of saturated stars in bright time, and solid lines represent the number of saturated stars in dark time. The plot has been truncated at 100 objects to increase the visibility of the low exposure time range. In each relevant plot the black lines represent selecting only the dwarf type stars, the grey lines represent giant class stars.

Table 5.3: Coordinates for the three simulations run for the NGTS planet catch.

Field	α	δ	l	b
1	05:21:06.78	-24:29:05.72	227	-30
2	03:58:09.67	-44:43:02.68	251	-49
3	11:57:17.00	-44:49:03.60	293	+17

Table 5.4: Coefficients for Eq. 5.10.

$R_p(R_E)$	k_P	β	P_0 (days)	γ
2 – 4	0.064 ± 0.040	0.27 ± 0.27	7.0 ± 1.9	2.6 ± 0.3
4 – 8	0.0020 ± 0.0012	0.79 ± 0.50	2.2 ± 1.0	4.0 ± 1.2
8 – 32	0.0025 ± 0.0015	0.37 ± 0.35	1.7 ± 0.7	4.1 ± 2.5
2 – 32	0.035 ± 0.0023	0.52 ± 0.25	4.8 ± 1.6	2.4 ± 0.3

where S is the number of photons counted, and $\sigma_w(S)$ is a signal-dependant white noise component. This parameter was chosen to match the results from the noise models calculated for NGTS.

I took an assumed underlying distribution of planets from the Kepler project, from which planets were randomly drawn. Howard et al. [2012] predict 0.165 ± 0.008 planets per star incorporating orbital periods up to 50 days and planetary radii from $2R_E$ to $32R_E$. For each star in the simulation results a random number was computed and compared it to the global occurrence rate 0.165, if the number is less then a period and radius value were chosen. Eq. 5.10 describes the predicted occurrence in orbital period, and Eq. 5.11 describes the predicted occurrence in planetary radius, Table 5.4 gives the constants used for Eq. 5.10, $\alpha = -1.92 \pm 0.11$ and $k_R = 2.9^{+0.5}_{-0.4}$.

$$\frac{df(P)}{d \log P} = k_P P^\beta (1 - e^{-(P/P_0)^\gamma}) \quad (5.10)$$

$$\frac{df(R)}{d \log R} = k_R R^\alpha \quad (5.11)$$

We use the transit probability $p_{tr} = (R_\star + R_p)/a$ for circular orbits and including grazing transits to assess whether the synthetic object is transiting or not. The detectability of a transiting planet based on its orbital period was simulated from Paranal, incorporating likely weather conditions and the number of transits required before a detection is believed [Bento, 2012]. By taking the average period window and simplifying, we used the assumed period window profile shown in Fig. 5.16, which is used to reject planets based on their orbital period.




Figure 5.16: Assumed period window function for the NGTS project.

To determine the significance of a transit-like signal in a noisy source, the total signal to noise for the duration of a single transit was calculated from the NGTS noise model. The simulation results had a fixed exposure time so the number of exposures N_{exp} would be taken during transit is given as

$$N_{exp} = T_{tr}/t_{exp} \quad (5.12)$$

where T_{tr} is the transit duration and t_{exp} the exposure time of the simulation. The transit duration is given as

$$T_{tr} = \frac{P}{\pi} \sin^{-1} \left[\frac{R_{\star} \sqrt{(1+k)^2 - b^2}}{a \sin i} \right] \quad (5.13)$$

where P is the orbital period, R_{\star} is the stellar radius, a is the orbital separation, $k = R_p/R_{\star}$ the ratio of the object radii, $b = a \cos i/R_{\star}$ is the impact parameter for a circular orbit, and i is the inclination [Seager, 2011]. The signal to noise ratio therefore increases by $\sqrt{N_{exp}}$.

The signal to noise ratio calculated from the galaxy model simulations did not account for the scintillation of the atmosphere, so photometric quality of each star was degraded based on the assumed exposure time in the simulation and parameters from the Paranal observatory. I used the scintillation equation from Dravins et al.

[1998] as described in Section 4.3.1 reprinted here:

$$M_{scin} = 0.004D^{-2/3}X^{7/4}e^{-h/H}(2t_{exp})^{-1/2} \quad (5.14a)$$

$$\sigma_{scin} = 1 - \left(10^{-M_{scin}/2.5}\right) \quad (5.14b)$$

For an exposure time of 600 seconds the scintillation value for NGTS is $\sigma_{scin} = 0.000224$. This white noise component was added before the signal to noise is scaled by the number of exposures during transit, as it behaves as white noise and so is reduced by combining measurements. At this point a red noise component is added, which does not decrease with combining observations. Since we find no hint of red noise from the system tests in Geneva (see Section 5.2.2) down to 0.3 mmag we set the red noise floor to this value, as otherwise it would have been apparent in the test data.

Transit significance

The significance of a transit is

$$S = \frac{\delta}{\frac{S}{N}} = \left(\frac{R_p}{R_\star}\right)^2 \frac{N}{S} \quad (5.15)$$

which was calculated for each planet-star system. We next calculated the number of transits which would be observable, based on the observing strategy chosen. The colour scale for Fig. 4.27 represents the number of yearly hours that particular point of sky is visible for including weather effects. Each NGTS field is only observable for about two thirds of the year as it is only in the sky for this time. We used simulations of the NGTS yearly coverage (Fig. 4.27, Bento [2012]) for the chosen fields (Table 5.3) and assumed 800 hours per field. The number of transits possible to observe n_{tr} is therefore

$$N_{tr} = \frac{N_{obs}}{P} \quad (5.16)$$

where N_{obs} represents the number of hours the object is observed for. Fisher's method was used for combining p -values as we have N_{tr} measurements of transits with significance S . The p -value is related to the significance of a measurement via

$$p = 2(1 - \Phi(S)) \quad (5.17)$$

where $\Phi(S)$ is the cumulative distribution function of a normal distribution function with zero mean and unit standard deviation, tested at point S . Fisher's method

measures a parameter X where

$$X^2 = -2 \sum_i^k \ln p_i \quad (5.18)$$

This X parameter is distributed as a χ^2 distribution with $2k$ degrees of freedom. For this study $k = N_{obs}$, so the total probability of detection can be established. We define the detection threshold as 6σ and if a synthetic system has a detection significance greater than this counts as detected.

The final stage of planet confirmation is to observe the radial velocity signature to rule out false positives, and for an independent confirmation of the planet signal. A colleague Tom Loudon took the parameters of the systems calculated to be detectable by NGTS, and assessed whether they would be conformable using HARPS on La Silla. To calculate planetary masses, a mass-radius relation was assumed of the form $(m_p / M_\oplus) = (R_p / R_\oplus)^2$ and calculated the velocity semi-amplitude K

$$\frac{m_p}{(m_p + M_\star)^{2/3}} = \frac{K \sqrt{1 - e^2}}{\sin i} \left(\frac{P}{2\pi G} \right)^{1/3} \quad (5.19)$$

[Seager, 2011]. Figure 5.17 shows the planets which are detectable by NGTS along with the Kepler candidates and other confirmed exoplanets; the simulation was oversampled by 12 times so we show 1/12 of the points in the figure. The dashed lines represent the detection limits for HARPS (dashed) and ESPRESSO (dotted) in ten hours of observations. The majority of the Kepler planets were found to be un-confirmable due to their low predicted mass and faint host stars, whereas most of the simulated NGTS planets were found to be confirmable. This reflects the optimal design of NGTS for follow up potential: by optimising the instrument to target brighter stars, the survey was designed for synergy with current and future radial velocity facilities.

We used the detection limits to determine whether the candidates were confirmable, objects fainter than the limiting V magnitude for radial velocity confirmation were rejected. ESPRESSO will be available to NGTS so we take this limit for the NGTS candidate planets. We use the HARPS radial velocity limit to determine the confirmability of the Kepler candidates assuming a similar performance between HARPS and HARPS-N. To independently study the population of detectable planets independently of their host stars, the planets were plotted in radius-magnitude space (Fig. 5.18). We plot the I magnitude of the stars hosting the planets to place the population in context with future red or near infra-red characterisation missions




Figure 5.17: Red points: random sample of the simulated planet catch for NGTS, blue points: Kepler candidates, green points: confirmed planets detected through the radial velocity method. Black lines indicate (dashed) the detection limit for HARPS and (dotted) the predicted detection limit for ESPRESSO. Data courtesy of Tom Loudon.




Figure 5.18: Kepler candidates (blue), NGTS predicted planets (red) and RV detected planets (green), which are confirmable plotted in more natural units for NGTS. Vertical lines denote the planet classes defined in Howard et al. [2012].

(e.g. CHEOPS [Broeg et al., 2013] or JWST). We see that NGTS is predicted to be able to detect super Earths and Neptunes around bright ($I \lesssim 11$) stars. The predicted numbers of planets detectable by NGTS and confirmable by HARPS or ESPRESSO are given in Table 5.5. We predict on order 200 Neptunes and 30 super Earths with the next generation follow up instrument ESPRESSO.

5.6 Summary

In this section I updated the noise model of the previous chapter to account for the total integration time, by binning measurements to reduce the noise. The measured behaviour from the Geneva instrument was found to be white noise dominated down to the sub-mmag level, which is likely due to the low altitude of the site causing

Table 5.5: Predicted planet catch for NGTS

Planet class	Super-Earths (2-4 R_{\oplus})			Neptunes (4-8 R_{\oplus})	Jupiters (8-32 R_{\oplus})
Confirmed transiting planets	3			4	120
Kepler candidates confirmable with HARPS-N in 10h	0			7	105
NGTS candidates confirmable with HARPS in 10h	1			21	315
Kepler candidates confirmable with HARPS-N in 20h	1			21	130
NGTS candidates confirmable with HARPS in 20h	3			42	429
NGTS candidates confirmable with ESPRESSO in 10h	28			174	657
NGTS candidates confirmable with ESPRESSO in 20h	62			323	725
Bright ($I < 11$) NGTS candidates confirmable with ESPRESSO in 10h	17			17	21

increased scintillation noise. We are confident in our ability to combine images to reduce the noise and determine that the level of red noise must be < 0.6 mmag.

On sky data from Geneva showed the bright stars leaving trails of charge during readout of the CCD. By collecting dark images and waiting for cosmic rays to saturate pixels, the count rate at which the streaking occurred was characterised. It was found that the streaks only occurred inside an unusual region of the CCD caused by incorrectly set voltages. The cameras were returned to the manufacturer, and the saturated stars no longer caused streaks. The dark frames collected were used to characterise the dark current behaviour, which showed a good match in the temperature behaviour with the camera manufacturer measurements and model performance.

The noise model was used to calculate the magnitude at which stars would have a precision of 1mmag in one hour, and the magnitude of saturation to determine the number of priority objects for a range of science exposure times. We applied this to three potential fields for NGTS, each with a high number of nights with favourable weather conditions, and presented three considerations for the project: the number of stars we can achieve sub-mmag photometry, the number of saturated stars and the number of shutter operations. We find a balance to be had at around 10s. The noise model was used to determine the planet catch for NGTS. We drew planetary systems from the Howard et al. [2012] model of the underlying occurrence rate of planets from Kepler, with stars from the Besançon galaxy model, applied signal to noise detectability arguments and radial velocity signal significance arguments to determine the number and types of planets which are detectable by NGTS. A total of 200 Neptunes and 30 super Earths were found to be confirmable by the next generation spectrograph ESPRESSO.

Chapter 6

Conclusions and future work

In this thesis I present a study of ground based surveys searching for transiting extrasolar planets, both current, with the Wide Angle Search for Planets (WASP) the most successful ground based transiting exoplanet survey to date, and future, with the Next Generation Transit Survey (NGTS) designed to find small planets around bright stars.

I determine the underlying population of hot Jupiters by correcting for selection biases in the WASP project, and propose a new model for the underlying period distribution. Looking to the future, I introduce NGTS, discuss my role in the development of the instrument, and present predictions of the number and types of planets that will be detected.

6.1 Determining the underlying population of hot Jupiters with the WASP project

Jupiter sized planets form in protoplanetary disks beyond the snow line, where the density of disk matter is high. The presence of such planets within 0.1 AU of their host stars, the hot Jupiters, suggests that these planets must have migrated inwards to their current locations, but the primary mechanism for this is currently unknown. These hot Jupiters are markers for the dominant migration mechanism. The WASP project is ideally suited to study these objects as it has detected the largest sample of hot Jupiters to date, but biases exist in the observed population which must be corrected for before conclusions can be drawn. We have corrected for these selection biases by simulating transiting systems and calculating the fraction recovered, to invert the observed population of hot Jupiters to the underlying distribution.

A clear pile up in orbital period is observed in the underlying hot Jupiter

distribution at 3-5 days, and we place a joint constraint with data from the Kepler project on the underlying distribution of orbital period, to model the period distribution of giant planets up to orbital periods of 50 days. We propose a new model for the orbital period distribution as a rising power law with index 1.0 ± 0.3 and a Gaussian excess at 3.7 ± 0.1 days with width 0.7 ± 0.2 days.

When considering the ratio of star-planet separation to the Roche limit, the corrected distribution exhibits a peak at a ratio of 3.5 with an inner edge at 2, which is consistent with the dynamical scattering migration scenario, where planets are excited onto eccentric orbits through dynamical interactions with other bodies in the system, and their orbits subsequently circularise through tidal interactions with their host stars. This observed pile up in orbital period alternatively may be due to truncation at the inner edge of protoplanetary disks halting disk migration, where the planets no longer experience an inward migration torque. In this thesis, I have shown undeniable evidence for a pile up at orbital periods of 3.7 days in the underlying population of hot Jupiters, and any proposed theory for planetary migration must incorporate this feature.

6.2 The Next Generation Transit Survey

The Next Generation Transit Survey (NGTS) is a new wide-field ground-based transiting exoplanet survey aimed at discovering super Earth and Neptune-sized exoplanets around nearby G, K and M stars. Ground based transiting exoplanet surveys have discovered a large sample of planets, typically Jupiter or Saturn-sized which have challenged our understanding of how such planets form. Space based surveys such as Kepler have detected smaller planets, typically super Earth objects. Conversely, intermediate Neptune-sized planets have not been found in nearly the same number as the super Earths or Jupiters, as they are fewer in number than the super Earths, and harder to detect than the Jupiters. With NGTS we aim to find a significant population of Neptune sized planets and smaller, orbiting around nearby bright stars.

The transit technique is powerful when coupled with radial velocity estimates, as the radius and mass are independently determined which together provide the planetary density. In the current population of planets however, the bulk densities of Neptunes are not well understood, so NGTS will provide targets for current and next generation spectrographs such as HARPS or ESPRESSO to perform radial velocity analysis and determine the masses. The atmospheres of smaller planets are also not well understood, as they orbit stars too faint for atmospheric characterisation. By

optimising the survey for bright stars, we increase the precision of mass estimates possible from radial velocity surveys, and the planets detected by NGTS can be used as targets for future dedicated exoplanetary characterisation missions such as CHEOPS or EChO to determine the atmospheric compositions.

In this thesis I present a study of the noise characteristics of NGTS through the analysis of two prototype instruments, and find that bright stars are limited by scintillation noise in the atmosphere, and that after combining exposures we show white noise behaviour down to $< 1\text{mmag}$ on the bright targets, which is the required precision to detect Neptunes around K stars, and super Earths around M dwarfs. Based on this characterisation I present a simulation of the predicted planet catch of NGTS and show that 200 Neptunes and 30 super Earths will be detected, and 20 of each category will orbit stars $I < 11$ and are therefore bright enough for atmospheric characterisation.

6.3 Future work

The observing strategy for WASP was changed to track single fields rather than sampling a group of fields per night, and so the selection biases in this new data will likely be different, most probably the sensitivity to longer periods and smaller radii will be increased. This will require further analysis of the sensitivity, as the occurrence rate calculated whilst including these planets, which are likely subject to different selection biases, will be different. Similarly the mapping of the selection effects described in this thesis can be applied to other projects, in particular NGTS due to similarities in the analysis method, where it will be useful at two points of the project: it can be used for quality control during the initial stages of the project, as a tool to drive further development; and after planets have been detected these methods can again be used to determine the underlying distribution in an expanded parameter space.

We can also expand the parameter space of the WASP project, as one of the selection cuts designed to reject likely false positive systems, $P(R_p)$, potentially restricts the number of inflated planets detectable by WASP. The existing WASP database could be queried for objects which explicitly fail this selection cut, but otherwise would be selected. A large number of false positive objects such as eclipsing binary systems would likely be found, but there may be some inflated planets that have been missed by the consortium.

An extension to the project would be to test the sensitivity of WASP to eccentric planets, by repeating the sensitivity analysis but allowing elliptical or-

bits, either proposing them with an even distribution (and rejecting planets with simulated orbits inside their minimum allowed periastron distance), or by drawing randomly from the expected distribution. The eccentricity distribution of the recovered synthetic models would then allow the sensitivity of WASP to eccentric planets to be determined. Similarly the sensitivity of WASP to circumbinary planets can be tested. Planets orbiting a binary star usually experience transit timing variations and transit duration variations, which would likely cause a non-detection by the WASP planet search pipeline. The WASP planet search implementation ORION searches for exact periodicity, so deviations from this behaviour such as periodic variations in the transit times may affect the sensitivity to these systems. The inclusion of simulated systems with transit timing variations (TTVs) would allow the sensitivity of WASP to these systems to be determined. This could lead to further analysis of the current WASP dataset with an algorithm which is sensitive to TTVs and possible WASP circumbinary planets. NGTS should be sensitive to circumbinary planets, as with high cadence observations TTVs can be measured with high accuracy. This will allow NGTS to detect multiple planet systems and map their population.

Our analysis involved inverting the distribution of observed planets to calculate the underlying distribution, but this process can be reversed, by applying the sensitivity maps to a model underlying distribution, or synthesised planetary population, and comparing the observed planets with the sampled planets with selection biases applied. To this end we will be making the sensitivity maps available, to help constrain population synthesis models.

It is commonly thought that to achieve high enough precision to detect small planets, observations must be made from space, but these high precision measurements come at a cost: both financially, as space missions are significantly more expensive than NGTS, and in the stars they target, which are often too faint for further characterisation. The limits of precision for ground-based wide-field surveys are being pushed with NGTS, but there may be further to go with the limiting precision. For example scintillation noise can be decreased with larger apertures, or longer exposure times and so the next limit to precision will be reached. To continue to produce high precision observations in wide-field surveys from the ground, the fundamental limit of detection needs to be determined.

Appendix A

Sensitivity map trends

Figures A.1 to A.4 show sensitivity maps for each of the planets used to determine the sensitivity of WASP for the planet hosting stars, ordered by parameters assessed for trends (see Section 2.5.1).

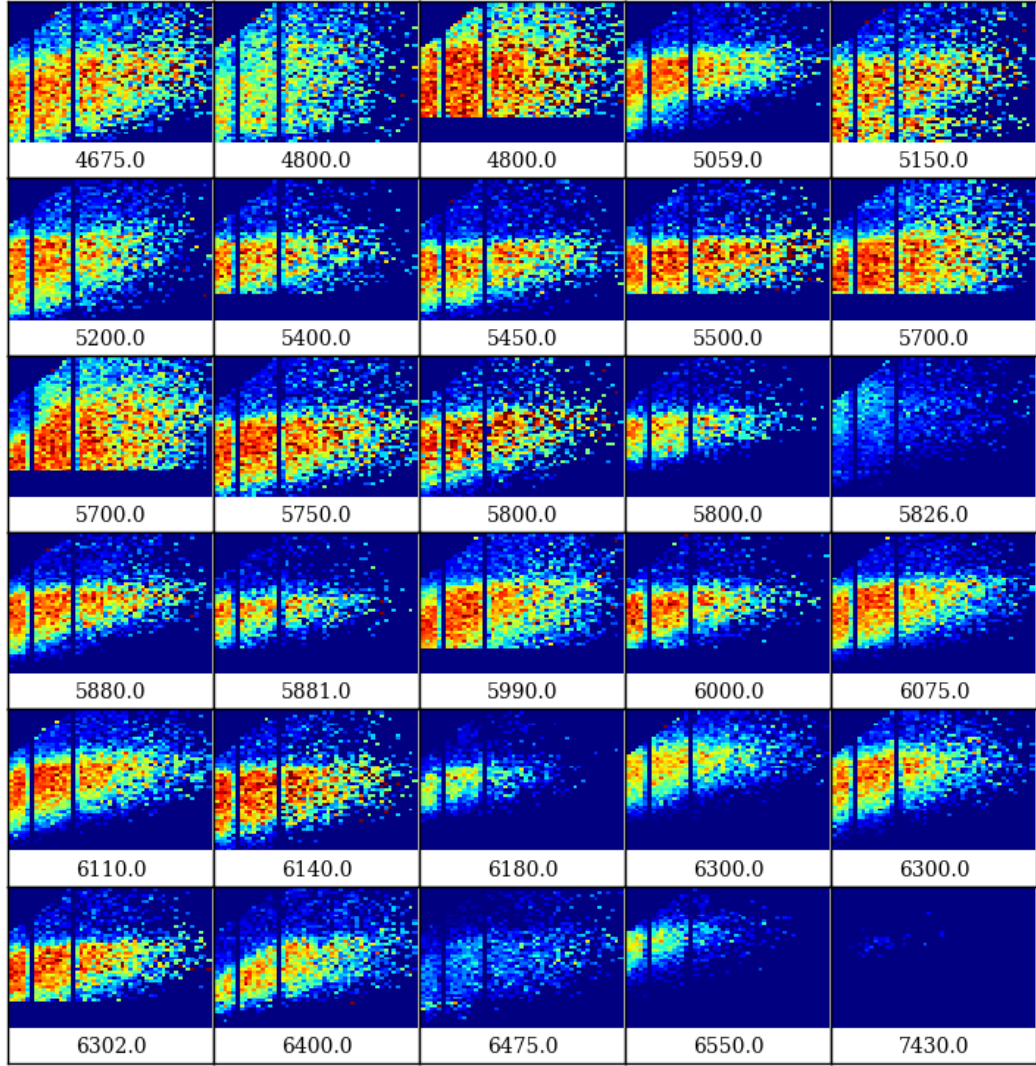


Figure A.1: Sensitivity maps in order of T_{eff} . See Fig. 2.12 for a description.

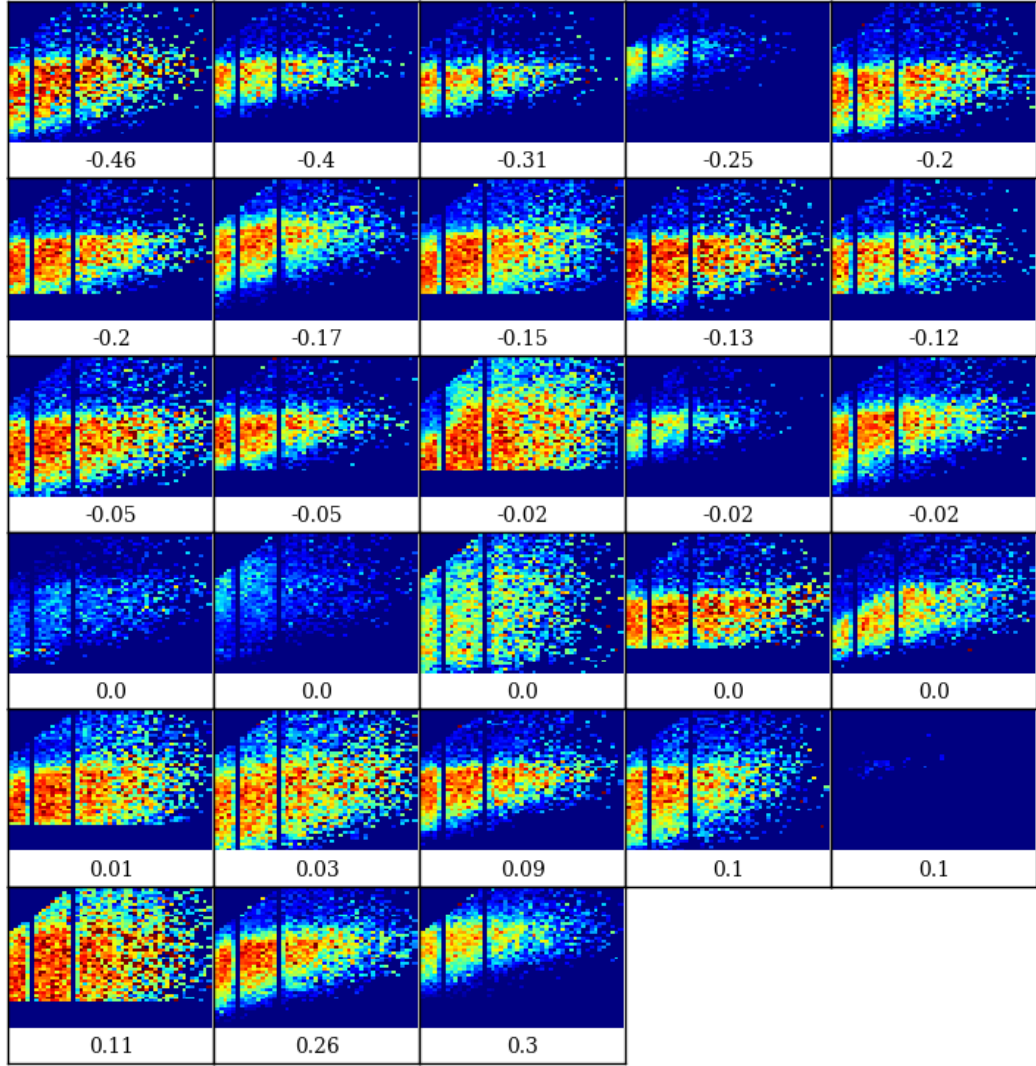


Figure A.2: Sensitivity maps in order of Metallicity. See Fig. 2.12 for a description.

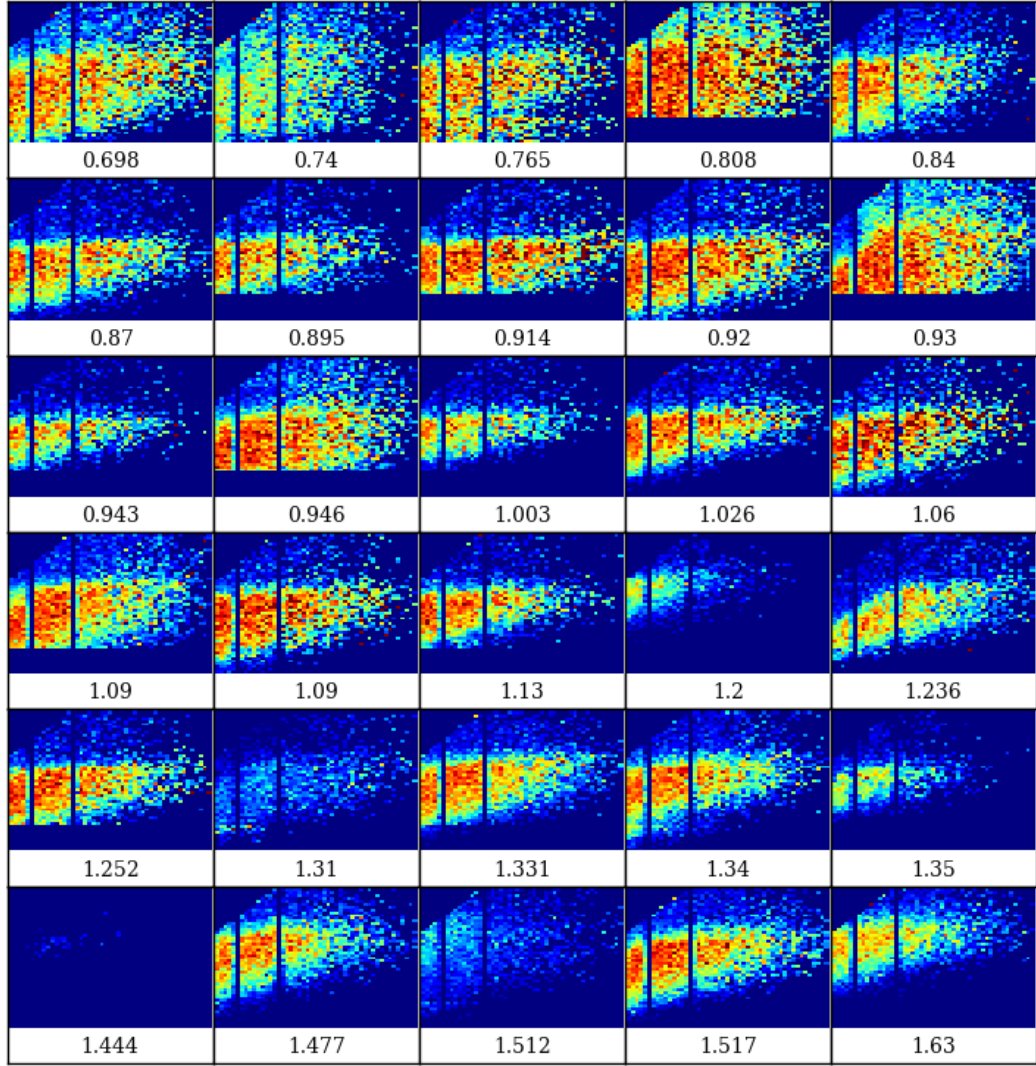


Figure A.3: Sensitivity maps in order of R_{\star} . See Fig. 2.12 for a description.

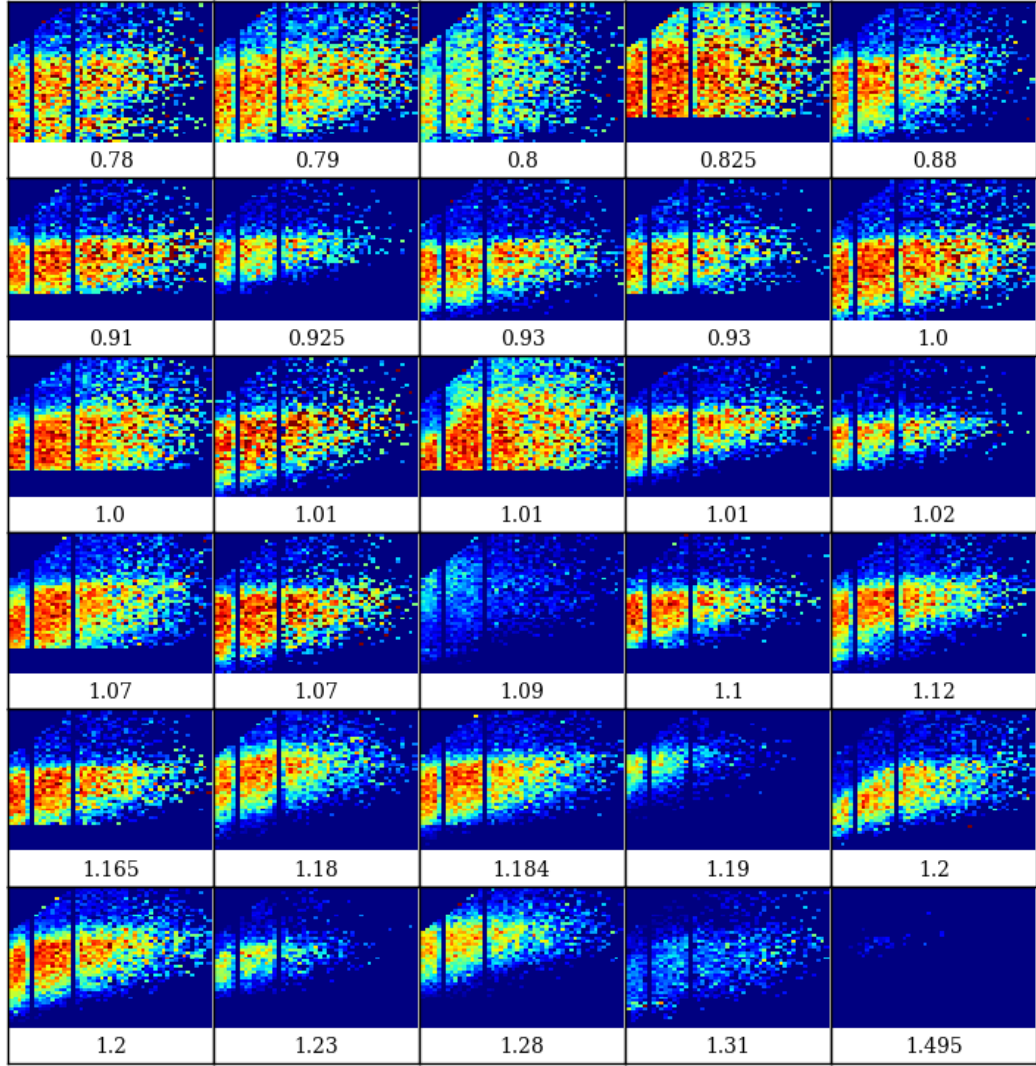


Figure A.4: Sensitivity maps in order of M_{\star} . See Fig. 2.12 for a description.

Bibliography

- Albrecht, S., et al., 2012, ApJ, 757, 18
- Alibert, Y., Mordasini, C., Benz, W., Winisdoerffer, C., 2005, A&A, 434, 343
- Ammons, S. M., Robinson, S. E., Strader, J., Laughlin, G., Fischer, D., Wolf, A., 2006, ApJ, 638, 1004
- Anderson, D. R., et al., 2010, ApJ, 709, 159
- Anderson, D. R., et al., 2013, ArXiv e-prints
- Armitage, P. J., 2010, Astrophysics of Planet Formation, Cambridge University Press
- Bakos, G. Á., et al., 2011, ApJ, 742, 116
- Baraffe, I., Selsis, F., Chabrier, G., Barman, T. S., Allard, F., Hauschildt, P. H., Lammer, H., 2004, A&A, 419, L13
- Barbieri, C., 2007, Fundamentals of Astronomy, Taylor & Francis
- Barnes, R., 2010, Formation and Evolution of Exoplanets
- Baruteau, C., Meru, F., Paardekooper, S.-J., 2011, MNRAS, 416, 1971
- Beaulieu, J.-P., et al., 2006, Nat, 439, 437
- Benedict, G. F., et al., 2006, ApJ, 132, 2206
- Bento, J., 2012, Research and development of ground-based transiting extrasolar planet projects, Ph.D. thesis
- Beuermann, K., et al., 2010, A&A, 521, L60
- Birney, D. S., Gonzalez, G., Oesper, D., 2006, Observational Astronomy - 2nd Edition, Cambridge University Press

- Borucki, W., et al., 2009, in Pont, F., Sasselov, D., Holman, M. J., eds., IAU Symposium, vol. 253 of *IAU Symposium*, p. 289
- Borucki, W. J., et al., 2010, *Science*, 327, 977
- Boss, A. P., 2002, *ApJ*, 567, L149
- Broeg, C., et al., 2013, in European Physical Journal Web of Conferences, vol. 47 of *European Physical Journal Web of Conferences*, p. 3005
- Brown, D. J. A., Collier Cameron, A., Hall, C., Hebb, L., Smalley, B., 2011a, *MNRAS*, 415, 605
- Brown, D. J. A., et al., 2012, *MNRAS*, 423, 1503
- Brown, T. M., Latham, D. W., Everett, M. E., Esquerdo, G. A., 2011b, *ApJ*, 142, 112
- Casertano, S., et al., 2008, *A&A*, 482, 699
- Cash, W., 1979, *ApJ*, 228, 939
- Chabrier, G., Baraffe, I., Leconte, J., Gallardo, J., Barman, T., 2009, in Stempels, E., ed., 15th Cambridge Workshop on Cool Stars, Stellar Systems, and the Sun, vol. 1094 of *AIP Conf. Proc.*, p. 102
- Chan, T., Ingemr, M., Winn, J. N., Holman, M. J., Sanchis-Ojeda, R., Esquerdo, G., Everett, M., 2011, *ApJ*, 141, 179
- Chaplin, W. J., et al., 2011, *Science*, 332, 213
- Charbonneau, D., Brown, T. M., Latham, D. W., Mayor, M., 2000, *ApJ*, 529, L45
- Claret, A., 2000, *A&A*, 363, 1081
- Collier Cameron, A., et al., 2006, *MNRAS*, 373, 799
- Collier Cameron, A., et al., 2007a, *MNRAS*, 380, 1230
- Collier Cameron, A., et al., 2007b, *MNRAS*, 375, 951
- Collier Cameron, A., et al., 2010, *MNRAS*, 407, 507
- Cumming, A., 2004, *MNRAS*, 354, 1165
- Cumming, A., Marcy, G. W., Butler, R. P., 1999, *ApJ*, 526, 890

- Cumming, A., Butler, R. P., Marcy, G. W., Vogt, S. S., Wright, J. T., Fischer, D. A., 2008, *PASP*, 120, 531
- Dawson, R., 2013, Presented at the IAU symposium 299
- Dawson, R. I., Johnson, J. A., 2012, *ApJ*, 756, 122
- Dawson, R. I., Murray-Clay, R. A., 2013, *ApJ*, 767, L24
- Désert, J.-M., Vidal-Madjar, A., Lecavelier Des Etangs, A., Sing, D., Ehrenreich, D., Hébrard, G., Ferlet, R., 2008, *A&A*, 492, 585
- Désert, J.-M., et al., 2011, *ApJS*, 197, 14
- Dodson-Robinson, S. E., Veras, D., Ford, E. B., Beichman, C. A., 2009, *ApJ*, 707, 79
- Dolphin, A. E., 2000, *PASP*, 112, 1397
- Dong, S., et al., 2006, *ApJ*, 642, 842
- Dravins, D., Lindegren, L., Mezey, E., Young, A. T., 1998, *PASP*, 110, 610
- Dressing, C. D., Charbonneau, D., 2013, *ApJ*, 767, 95
- Dumusque, X., et al., 2012, *Nat*, 491, 207
- Enoch, B., Haswell, C. A., Norton, A. J., Collier-Cameron, A., West, R. G., Smith, A. M. S., Parley, N. R., 2012, *A&A*, 548, A48
- Fischer, D. A., Valenti, J., 2005, *ApJ*, 622, 1102
- Ford, E. B., Rasio, F. A., 2006, *ApJ*, 638, L45
- Fortney, J. J., Saumon, D., Marley, M. S., Lodders, K., Freedman, R. S., 2006, *ApJ*, 642, 495
- Fortney, J. J., Marley, M. S., Barnes, J. W., 2007, *ApJ*, 659, 1661
- Fortney, J. J., et al., 2011, *ApJS*, 197, 9
- Fressin, F., Guillot, T., Morello, V., Pont, F., 2007, *A&A*, 475, 729
- Fressin, F., et al., 2012, *Nat*, 482, 195
- Fressin, F., et al., 2013, *ApJ*, 766, 81

- Galicher, R., Marois, C., Zuckerman, B., Macintosh, B., 2013, *ApJ*, 769, 42
- Gilliland, R. L., et al., 2011, *ApJS*, 197, 6
- Gilman, D., Fuglister, F., Mitchell Jr, J., 1963, *Journal of the Atmospheric Sciences*, 20, 182
- Goldreich, P., Tremaine, S., 1979, *ApJ*, 233, 857
- Gomes, R., Levison, H. F., Tsiganis, K., Morbidelli, A., 2005, *Nat*, 435, 466
- Gould, A., Morgan, C. W., 2003, *ApJ*, 585, 1056
- Gould, A., Dorsher, S., Gaudi, B. S., Udalski, A., 2006, *Acta Astron.*, 56, 1
- Gray, D. F., 1992, *The observation and analysis of stellar photospheres*.
- Greiss, S., et al., 2012, *ApJ*, 144, 24
- Hartman, J. D., et al., 2011, *ApJ*, 742, 59
- Haswell, C. A., 2010, *Transiting Exoplanets*
- Hellier, C., et al., 2009, *Nat*, 460, 1098
- Hellier, C., et al., 2012, *MNRAS*, 426, 739
- Herbst, W., Mundt, R., 2005, *ApJ*, 633, 967
- Høg, E., et al., 2000, *A&A*, 355, L27
- Howard, A. W., et al., 2012, *A&AS*, 201, 15
- Howell, S. B., 2006, *Handbook of CCD Astronomy*, Cambridge University Press
- Husnoo, N., Pont, F., Mazeh, T., Fabrycky, D., Hébrard, G., Bouchy, F., Shporer, A., 2012, *MNRAS*, 422, 3151
- Ida, S., Lin, D. N. C., 2004, *ApJ*, 604, 388
- Innanen, K. A., Zheng, J. Q., Mikkola, S., Valtonen, M. J., 1997, *ApJ*, 113, 1915
- Janesick, J., Blouke, M., 1987, *Sky and Telescope*, 74, 238
- Kalas, P., et al., 2008, *Sci*, 322, 1345
- Kane, S. R., Collier Cameron, A., Horne, K., James, D., Lister, T. A., Pollacco, D. L., Street, R. A., Tsapras, Y., 2004, *MNRAS*, 353, 689

- Kley, W., Crida, A., 2008, *A&A*, 487, L9
- Knutson, H. A., Charbonneau, D., Allen, L. E., Burrows, A., Megeath, S. T., 2008, *ApJ*, 673, 526
- Kokubo, E., Ida, S., 2002, *ApJ*, 581, 666
- Kovács, G., Zucker, S., Mazeh, T., 2002, *A&A*, 391, 369
- Kovács, G., Bakos, G., Noyes, R. W., 2005, *MNRAS*, 356, 557
- Kozai, Y., 1962, *ApJ*, 67, 591
- Kuchner, M. J., Lecar, M., 2002, *ApJ*, 574, L87
- Kurucz, R. L., 1979, *ApJS*, 40, 1
- Lin, D. N. C., Bodenheimer, P., Richardson, D. C., 1996, *Nat*, 380, 606
- Lissauer, J. J., et al., 2011, *ApJS*, 197, 8
- Lucas, P. W., Roche, P. F., 2000, *MNRAS*, 314, 858
- Mandel, K., Agol, E., 2002, *ApJ*, 580, L171
- Marcy, G., Butler, R. P., Fischer, D., Vogt, S., Wright, J. T., Tinney, C. G., Jones, H. R. A., 2005, *Progress of Theoretical Physics Supplement*, 158, 24
- Marois, C., Zuckerman, B., Konopacky, Q. M., Macintosh, B., Barman, T., 2010, *Nat*, 468, 1080
- Marsh, T. R., et al., 2014, *MNRAS*, 437, 475
- Martinez, P., Klotz, A., 1998, *A practical guide to CCD astronomy*
- Matsumura, S., Thommes, E. W., Chatterjee, S., Rasio, F. A., 2010, *ApJ*, 714, 194
- Mayer, L., Lufkin, G., Quinn, T., Wadsley, J., 2007, *ApJ*, 661, L77
- Mayor, M., Queloz, D., 1995, *Nat*, 378, 355
- Mayor, M., et al., 2003, *The Messenger*, 114, 20
- Mayor, M., et al., 2009, *A&A*, 507, 487
- Mayor, M., et al., 2011, *ArXiv e-prints*

- McCormac, J., Pollacco, D., Skillen, I., Faedi, F., Todd, I., Watson, C. A., 2013, *PASP*, 125, 548
- Miller, N., Fortney, J. J., Jackson, B., 2009, *ApJ*, 702, 1413
- Monet, D. G., et al., 2003, *ApJ*, 125, 984
- Morbidelli, A., Crida, A., Masset, F., Nelson, R. P., 2008, *A&A*, 478, 929
- Morton, T. D., Johnson, J. A., 2011, *ApJ*, 738, 170
- Naef, D., et al., 2007, *A&A*, 470, 721
- Nagasawa, M., Ida, S., Bessho, T., 2008, *ApJ*, 678, 498
- Newton, I., 1726, *Philosophi Naturalis Principia Mathematica*
- Paczynski, B., 1996, *ARA&A*, 34, 419
- Penev, K., Jackson, B., Spada, F., Thom, N., 2012, *ApJ*, 751, 96
- Perryman, M., 2011, *The Exoplanet Handbook*
- Perryman, M. A. C., et al., 2001, *A&A*, 369, 339
- Pickles, A. J., 1998, *PASP*, 110, 863
- Plavchan, P., Bilinski, C., 2013, *ApJ*, 769, 86
- Pollacco, D., Queloz, D., West, R., Wheatley, P., 2011
- Pollacco, D. L., et al., 2006, *PASP*, 118, 1407
- Pont, F., 2009, *MNRAS*, 396, 1789
- Pont, F., Zucker, S., Queloz, D., 2006, *MNRAS*, 373, 231
- Pont, F., Husn00, N., Mazeh, T., Fabrycky, D., 2011, *MNRAS*, 414, 1278
- Pont, F., Sing, D. K., Gibson, N. P., Aigrain, S., Henry, G., Husn00, N., 2013, *MNRAS*, 432, 2917
- Press, W. H., 2007, *Numerical recipes Third Edition*, Cambridge University Press
- Protopapas, P., Jimenez, R., Alcock, C., 2005, *MNRAS*, 362, 460
- Queloz, D., Eggenberger, A., Mayor, M., Perrier, C., Beuzit, J. L., Naef, D., Sivan, J. P., Udry, S., 2000, *A&A*, 359, L13

- Queloz, D., et al., 2009, A&A, 506, 303
- Rice, W. K. M., Veljanoski, J., Collier Cameron, A., 2012, MNRAS, 425, 2567
- Robin, A. C., Reyl  , C., Derri  re, S., Picaud, S., 2003, A&A, 409, 523
- Santerne, A., et al., 2012, A&A, 545, A76
- Santos, N. C., Israelian, G., Mayor, M., Rebolo, R., Udry, S., 2003, A&A, 398, 363
- Seager, S., 2011, Exoplanets
- Seager, S., Richardson, L. J., Hansen, B. M. S., Menou, K., Cho, J. Y.-K., Deming, D., 2005, ApJ, 632, 1122
- Seager, S., Kuchner, M., Hier-Majumder, C. A., Militzer, B., 2007, ApJ, 669, 1279
- Skrutskie, M. F., et al., 2006, ApJ, 131, 1163
- Smith, A. M. S., et al., 2006, MNRAS, 373, 1151
- Socrates, A., Katz, B., Dong, S., Tremaine, S., 2012, ApJ, 750, 106
- Sumi, T., et al., 2011, Nat, 473, 349
- Tamuz, O., Mazeh, T., Zucker, S., 2005, MNRAS, 356, 1466
- Tanaka, H., Takeuchi, T., Ward, W. R., 2002, ApJ, 565, 1257
- Thommes, E. W., Duncan, M. J., Levison, H. F., 2002, ApJ, 123, 2862
- Tinetti, G., et al., 2012, Experimental Astronomy, 34, 311
- Torres, G., Konacki, M., Sasselov, D. D., Jha, S., 2004, ApJ, 614, 979
- Torres, G., Konacki, M., Sasselov, D. D., Jha, S., 2005, ApJ, 619, 558
- Triaud, A. H. M. J., et al., 2010, A&A, 524, A25
- Trilling, D. E., Benz, W., Guillot, T., Lunine, J. I., Hubbard, W. B., Burrows, A., 1998, ApJ, 500, 428
- Welsh, W. F., Orosz, J. A., Seager, S., Fortney, J. J., Jenkins, J., Rowe, J. F., Koch, D., Borucki, W. J., 2010, ApJ, 713, L145
- Winn, J. N., 2010, ArXiv e-prints
- Wolfgang, A., Laughlin, G., 2012, ApJ, 750, 148

- Wolszczan, A., Frail, D. A., 1992, *Nat*, 355, 145
- Wolszczan, A., Hoffman, I. M., Konacki, M., Anderson, S. B., Xilouris, K. M., 2000, *ApJ*, 540, L41
- Wright, J. T., Upadhyay, S., Marcy, G. W., Fischer, D. A., Ford, E. B., Johnson, J. A., 2009, *ApJ*, 693, 1084
- Wright, J. T., Marcy, G. W., Howard, A. W., Johnson, J. A., Morton, T. D., Fischer, D. A., 2012, *ApJ*, 753, 160
- Wright, J. T., et al., 2011, *PASP*, 123, 412
- Wu, Y., Lithwick, Y., 2011, *ApJ*, 735, 109
- Wu, Y., Murray, N. W., Ramsahai, J. M., 2007, *ApJ*, 670, 820
- Zacharias, N., et al., 2010, *ApJ*, 139, 2184