

Observations of exoplanet atmospheres

by

John William Rostron

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Department of Physics

March 2015

THE UNIVERSITY OF
WARWICK

Contents

List of Tables	vi
List of Figures	viii
Acknowledgments	xii
Declarations	xiii
Abstract	xiv
Chapter 1 Introduction	1
1.1 Initial exoplanet discoveries	1
1.2 Exoplanet detection methods	3
1.2.1 The radial velocity technique	4
1.2.2 The transit method	6
1.2.2.1 The WASP survey	8
1.2.2.2 Other results from transit surveys	10
1.2.3 Other detection methods	11
1.3 Atmospheric characterisation	13
1.3.1 Atmospheric detection methods	14
1.3.1.1 Secondary eclipses	14
1.3.1.2 Phase curves	16
1.3.1.3 Transmission spectroscopy	16
1.3.1.4 Direct imaging	18
1.3.2 Results from atmospheric studies	18
1.3.2.1 First atmosphere detections	18
1.3.2.2 Temperatures, albedos and heat redistribution	21
1.3.2.3 Vertical temperature structures	24
1.3.2.4 Compositions	28

1.4	My work	31
Chapter 2 Instrumentation and Methods		32
2.1	The <i>Spitzer Space Telescope</i> and the Infra-Red Array Camera (IRAC)	32
2.1.1	Telescope description	32
2.1.2	IRAC description	34
2.1.3	The Basic Calibrated Data (BCD) pipeline	38
2.1.4	IRAC systematic features	39
2.2	ULTRACAM on the William Herschel Telescope	43
2.2.1	The William Herschel Telescope (WHT)	43
2.2.2	ULTRACAM description	43
2.3	Aperture Photometry	46
2.3.1	The ULTRACAM pipeline	47
2.3.2	Aperture photometry with IDL	50
2.4	System modelling for <i>Spitzer</i> IRAC data	51
2.4.1	The system model	51
2.4.2	Parameter and uncertainty estimation	57
2.4.2.1	Figure of merit for parameter optimisation	57
2.4.2.2	Markov Chain Monte Carlo	58
2.4.2.3	Optimal parameter and uncertainty estimation	59
2.5	The Bayesian Information Criterion (BIC)	61
2.6	Temperature estimation	62
2.6.1	Brightness temperatures	62
2.6.2	Effective temperatures, albedos and heat redistribution	63
Chapter 3 The thermal emission of the exoplanet WASP-3b		65
3.1	Introduction	65
3.2	Observations	66
3.3	Analysis	66
3.3.1	Basic Calibrated Data	66
3.3.1.1	Channel 1 (3.6 μm) Sub-array Warm Mission Data	67
3.3.1.2	Channel 2 (4.5 μm) Cryogenic Mission Data	69
3.3.1.3	Channel 4 (8.0 μm) Cryogenic Mission Data	71
3.3.2	Aperture Photometry	72
3.3.2.1	Source Aperture	72
3.3.2.2	Estimating the Background	73
3.3.2.3	Photometric errors	75
3.3.2.4	Centroiding	75

3.3.2.5	Frame rejection	75
3.3.3	IRAC systematics	76
3.3.3.1	Channels 1 and 2	76
3.3.3.2	Channel 4	77
3.3.4	Markov Chain Monte Carlo Parameter Fitting	78
3.3.5	Light travel time corrections	80
3.3.6	Detrending model selection	81
3.3.7	Source aperture size	85
3.3.8	Centroiding method	85
3.3.9	Channel 1 residual systematics	88
3.3.9.1	Detrending function	89
3.3.9.2	Effect of the source aperture radius	89
3.3.9.3	Centroid position	89
3.3.9.4	Centroiding method	89
3.3.9.5	Trend with FWHM	92
3.3.9.6	Stellar activity	92
3.3.10	Final analysis and prayer bead errors	93
3.4	Results	95
3.5	Discussion	99
3.5.1	Albedo and energy redistribution	99
3.5.2	Atmospheric modelling	99
3.5.3	Stellar activity correlation	103
3.6	Conclusion	105

Chapter 4 Testing the effects of metallicity on the temperature structure of exoplanet atmospheres 107

4.1	Introduction	107
4.2	Observations	110
4.3	Analysis of WASP-21	111
4.3.1	WASP-21 pre-aperture photometry	112
4.3.2	WASP-21 aperture photometry	115
4.3.3	WASP-21 MCMC set up	118
4.3.4	WASP-21 detrending functions tests	120
4.3.5	WASP-21 aperture radius tests	122
4.3.6	WASP-21 additional tests	123
4.3.7	WASP-21 optimal MCMC analysis	127
4.4	Analysis of WASP-28	128

4.4.1	WASP-28 pre-aperture photometry	128
4.4.2	WASP-28 aperture photometry	128
4.4.3	WASP-28 MCMC set up	130
4.4.4	WASP-28 detrending functions tests	131
4.4.5	WASP-28 aperture radius tests	134
4.4.6	WASP-28 optimal MCMC analysis	135
4.5	Analysis of WASP-37	136
4.5.1	WASP-37 pre-aperture photometry	136
4.5.2	WASP-37 aperture photometry	136
4.5.3	WASP-37 MCMC set up	137
4.5.4	WASP-37 detrending functions tests	139
4.5.5	WASP-37 aperture radius tests	141
4.5.6	WASP-37 optimal MCMC analysis	142
4.6	Results	143
4.7	Discussion	149
4.8	Conclusions	153

Chapter 5 Ground-based secondary eclipse observations of WASP-

33		154
5.1	Introduction	154
5.2	Observations	156
5.3	Analysis	158
5.3.1	ULTRACAM Pipeline processing	158
5.3.2	Post-aperture photometry processing	162
5.3.3	Reduced light curves	163
5.3.4	Testing systematic effects	166
5.3.5	The light curve model	173
5.3.6	Light curve fitting	174
5.3.7	Eclipse depth uncertainty estimates	178
5.4	Results	181
5.5	Discussion	182
5.5.1	Comparison to previous measurements	182
5.5.2	Temperature estimates	184
5.5.3	Thermal emission or reflected light?	187
5.6	Conclusions	189

Chapter 6	Conclusions and future work	190
6.1	Chapter summaries	190
6.1.1	The thermal emission of the exoplanet WASP-3b	190
6.1.2	Testing the effects of metallicity on the temperature structure of exoplanet atmospheres	191
6.1.3	Ground-based secondary eclipse observations of WASP-33 . .	191
6.2	Future work	192
Appendix A	Appendix	194
A.1	WASP-33 airmass trend plots	194

List of Tables

2.1	Definitions of the proposal parameters used in MCMCTRANSIT	52
2.2	Definitions of derived parameters used in MCMCTRANSIT	53
2.3	Qualitative summary of evidence provided by Δ BIC values	62
3.1	Comparison of IRAC background variances	74
3.2	Number of rejected frames for WASP-3 <i>Spitzer</i> IRAC observations	76
3.3	Eclipse depth and BIC values for different detrending functions applied to WASP-3's IRAC channel 1 data	82
3.4	Eclipse depth and BIC values for different detrending functions applied to WASP-3's IRAC channel 2 data	83
3.5	Eclipse depth and BIC values for different detrending functions applied to WASP-3's IRAC channel 4 data	83
3.6	Eclipse depth and BIC values for different detrending functions applied to simultaneous fits to WASP-3's IRAC channel 2 and 4 data	84
3.7	Eclipse depths and RMS values for different centroiding methods	88
3.8	Results table for the MCMC analysis of WASP-3 (fitted parameters)	97
3.9	Results table for the MCMC analysis of WASP-3b (derived parameters)	98
4.1	Observation details for WASP-21, WASP-28 and WASP-37	111
4.2	Eclipse depths and BIC values for different detrending functions fitted to WASP-21 IRAC data	121
4.3	Eclipse depths and BIC values for different detrending functions fitted to WASP-28 IRAC data	133
4.4	Eclipse depths and BIC values for different detrending functions fitted to WASP-37 IRAC data	140
4.5	Results table for the MCMC analyses of WASP-21, WASP-28 and WASP-37	143
5.1	Frequencies used in the modelling of WASP-33's pulsations	177

5.2	MCMC and prayer bead uncertainty estimates for the eclipses of WASP-33	180
5.3	Eclipse depth results for WASP-33	181

List of Figures

1.1	Radial velocity of 51 Peg	2
1.2	HD 209458 transit light curve	2
1.3	Summary of exoplanet detections	3
1.4	Radial velocity of 70 Vir	5
1.5	Schematic of transiting planet	7
1.6	The SuperWASP instrument	8
1.7	WASP light curve of WASP-17	9
1.8	Images of HR 8799	12
1.9	<i>Spitzer</i> IRS secondary eclipse of HD 189733b	15
1.10	Sodium detection in HD 209458b	19
1.11	Transmission spectra models for HD 209458b	19
1.12	<i>Spitzer</i> MIPS secondary eclipse detection of HD 209458b	20
1.13	<i>Spitzer</i> 4.5 μm phase curve of HD 189733b	22
1.14	Exoplanet effective temperatures vs irradiation levels	24
1.15	HD 209458b spectrum and model fits	25
1.16	XO-1b spectrum and model fits	26
1.17	Host star activity against planetary 3.6 – 4.5 μm spectral slope	28
1.18	Optical transmission spectrum of HD 209458b	29
1.19	Optical transmission spectrum of HD 189733b	30
2.1	<i>Spitzer's</i> Cryogenic Telescope Assembly	33
2.2	Optical layout of <i>Spitzer</i> IRAC	35
2.3	<i>Spitzer</i> IRAC transmission curves	36
2.4	Schematic describing Fowler sampling	37
2.5	<i>Spitzer</i> IRAC secondary eclipse light curves of TrES-4b	41
2.6	ULTRACAM ray trace diagram	44
2.7	ULTRACAM response curves	45
2.8	<i>Spitzer</i> IRAC channel 2 image of WASP-3	46

2.9	ULTRACAM red channel image of WASP-33	47
2.10	Schematic of the geometry used to model the visible fraction of a planet	55
2.11	MCMC jump parameters for eccentricity constraints	56
3.1	<i>Spitzer</i> IRAC channel 1 image of WASP-3	67
3.2	<i>Spitzer</i> IRAC channel 2 image of WASP-3	68
3.3	<i>Spitzer</i> IRAC channel 4 image of WASP-3	69
3.4	Autocorrelation plot for WASP-3's flux and x and y positions in <i>Spitzer</i> IRAC channel 1	70
3.5	Background trends in <i>Spitzer</i> IRAC channel 1 sub-array frames for WASP-3	71
3.6	Column pull-down in <i>Spitzer</i> IRAC channel 2	72
3.7	Row pull-up in <i>Spitzer</i> IRAC channel 4	73
3.8	Comparison of background estimation methods for WASP-3	74
3.9	Flux trends in x , y and time for WASP-3 in IRAC channel 1	77
3.10	Flux trends in x , y and time for WASP-3 in IRAC channel 2	78
3.11	Signal-to-noise, residual RMS and eclipse depths of WASP-3 against source aperture radius	86
3.12	Raw flux, eclipse and residual light curves of WASP-3	87
3.13	Red noise plot for WASP-3 for different source aperture sizes	90
3.14	Median filtered centroid positions of WASP-3 in IRAC channel 1	91
3.15	Red noise plot for WASP-3 for different centroiding methods	91
3.16	Light curve residuals against FWHM for WASP-3 in IRAC channel 1	92
3.17	Prayer bead errors for WASP-3	94
3.18	2D PDF for Bond albedo and ϵ values for WASP-3	100
3.19	WASP-3 spectrum and model fits	101
3.20	Atmospheric temperature-pressure profiles and contribution functions for WASP-3	103
3.21	Planetary 3.6 – 4.5 μm spectral slope against host star activity (in- cluding WASP-3)	104
4.1	3.6–4.5 μm spectral slope against host star metallicity	109
4.2	Background trends in <i>Spitzer</i> IRAC channel 1 sub-array frames for WASP-21	113
4.3	Background trends in <i>Spitzer</i> IRAC channel 2 sub-array frames for WASP-21	114
4.4	Flux trends in x , y and time for WASP-21 in IRAC channel 1	116
4.5	Flux trends in x , y and time for WASP-21 in IRAC channel 2	117

4.6	Residual RMS values and eclipse depths of WASP-21 against source aperture radius	122
4.7	Red noise plots for WASP-21 for different source aperture sizes . . .	124
4.8	Comparison of light curves of WASP-21 in IRAC channel 1 using source apertures of 2.5 and 5 pixels	125
4.9	Difference image for WASP-21 in IRAC channel 1	126
4.10	Difference light curve for the background flux in IRAC channel 1 for WASP-21 assessed in apertures of 2.5 and 5.0 pixels	127
4.11	Flux trends in x , y and time for WASP-28 in IRAC channel 1	129
4.12	Flux trends in x , y and time for WASP-28 in IRAC channel 2	130
4.13	Residual RMS values and eclipse depths of WASP-28 against source aperture radius	134
4.14	Red noise plots for WASP-28 for different source aperture sizes . . .	135
4.15	Flux trends in x , y and time for WASP-37 in IRAC channel 1	137
4.16	Flux trends in x , y and time for WASP-37 in IRAC channel 2	138
4.17	Residual RMS values and eclipse depths of WASP-37 against source aperture radius	141
4.18	Red noise plots for WASP-37 for different source aperture sizes . . .	142
4.19	Raw flux, eclipse and residual light curves of WASP-21	144
4.20	Raw flux, eclipse and residual light curves of WASP-28	145
4.21	Raw flux, eclipse and residual light curves of WASP-37	146
4.22	MCMC chain values for e against ω for WASP-21, WASP-28 and WASP-37	148
4.23	2D PDF for Bond albedo and ε values for WASP-28	150
4.24	Planetary 3.6 – 4.5 μm spectral slope against host star metallicity (including values for WASP-21, WASP-28 and WASP-37)	151
4.25	Planetary 3.6 – 4.5 μm spectral slope against host star activity (including WASP-21, WASP-28 and WASP-37)	152
5.1	Secondary eclipse of WASP-33b in the SIII band	156
5.2	ULTRACAM response curves used in the observations of WASP-33 .	157
5.3	Signal-to-noise ratios for typical ULTRACAM images of WASP-33 .	161
5.4	Raw flux light curves of WASP-33 with ULTRACAM	163
5.5	Sky background light curves for the WASP-33 observations with ULTRACAM	164
5.6	Differential flux light curves of WASP-33 with ULTRACAM	165

5.7	Flux trends for comparison 1 and comparison 2 in airmass and time for the red channel of ULTRACAM	167
5.8	Flux and position trends for comparison 1 and comparison 2 in the red channel of ULTRACAM	168
5.9	Flux and position trends for WASP-33 and comparison 1 in the red channel of ULTRACAM	169
5.10	Flux and position trends for WASP-33 and comparison 2 in the red channel of ULTRACAM	170
5.11	Flux trends for WASP-33 and comparison 2 in airmass and time for the red channel of ULTRACAM	172
5.12	Binned RMS, reduced χ^2 and eclipse depths as a function of the number of modelled frequency modes for WASP-33, in the three channels of ULTRACAM	176
5.13	Red noise plots for WASP-33 for different numbers of modelled frequency modes	179
5.14	Normalised differential flux, eclipse and residual light curves of WASP-33 in the three ULTRACAM channels	183
5.15	WASP-33 spectrum and blackbody model fits	186
5.16	2D PDF for Bond albedo and ε values for WASP-33	187
A.1	Flux trends for comparison 1 and comparison 2 in airmass and time for the green channel of ULTRACAM	195
A.2	Flux trends for comparison 1 and comparison 2 in airmass and time for the blue channel of ULTRACAM	196
A.3	Flux trends for WASP-33 and comparison 2 in airmass and time for the green channel of ULTRACAM	197
A.4	Flux trends for WASP-33 and comparison 2 in airmass and time for the blue channel of ULTRACAM	198

Acknowledgments

First, I would like to thank my supervisor Dr. Peter Wheatley for his excellent support and guidance over the past 4 years. His ideas and insights have been invaluable and our detailed discussions about all aspects of our work (and beyond) are something I will certainly look back on fondly. I thank him for giving me the opportunity to work in such an exciting area of research.

I would also like to thank all of the PhD students in the Astronomy group with whom I have shared my time at Warwick. In particular Joao, Penelope, Phil, Rachel, Sandra, Simon and Steve for helping me settle into the group.

I thank Mum, Dad, Dave and Rob for their wonderful support throughout my time at Warwick (and before!). Finally, I would like to thank my girlfriend Rosie for putting up with me and for her love, encouragement and silliness.

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 2010 to September 2014, under the supervision of Dr. Peter Wheatley. 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.

The work in Chapter 3 is based on a paper accepted by the *Monthly Notices of the Royal Astronomical Society (MNRAS)* entitled ‘The thermal emission of the exoplanet WASP-3b’ [Rostron et al., 2014].

The *Spitzer* observations presented in Chapter 4 were part of a *Spitzer* program entitled ‘Testing the effects of metallicity on the atmospheric chemistry of exoplanets’ (*Spitzer* program ID 80164, PI Peter Wheatley).

The observations presented in Chapter 5 were from an observing program entitled ‘The thermal emission of the hottest known exoplanet using ULTRACAM/WHT’ (WHT program ID P26, PI Joao Bento). These observations were carried out by Chris Copperwheat and Madelon Bours.

Other contributions based on this thesis are:

- National Astronomy Meeting, Manchester, UK, March 2012, Oral presentation: *Detection of thermal emission from WASP-3b*.
- International Astronomical Union Symposium 299, Victoria, Canada, June 2013, Poster presentation: *Detection of thermal emission from WASP-3b*.
- European Planetary Science Congress, London, UK, September 2013, Oral presentation: *Testing the effects of metallicity on exoplanets*.

Abstract

In the past decade the field of extra-solar planetary science has moved beyond simple detections of planets outside of our Solar System into more detailed characterisations of these objects. One avenue that is at the forefront of current research is the study of extra-solar planet atmospheres. This work has focused on a subset of the current exoplanet population: transiting hot Jupiters. The large sizes, high temperatures and bright stellar hosts of these planets make them particularly amenable to atmospheric studies through techniques such as secondary eclipse observations, which can sample the planetary thermal and reflected light.

Atmospheric detections have now been made for over 50 extra-solar planets. With such a population, we can begin to look for trends in the atmospheric properties of these planets, in order to shed more light on the physical processes that affect their atmospheres. It is in this context that I present the work in this thesis, which comprises secondary eclipse observations of five transiting hot Jupiter exoplanets.

Secondary eclipses of WASP-3b were observed using the *Spitzer Space Telescope*, giving estimates of the planets thermally emitted flux at 3.6 μm , 4.5 μm and 8.0 μm . These estimates imply the planet is very hot ($T_{\text{eff}} = 2280_{-150}^{+200}$ K) and that it may host an inverted vertical temperature profile. This system probes a cut-off in a proposed correlation between the vertical temperature structures of hot Jupiters and the chromospheric activity of their host stars. I find that my measurements for WASP-3b imply this cut-off is more complex than initial data has suggested.

Secondary eclipses of the planets WASP-21b, WASP-28b and WASP-37b, all with low metallicity host stars, were also observed with the *Spitzer Space Telescope*, at 3.6 μm and 4.5 μm . These systems were studied to explore potential correlations between planetary spectral properties and host star metallicities. In existing data, a hint of a trend between the vertical temperature structures of planets and the metallicities of their host stars is found. However this trend is not supported by the planetary flux estimates derived for the three systems I have studied.

Ground-based optical and near infra-red secondary eclipses were also observed for WASP-33b, using the ULTRACAM instrument on the William Herschel Telescope. The emission of WASP-33b, detected in the z' band, confirms that the planet is extremely hot ($T_z = 3170_{-190}^{+90}$ K) and supports a trend found for highly irradiated planets to have systematically low albedos and poor heat redistribution properties.

Chapter 1

Introduction

1.1 Initial exoplanet discoveries

Extra-solar planets, or ‘exoplanets’ are planets that have been found outside of our Solar System. The first discovery of an exoplanetary system was by Wolszczan & Frail [1992] who detected a pair of planets around the millisecond pulsar PSR1257+12. The planets were found by monitoring quasi-periodic changes to the pulsar period as it orbited the system’s barycentre. Modelling of these variations revealed two planets with masses $2.8 M_{\oplus}$ and $3.4 M_{\oplus}$ orbiting with periods of 98.2 d and 66.6 d, respectively.

The first discovery of a planet around a main-sequence star came in 1995, with the detection of a Jupiter-mass planet in a very tight 4.2 d orbit around the star 51 Pegasi [Mayor & Queloz, 1995]. The discovery of such a large planet so close to its host star ($a = 0.05$ au) was not expected and motivated new work into the formation and evolution of planetary systems. The planet ‘51 Pegasi b’ was found using the radial velocity method (see Section 1.2 and Figure 1.1) which has been used in many subsequent exoplanet discoveries [e.g. Butler & Marcy, 1996; Marcy & Butler, 1996; Butler et al., 1997; Cochran et al., 1997; Noyes et al., 1997].

The next major advancement in the field came in 2000 with the detection of the first transiting exoplanet, HD 209458b [Henry et al., 2000; Charbonneau et al., 2000]. Photometric monitoring of the light from the system around the predicted transit time (determined from radial velocity observations) revealed a 1.6% dip in light as the planet passed in front of the star (see Figure 1.2). By analysing the system’s transit and radial velocity data simultaneously both Henry et al. [2000] and Charbonneau et al. [2000] were able to obtain mass and radius estimates for the planet and hence constrain its bulk composition. Both studies concluded that




Figure 1.1: The radial velocity signature of the star 51 Pegasi measured with the ELODIE echelle spectrograph [Mayor & Queloz, 1995]. From the sinusoidal variations the presence of the planet, 51 Pegasi b, was inferred with a minimum mass of $M_p \sin i = 0.47 M_J$.




Figure 1.2: Light curve of HD 209458 by Charbonneau et al. [2000], normalised to the out-of-transit flux level. The transit of HD 209458b is clearly seen, with a depth of 1.6%.

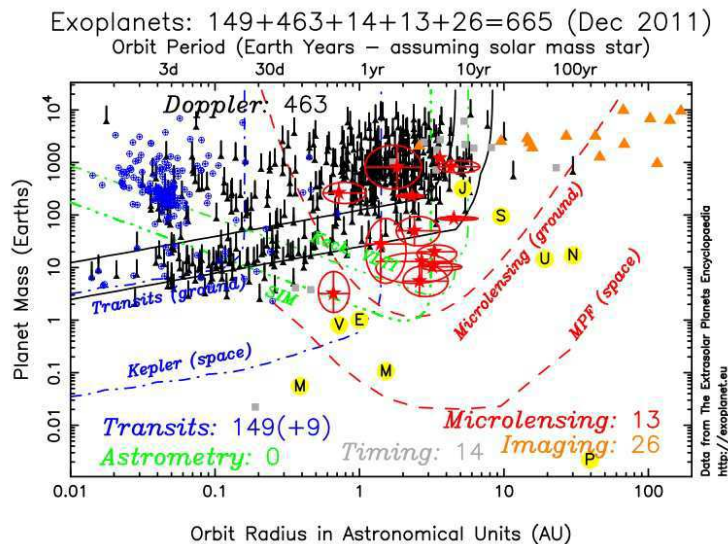


Figure 1.3: Plot of planets discovered (as of 2011) in planet mass - orbital separation space. The radial velocity and transit techniques are the dominant methods and are typically sensitive to higher mass planets at small separations, though the radial velocity technique does find many planets beyond 1 au. Other methods e.g. imaging and microlensing (Section 1.2.3) are sensitive to other regions of this parameter space. Figure courtesy of Keith Horne.

HD 209458b is a gas giant planet, with a mass of $0.6 M_J$ and a radius of $1.3\text{--}1.4 R_J$

The discovery of transiting exoplanets also opened up opportunities for detailed studies beyond estimates of bulk composition. In this thesis I focus on one of these avenues: constraining exoplanetary atmospheric properties through secondary eclipse observations.

1.2 Exoplanet detection methods

Over 1700¹ exoplanets have been discovered so far, using a variety of techniques. Figure 1.3 shows how these different techniques are sensitive to different parts of the planet mass–orbital separation parameter space. For example, direct imaging tends to find very massive, young planets ($M_p > 4 M_J$) at large separations, while the radial velocity and transit methods are sensitive to planets in tighter orbits and with a range of masses. The latter two methods have been by far the most successful, accounting for around 95% of discoveries. These methods are discussed

¹<http://exoplanetarchive.ipac.caltech.edu>, accessed 22/08/14

in the following sections.

1.2.1 The radial velocity technique

In the radial velocity technique the presence of a planet (or planets) orbiting a star is inferred from variations in the star's Doppler shift as it orbits the system's barycentre. The corresponding variation in the radial velocity of the star is then given by:

$$V_{\text{rad}} \simeq c \frac{\Delta\lambda}{\lambda} \quad (1.1)$$

A Jupiter-analogue will produce a reflex motion of 2.4 ms^{-1} , while for an Earth-analogue this value will be 0.09 ms^{-1} , so precise radial velocity measurements are needed. This is done by measuring high resolution spectra of the star and cross-correlating them with a template spectrum. The resulting cross-correlation function combines the information from all of the spectral lines present, from which Doppler shifts, and therefore radial velocities, can be measured very accurately [e.g. $< 1 \text{ ms}^{-1}$ with the HARPS instrument; Mayor et al., 2003].

By monitoring and modelling the variations in the radial velocity signal (see Figure 1.1) the presence of a planet-mass object can be inferred. For a two-body system in a circular orbit the radial velocities are modelled as

$$V_{\text{rad}} = -K_1 \sin(2\pi\phi), \quad (1.2)$$

where ϕ is the orbital phase (see equation 2.22 for the equation accounting for eccentricity). An important quantity here is the radial-velocity semi-amplitude (K_1), which is related to the masses of the objects through

$$K_1 = \left(\frac{2\pi G}{P}\right)^{1/3} \frac{M_p \sin i}{(M_\star + M_p)^{2/3} \sqrt{1-e^2}}. \quad (1.3)$$

Here M_p and M_\star are the planetary and stellar masses, P is the orbital period and i is the inclination of the system - the angle between the orbital axis of the planet and our line of sight. e is the eccentricity of the orbit and is determined by the shape of the radial velocity curve (equation 2.22). Figures 1.1 and 1.4 show examples of both circular and eccentric systems. The effect of a non-zero eccentricity is obvious for the planet around 70 Virginis. Using the above equation, along with a value for M_\star (estimated from the spectral type of the star), the quantity $M_p \sin i$ can be estimated from K_1 (assuming $\frac{M_p}{M_\star} \ll 1$, where $\frac{M_1}{M_\odot} = 0.001$).

$M_p \sin i$ represents a minimum mass estimate for the planet. The degeneracy




Figure 1.4: The radial velocity signature of 70 Virginis, showing the highly eccentric reflex motion of the star around the system’s barycentre. A planet with minimum mass $M_p \sin i = 6.6 M_J$ and eccentricity $e = 0.4$ was inferred from these data [Marcy & Butler, 1996].

with i cannot be resolved with simple radial velocity measurements of the star alone. The radial velocity signature of a system with an edge-on orbit ($i = 90^\circ$) is equivalent to that of a system with an inclined orbit and a larger planet mass (all other quantities being equal). While a definitive mass measurement cannot be made, the probability of a signal being due to a planet can be assessed [Mayor & Queloz, 1995]. For example, if $M_p \sin i = M_J$ then for the true planet mass to be greater than $10 M_J$, i would need to be less than 5.8° . For random orbit orientations this corresponds to a probability of 0.5% that $M_p > 10 M_J$.

Equation 1.3 shows that larger radial velocity signals are given by short period planets with high planet-to-star mass ratios. The bias towards shorter periods is compounded because, in a given amount of time, the radial velocity signature can be measured multiple times and the impacts of the long-term changes in stability of the spectrograph will be suppressed. However, the radial velocity technique is probably the most mature method for finding exoplanets, with surveys having run for many years [e.g. Campbell et al., 1988] and now achieving precisions of $< 1 \text{ ms}^{-1}$ [Cosentino et al., 2012]. Therefore, while surveys are more sensitive to lower period, higher mass planets, they do cover a significant mass and orbital separation range, as shown in Figure 1.3.

Result highlights

After the initial surprise of finding a Jupiter-mass planet orbiting so close to its host star [Mayor & Queloz, 1995], several other examples of these ‘hot Jupiters’ were discovered [Butler et al., 1997, 1998; Fischer et al., 1999; Butler et al., 1999]. Other orbital configurations were also found: a longer period ($P = 3$ yr) Jupiter-mass planet [Butler & Marcy, 1996]; planets in eccentric orbits [Marcy & Butler, 1996; Cochran et al., 1997] and a multi-planet system [Butler et al., 1999]. More recent work has pushed to progressively lower mass objects, culminating in the claim of an $M_p \sin i = 1.1 M_\oplus$ planet orbiting the nearby star α Centauri B with a period of 3.24 d [Dumusque et al., 2012].

Due to the large number of planets discovered, statistical analyses of these systems can also be made. For example, giant planets have been found to be more common around metal-rich stars than metal-poor stars [Udry & Santos, 2007; Mayor et al., 2011], lending support to the core-accretion theory of planet formation.

1.2.2 The transit method

The second of the two dominant methods for exoplanet discoveries is the transit method. As shown in Figure 1.5, if a planet’s orbit is inclined at $\sim 90^\circ$ so that we view it edge-on, then the planet will periodically block a fraction of the light we receive from the star, giving a dip in light or ‘transit’.

One of the many advantages of studying transiting planets is the relative wealth of information that can be derived from them. From transit light curves alone (see Figure 1.2) the most useful parameter one can obtain is the planet-to-star radius ratio, which can be estimated from the fractional depth of the transit (δ) as:

$$\frac{R_p}{R_\star} = \sqrt{\delta} \quad (1.4)$$

Transit signals are small - for a Jupiter-sized planet around a Sun-like star the transit depth is 1%, while for an Earth-sized planet the depth is 0.01%. For the light curve in Figure 1.2, the transit depth of the planet HD 209458b is 1.6%, implying a planet-to-star radius ratio of $\frac{R_p}{R_\star} = 0.13$.

Modelling of transit light curves also allows estimates of the planet’s orbital inclination to be made. Therefore, when analysed together with radial velocity data for the system, the degeneracy described in equation 1.3 can be broken. With an estimate for the mass of the star, the scale of the system (the semi-major axis, a) can be set and absolute values for M_p and R_p can be estimated². A combined analysis

²See Section 2.4.1 for a more detailed description of the equations used, in the context of the




Figure 1.5: A schematic diagram from Seager [2011] showing the orbital geometry and light curve for a transiting planet. When a planet passes in front of its host star a dip in light, or transit, occurs. If the planet contributes light to the system there may be variations in flux as different fractions of the planet’s day/night side are presented to us - a phase curve. Additionally the planet can disappear behind the star, causing a smaller dip in light - a secondary eclipse. The latter can occur due to the thermal or reflected light from the planet.

of transit and radial velocity data therefore allows the bulk density of planets to be measured and their classification (e.g. terrestrial or gas giant) determined, providing crucial tests for theoretical work on the interior structures of planets.

Most planetary systems will not be transiting. The probability for a randomly orientated orbit to be transiting is given by

$$P_{\text{trans}} \simeq \frac{R_{\star}}{a}, \quad (1.5)$$

where R_{\star} is the radius of the star and a is the semi-major axis of the planet’s orbit. For a close-in planet at $a = 0.05$ au around a Sun-like star, $P_{\text{trans}} \sim 10\%$, while for a planet at $a = 1$ au, P_{trans} drops to 0.5%. While early searches for transiting planets focused on the known radial velocity planets (e.g. HD 209458b), in the past decade several transit surveys have been set up. In order to detect transiting planets, many stars must be monitored photometrically for long periods of time to a level of at least 1%, depending on the objects targeted. These surveys (e.g. WASP, Kepler) have,

modelling for the subsequent science chapters.




Figure 1.6: One of the two SuperWASP instruments [Pollacco et al., 2006], here shown at the South African Astronomical Observatory.

on the whole, had great success and now account for over 65% of known planets. Highlights from these surveys are described below, starting with the WASP project.

1.2.2.1 The WASP survey

The ‘Wide Angle Search for Planets’ [WASP; Pollacco et al., 2006] is a transiting planet survey that searches for short-period, Jupiter- and Saturn-sized planets around bright stars. The survey utilises two instruments. One is based in the northern hemisphere at the Roque de los Muchachos Observatory on La Palma and the other is based in the southern hemisphere at the South African Astronomical Observatory. Each instrument comprises a set of eight 200 mm $f/1.8$ telephoto lenses (see Figure 1.6) and can observe large fields of 64 square degrees. Each night the instruments observe a set of pre-chosen fields with a cadence of around 7 minutes, in a cyclic manner. In this way each field is observed for 120–150 nights per year [Collier Cameron et al., 2009]. Combined, the two instruments cover a declination range from -90° to $+60^\circ$ and monitor the brightness of over 30 million individual objects with $V < 15$.

For stars with $7 < V < 13$, a photometric accuracy of 1% is achieved on a time-scale of a few hours, which is sensitive enough to detect Jupiter- and Saturn-sized planets around solar-type stars. Indeed the WASP project has found many




Figure 1.7: Light curve of the transiting planet system WASP-17 as taken by the SuperWASP instrument [Anderson et al., 2010]. A transit depth of 1.7% is seen, revealing a planet with $R_p = 2.0 R_J$

such planets. Over 90 have been discovered and published, with many more awaiting publication. These planets tend to be large ($R_p = 0.7\text{--}2.0 R_J$) and in short-period orbits ($P = 0.8\text{--}10$ d) i.e. they are ‘hot Jupiters’. The tendency towards large planets is a direct result of equation 1.4: larger planets produce bigger transit signals. Short-period planets are favoured for a variety of reasons: they are more likely to transit (equation 1.5); their transits typically last a few hours (short enough to observe in a single night) and the transit signal can be sampled over many transit events. WASP candidates are also amenable to radial velocity follow-up for confirmation as planet-mass objects as they are bright systems.

These factors ensure WASP planets are good targets for further follow-up work. In particular, they are well suited to atmospheric studies due to their low densities and high equilibrium temperatures (see Section 1.3.1). These atmospheric follow-up observations of WASP planets are the focus of this thesis.

Highlights from the WASP survey include WASP-17b [see Figure 1.7; Anderson et al., 2010], a bloated ($R_p = 1.9 R_J$), low density planet ($\rho_p = 0.19 \text{gcm}^{-3}$). This planet was the first to have been found to orbit its star in a retrograde fashion (using Rossiter-McLaughlin measurements) i.e. in the opposite direction to the rotation of the star, suggesting planet-star or planet-planet scattering processes may have been involved in the orbital evolution of the planet.

WASP-19b is a Jupiter-like planet in a very tight orbit around its host [Hebb et al., 2010]. At the time of discovery it was the shortest period planet known, with $P = 0.79$ d. Brown et al. [2011] studied the tidal interactions of the system and

estimated the remaining lifetime of the planet (before infall onto the star) to be $0.0067^{+1.1073}_{-0.0061}$ Gyr, raising the possibility that the planet could be in the final stages of its life.

The planet WASP-33b was the first (and is currently one of only two) transiting planet to have been found orbiting an A-type star [Collier Cameron et al., 2010]. Orbiting its host with a 1.2 d period it is also the most highly irradiated planet currently known. Herrero et al. [2011] subsequently found the star to be a delta-Scuti type variable - a unique feature in the current exoplanet population. This planet is the focus of Chapter 5 of this thesis.

1.2.2.2 Other results from transit surveys

Several other ground-based surveys have also been successful in finding hot Jupiter planets around bright stars e.g. the Trans-Atlantic Exoplanet Survey [Alonso et al., 2004], the HATNet project [Bakos et al., 2004] and the XO project [McCullough et al., 2005]. One of the most intriguing results from this population of well characterised planets is the discrepancies between the observed radii and those predicted from structural evolutionary models. They tend to be over-sized, implying an excess in thermal support for these planets. Explanations for this trend have focussed on the observed correlation between the radius residuals and the irradiation the planet experiences [Laughlin et al., 2011], with mechanisms such as Ohmic or kinetic heating being proposed.

Along with the ground-based surveys, there have also been space-based transiting surveys. The two most successful of these have been the ‘COncvection ROtation and planetary Transits’ (CoRoT) and Kepler missions. CoRoT [Auvergne et al., 2009] utilised a 27 cm telescope and observed two $2.7 \times 3.0^\circ$ fields with a photometric precision of better than 0.1% (for $R = 15$ over a 8.5 min exposure). To date the survey has discovered 27 planets³. A highlight amongst these has been CoRoT-7b, the first example of a transiting super-Earth [$R_p = 1.7 R_\oplus$; Léger et al., 2009] and possibly the first rocky exoplanet discovered [there is still debate over this due to differing planetary mass estimates from various studies: Pont et al., 2011; Hatzes et al., 2011; Haywood et al., 2014].

Kepler [Borucki et al., 2010] is a mission designed primarily to characterise the frequency of Earth-sized planets in the habitable zones of Sun-like stars. It utilises a 95 cm optical telescope observing a single 100 square degree field-of-view and provides simultaneous monitoring of the light curves of over 100,000 stars. The

³<http://exoplanetarchive.ipac.caltech.edu>; accessed 27/08/14.

precision of the measurements is 0.01 % for a 30 min exposure of a star with a Kepler magnitude of 12.5.

To date Kepler has discovered 978 planets⁴. Amongst the highlights from this survey are: Kepler-10b [Batalha et al., 2011; Dumusque et al., 2014], the first rocky planet discovered by Kepler; the Kepler-9 system [Holman et al., 2010], the first multi-planet transiting system and Kepler-16b [Doyle et al., 2011], the first circumbinary planet discovered. In 2011 the first transiting planet in the habitable zone of its host star was found [Borucki et al., 2012]. Kepler-22b is a $2.4 R_{\oplus}$ planet orbiting a G5 dwarf star with a period of 290 days and a radiative equilibrium temperature of 262 K (for comparison, the value for Earth is 255 K). However, the mass of the planet is still highly uncertain, so its composition relative to Earth is still unknown. Similar discoveries have subsequently been made e.g. Kepler-186f [Quintana et al., 2014], but again with an ambiguity in terms of their compositions.

1.2.3 Other detection methods

Direct Imaging

Direct imaging involves the detection of the point source image of an exoplanet, either in thermal emission or reflected light (see Fig. 1.8). The limiting factor in these observations is the planet-to-star contrast, which is $\sim 10^{-5}$ (in favourable cases). To overcome this several techniques are used: the use of large aperture telescopes, since resolution goes inversely with the telescope’s aperture size; adaptive optics, used to suppress the effects of atmospheric turbulence, allowing the telescope to approach its diffraction limit; use of coronagraphs, to physically block the host star’s light and ‘angular difference imaging’ [Perryman, 2011].

The direct imaging technique is most sensitive to nearby systems containing bright planets orbiting far from their host stars (see Fig 1.3). As a result, the systems discovered using this technique tend to be young, where newly formed giant planets are large and highly luminous due to on-going gravitational contraction [Baraffe et al., 2010]. Thus direct imaging samples a rather different region of parameter space to the transiting and radial velocity techniques.

Current estimates for the numbers of directly imaged planets vary from 8^5 – 51^6 , due to the uncertain mass estimates of many of the detections. But with new, dedicated instruments such as SPHERE (Spectro-Polarimetric High-contrast Exoplanet REsearch) on the VLT and GPI (Gemini Planet Imager) the future of the

⁴<http://kepler.nasa.gov>; accessed 27/08/14

⁵<http://exoplanets.org>; accessed 29/08/14

⁶<http://exoplanet.eu>; accessed 29/08/14




Figure 1.8: Images of the HR 8799 system taken using the Keck and Gemini telescopes [Marois et al., 2008]. Three planets (b, c and d) are seen, with masses estimated to be around 7, 10 and 10 M_J and (projected) orbital separations of 68, 38 and 24 au, respectively.

direct imaging technique looks promising.

Microensing

Planet detection through microlensing works because of the effects of the distortions of spacetime by the star-planet system. If a foreground star is aligned between the Earth and a background source (e.g. a more distant star in the galaxy) then it can act as a lens, magnifying the light of the background source over time-scales of a few days. If the star hosts a planet, an additional magnification on time-scales of hours can be superimposed onto the light curve. Surveys that photometrically monitor stars in high density regions (e.g. the galactic bulge) have discovered 26 planets in this way⁷.

The strengths of this method are its sensitivity to Earth-mass planets at moderate to large separations [$a > 1$ au; Bennett et al., 2007, also see Figure 1.3]. In addition, microlensing works most effectively for star-planet systems at distances of ~ 4 kpc, meaning the galactic population of planets at moderate separations from their host stars can be sampled. However, these systems are very faint (even invisible), and the transient nature of microlensing events does not make them available for subsequent follow-up studies.

⁷<http://exoplanetarchive.ipac.caltech.edu>, accessed 22/08/14

Astrometry

As a planet hosting star moves around the system's barycentre, not only will it undergo a periodic radial motion from our viewpoint (detected as radial velocity variations) but it will also undergo a periodic transverse motion. For a two-body system this motion will be an ellipse with an angular semi-major axis given by (for a circular orbit):

$$\alpha = \left(\frac{M_p}{M_\star}\right) \left(\frac{a}{1 \text{ au}}\right) \left(\frac{d}{1 \text{ pc}}\right)^{-1} \text{ arcsec}, \quad (1.6)$$

where d is the distance to the system [Perryman, 2011]. This elliptical orbit will be superimposed onto the proper motion signal of the system's barycentre, while the Earth's orbital motion around the Sun also has to be accounted for.

Precise astrometric monitoring of stars therefore gives another planet detection method. This technique is most amenable to the detection of nearby systems containing massive planets orbiting at large separations. No planets have so far been found using this technique, but with the launch of the GAIA mission this should soon change. With a single-measurement precision of $8 \mu\text{as}$, the expectation is that several thousand giant planets out to 3–4 au will be found within 200 pc [Casertano et al., 2008].

1.3 Atmospheric characterisation

The discovery of exoplanets, using the techniques described above, is only the first stage of study into these systems. Follow-up of these discoveries reveals a more detailed picture of exoplanetary systems. For example, bulk density calculations for transiting planets have shown a surprisingly diverse range of densities for gas giant planets, while Rossiter-McLaughlin measurements have shown cases of planets with retrograde orbits - another unexpected result.

Another avenue for follow-up work is the study of exoplanetary atmospheres. These studies can probe the temperatures, pressures and compositions of exoplanet atmospheres, allowing for a deeper understanding of the nature of these planets. Much of the work in this area has been done on tightly-orbiting transiting systems i.e. hot Jupiters, which are well suited to many of the techniques used (see Sections 1.3.1.1–1.3.1.3), with additional work coming from direct imaging (Section 1.3.1.4). While these systems are very different to our Solar System, the hope is that these techniques will one day be used to probe the atmospheres of potentially habitable planets in searches for the signatures of life.

In the following Section I describe the techniques used to probe these atmo-

spheres, while in Section 1.3.2 I review the literature in this field.

1.3.1 Atmospheric detection methods

1.3.1.1 Secondary eclipses

For a transiting planet system with a circular orbit, half an orbit after a transit occurs the planet will disappear behind the star. The light that the planet contributes to the system will be lost for the time it is hidden from view, causing another periodic ‘dip’ in the light curve of the system (see Figure 1.5). This dip is referred to as the secondary eclipse. By measuring the flux from the system outside and during secondary eclipse, and differencing these values, we can separate the light from the day-side of the planet from that of the star. Typically we do not measure the absolute flux, but rather the planet-to-star flux ratio, which is given by the depth of the eclipse normalised to the estimated stellar flux level. An example of a secondary eclipse light curve is shown in Figure 1.9, for the planet HD 189733b.

By measuring secondary eclipse depths at different wavelengths we can probe the planet’s day-side spectrum. This spectrum includes contributions from both reflected and thermal emission, which can be described using [Haswell, 2010]:

$$\frac{F_{p,\lambda}}{F_{\star,\lambda}} = A_{g,\lambda} \left(\frac{R_p}{a} \right)^2 + \frac{B_\lambda(T_{p,\lambda})}{B_\lambda(T_{\star,\lambda})} \left(\frac{R_p}{R_\star} \right)^2 \quad (1.7)$$

Here $F_{p,\lambda}$ is the day-side flux of the planet; $F_{\star,\lambda}$ is the stellar flux; $T_{p,\lambda}$ and $T_{\star,\lambda}$ are the brightness temperatures for the planet and star at wavelength λ and B_λ is the Planck function. $A_{g,\lambda}$ is the geometric albedo, which is defined as the ratio of the planet’s reflected day-side flux to the reflected flux from a Lambert disk (which scatters an equal intensity in all directions) with the same orbital separation and cross-sectional area as the planet. The first term on the right hand side of equation 1.7 describes the reflected light, while the second term describes the planet’s thermal emission.

Secondary eclipse studies (including this thesis) focus mainly on hot Jupiter planets. This can be understood through equation 1.7 because these large, hot planets in tight orbits around their hosts will give the biggest secondary eclipse signals in both reflected and thermal light. For a typical system with a Jupiter-like planet orbiting at 0.05 au from a Sun-like host, and a geometric albedo of $A_{g,\lambda} = 0.17$ [Rowe et al., 2008], the reflected light contribution dominates (i.e. is 10 times larger than the thermal emission) at $\lambda < 890$ nm, while the thermal emission dominates at $\lambda > 1.6 \mu\text{m}$. As a result, optical detections of secondary eclipses are assumed to probe




Figure 1.9: Secondary eclipse light curve of the planet HD 189733b, taken with *Spitzer's* IRS instrument at $\lambda = 16 \mu\text{m}$ [Deming et al., 2006]. The bottom panel here is a binned version of the top panel. The depth of the eclipse, a direct measure of the planet-to-star flux ratio, is $0.551 \pm 0.030 \%$ and implies a planetary brightness temperature of $T_{16\mu\text{m}} = 1120 \text{ K}$. Note that eclipses are deeper at longer wavelengths and with currently available instrumentation, which probe $\lambda < 10 \mu\text{m}$, eclipse depths tend to be $\sim 0.1\text{--}0.2 \%$.

the planet's reflected spectrum, while infra-red measurements probe the planet's thermal spectrum. The work in this thesis mainly focuses on secondary eclipses that probe the thermal spectra of planets (see Chapters 3–5). Typical eclipse depths are $\sim 0.1 \%$ in the infra-red (and increase with λ), while optical eclipses tend to be even smaller at $< 0.01 \%$. These signals are extremely small and require very stable instruments (e.g. *Spitzer*, Kepler). Even with the best current instrumentation, careful removal of systematic noise is crucial in the detection of these eclipses.

Infra-red secondary eclipses probe the day-side spectrum of the planet. From this, constraints can be placed on the abundances of species present in the planet's photosphere, through the identification of molecular features. The spectrum is also sensitive to vertical temperature structures because different wavelengths probe flux originating from different depths in the atmosphere. Features are typically expected

to be seen in absorption, but under certain conditions e.g. an atmospheric temperature that rises with altitude in the photosphere, molecular features can be driven into emission (see Figure 1.15). A caveat here is that the optical depths are set by the opacity sources, which will in turn depend on the temperature structure. Therefore a planet’s day-side spectrum probes a combination of the temperature structure and the composition. These are not easily separated with current eclipse measurements, which are typically made using broad band filters that integrate over multiple molecular features [e.g. Madhusudhan & Seager, 2010].

At optical wavelengths, secondary eclipses constrain the geometric albedo ($A_{g,\lambda}$, equation 1.7), which is useful for understanding the scattering properties of the planet [Evans et al., 2013]. Additionally, $A_{g,\lambda}$ can be used to constrain the Bond albedo (A_B) of the planet, which is the fraction of the incident stellar energy reflected back into space [Rowe et al., 2006]. This sets the energy budget for the planet and provides an important constraint for atmospheric models.

1.3.1.2 Phase curves

As an exoplanet orbits its host star there may be detectable variations in the system flux due to the varying illumination of the planet as seen from Earth (see Figure 1.5). This variation is known as a phase curve and could be a result of changes in the reflected or thermal flux we see from the planet.

The latter can occur due to differences in the temperature of the planet’s day and night-side. Phase curve observations have so far been performed exclusively on hot Jupiters [e.g. Cowan et al., 2012; Maxted et al., 2013], which are expected to be tidally locked to their hosts i.e. the same side of the planet always faces the star. This sets up a longitudinal temperature gradient, which could potentially be suppressed in the case of efficient advection of the irradiating energy around the planet. Alternatively, it could remain in the case of efficient re-radiation of this energy. Measuring the amplitude of the thermal phase curve probes the day-night contrast in thermal emission and can therefore reveal information on the global heat redistribution properties of the planet [Knutson et al., 2012].

Note that, while this method can be applied to non-transiting systems, the vast majority of phase curve observations are made for transiting systems.

1.3.1.3 Transmission spectroscopy

Another method for probing the atmospheres of transiting planets is that of transmission spectroscopy. This method makes use of the fact that a planet is not an

opaque disk with a sharp edge, but rather it has an extended atmosphere that becomes increasingly optically thin with altitude [Haswell, 2010]. The amount of light that can pass (transmit) through the atmosphere will vary with wavelength according to the absorption properties of the atmosphere at that wavelength. This in turn will depend on the species present in the planet’s atmosphere and their abundance. The varying amounts of absorption effectively cause a change to the size of the planet. Observationally a ‘transmission spectrum’ is given by the variation of the transit depth with wavelength (see Figure 1.11) and this is used to probe opacity sources in the planet’s atmosphere.

The strength of these absorption features is expected to scale with the atmospheric scale-height (H) where:

$$H = \frac{kT}{\mu g}, \quad (1.8)$$

where k is the Boltzmann constant, T is the atmospheric temperature, μ is the mean molecular weight of the atmosphere and g is the planet’s surface gravity [Madhusudhan et al., 2014]. The strength of the transmission spectrum can then be estimated by considering the size of the annulus extending from R_p to $R_p + H$ as a fraction of the area of the stellar disk [neglecting the squared term in H ; Brown, 2001]:

$$\delta_{\text{trans}} = \frac{2R_p H}{R_\star^2}, \quad (1.9)$$

so

$$\delta_{\text{trans}} = \frac{2R_p}{R_\star^2} \frac{kT}{\mu g} \quad (1.10)$$

Since $g \propto \frac{M_p}{R_p^2}$, equation 1.10 indicates that hot planets with low densities will produce the largest transmission signals. Therefore, hot Jupiters are ideal planets for transmission spectroscopy studies, but as with the secondary eclipses, signal strengths are very small (typically 0.01–0.1 %).

One of the advantages of transmission spectroscopy is that it probes the atmospheric composition and is not directly sensitive to the atmospheric temperature and pressure. Compositional information can be useful, for example, in constraining energy budgets through estimation of the abundances of principal absorbers. Additionally, transmission spectroscopy can be performed at a range of wavelengths, from the optical, where atmospheric hazes and alkali metals are probed [e.g. Huitson et al., 2012], to the infra-red, where features due to H_2O , CH_4 and CO are expected [e.g. Deming et al., 2013].

1.3.1.4 Direct imaging

Planets discovered using the direct imaging technique are also well suited to atmospheric studies. Detections are typically made at near infra-red wavelengths and often in multiple photometric bands, providing low resolution spectroscopic information. Bright planets that are well separated from their hosts can also be followed-up with spectroscopic observations at both low and moderate resolutions.

For example, the multi-planet system HR 8799 (shown in Figure 1.8) has been studied in this way. The near infra-red spectra of the planets HR 8977b and c show they are consistent with hydrogen-rich atmospheres that are dominated by H₂O absorption, but lacking in CH₄. Barman et al. [2011] suggest that models including significant cloud opacity, super-solar metallicities and non-equilibrium chemistry driven by atmospheric vertical mixing are required to explain the observed spectra. In addition, the observed ‘triangular’ shape of the H-band spectra is believed to be consistent with low surface gravities. H-band spectroscopy of other planets [e.g. Chilcote et al., 2014] all point towards low surface gravities, while deficiencies in CH₄ are also found (though for hotter planets this is expected from equilibrium chemistry).

Atmospheric studies of directly imaged planets are important for probing the regions where giant planets are expected to form and for putting constraints on possible formation mechanisms. For example, Konopacky et al. [2013] highlights that robust determinations of C/O ratios for directly imaged giant planets could distinguish between formation scenarios (i.e. core accretion vs gravitational instabilities). However, the planets studied in this thesis - hot Jupiters - are quite different from these directly imaged planets. As a result they will not be discussed further.

1.3.2 Results from atmospheric studies

In this Section I will review topics relating to atmospheric studies of hot Jupiter exoplanets. The vast majority of these systems are transiting and have all been studied through secondary eclipse, transmission spectroscopy and/or phase curve measurements. In this way I will place my work on secondary eclipse measurements of transiting hot Jupiters (Chapters 3–5) into the context of our current understanding of exoplanet atmospheres.

1.3.2.1 First atmosphere detections

The first detection of an exoplanet atmosphere was made by Charbonneau et al. [2002] who observed sodium in the transmission spectrum of the transiting planet




Figure 1.10: The detection of sodium absorption in the atmosphere of HD 209458b by Charbonneau et al. [2002]. Both panels show the difference in the transit light curves for a band centred on a sodium feature compared to bands adjacent to this feature. The negative values during the transit suggest the sodium band transit is deeper than the adjacent bands and this is interpreted as being due to significant sodium absorption in the planet’s atmosphere. The upper panel shows the raw flux, while the lower panel bins the data in sets of 42 observations.




Figure 1.11: Transmission spectra models from Charbonneau et al. [2002], in the same units as given in Figure 1.10. The sodium feature targeted in the study was at 589 nm, using a narrow 12 Å band. The standard model, using solar metallicity and a cloud deck at 0.04 bar is given by the dotted line. This model over-predicts the depth of the transmission signal seen in Figure 1.10. The solid black line is for an atmosphere with a high cloud layer, while the solid grey line has a depleted sodium abundance. Both of these provide potential explanations for the observation.

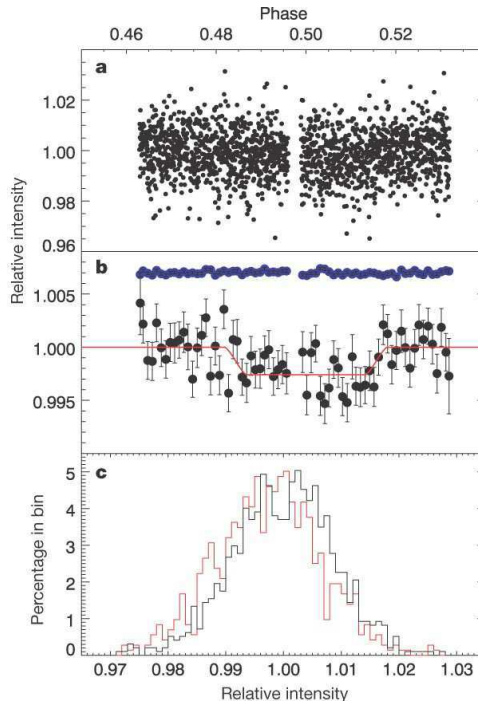


Figure 1.12: Secondary eclipse detection of HD 209458b with *Spitzer* MIPS [Deming et al., 2005]. (a) Raw flux normalised to the out-of-eclipse flux level against the orbital phase for the planet. (b) Raw fluxes binned in phase with a width of 0.001. The red model is the best fitting secondary eclipse model, with a depth of 0.260 % and central eclipse phase of 0.5. (c) Histograms of the out-of-eclipse (black) and in-eclipse fluxes (red) showing the reduction in flux from the system as the planet passes behind the star.

HD 209458b. These observations were made across three separate transits using the Space Telescope Imaging Spectrograph (STIS) onboard the Hubble Space Telescope (HST). The transit depth of the planet was measured to be $0.0232 \pm 0.0057\%$ deeper in a 12 \AA band centred on sodium than it was in adjacent bands (Figure 1.10). This measurement confirmed the prediction by Seager & Sasselov [2000] that a strong sodium doublet at 589 nm would be detectable in the atmospheres of hot Jupiter planets. Charbonneau et al. [2002] compared the detected signal to various models [Brown, 2001, see Figure 1.11] and found that it was smaller than the prediction from a standard model with solar metallicity and a cloud deck at 0.04 bar. Several causes for this were suggested, including a primordially depleted metal abundance or the presence of a high cloud deck.

The next milestone came in 2005 with the first detections of thermal emission

from exoplanets, through secondary eclipse observations. Two independent studies were made that represented the first ever measurement of light from planets outside of the solar system.

Deming et al. [2005] measured the thermal emission from the planet HD 209458b in a photometric band centred on $24\ \mu\text{m}$ ($\Delta\lambda = 8\ \mu\text{m}$) using the MIPS instrument onboard the *Spitzer Space Telescope*. The eclipse (shown in Figure 1.12) had a depth of $0.260 \pm 0.046\%$ corresponding to a planetary brightness temperature of $T_{24\mu\text{m}} = 1130 \pm 50\ \text{K}$. Charbonneau et al. [2005] then detected a secondary eclipse of the planet TrES-1b, using *Spitzer's* Infra-Red Array Camera (IRAC) instrument. Simultaneous observations of the secondary eclipse were made at $4.5\ \mu\text{m}$ ($\Delta\lambda = 1.1\ \mu\text{m}$) and $8.0\ \mu\text{m}$ ($\Delta\lambda = 3\ \mu\text{m}$), with eclipse depths of $0.066 \pm 0.013\%$ and $0.225 \pm 0.036\%$ being found. The derived brightness temperatures here were $T_{4.5\mu\text{m}} = 1010 \pm 60\ \text{K}$ and $T_{8.0\mu\text{m}} = 1230 \pm 110\ \text{K}$, suggesting marginal inconsistency with a blackbody spectrum.

These initial detections laid the ground work for subsequent observations of exoplanet atmospheres and the simultaneous development of theoretical atmosphere modelling. The current research themes within the field of exoplanetary atmospheres are explored in the following Sections.

1.3.2.2 Temperatures, albedos and heat redistribution

One of the most basic conclusions from early secondary eclipse detections, as noted by Seager & Deming [2010], was the confirmation that ‘hot’ Jupiters were indeed hot. Predictions of temperatures for these planets exceeded $1000\ \text{K}$, in agreement with the brightness temperatures derived from infra-red secondary eclipse observations [e.g. see Cowan & Agol, 2011, Table 1]. Early theoretical models also highlighted that the dominant energy source for hot Jupiters would be the irradiating flux (rather than energy released through gravitational contraction) and that this would significantly alter the temperature structure and emergent spectra compared with isolated planets [Seager & Sasselov, 1998].

These early theoretical studies also suggested that silicate clouds on the most highly irradiated planets could increase albedos significantly [e.g. Sudarsky et al., 2000]. However, this is not supported by the significant amounts of infra-red flux detected in secondary eclipse observations. Stronger evidence comes from direct measurements of the geometric albedo, through optical secondary eclipse and phase curve detections. For example, Rowe et al. [2008] found that for HD 209458b $A_{g,MOST} < 0.08$ at 1σ (where the MOST bandpass is $400 - 700\ \text{nm}$). The authors argue that $A_B < 0.12$, ruling out the possibility of reflective clouds. Other examples




Figure 1.13: $4.5\ \mu\text{m}$ phase curve for HD 189733b, taken using *Spitzer's* IRAC instrument [Knutson et al., 2012]. The eclipse signals for the planet are clear and have a depth of 0.18%, while the phase curve has an amplitude of 0.09%, implying a modest day-night temperature contrast of 260 K. Note that the troughs and peaks of the phase curve model are not centred on $\phi = 0$ and 0.5, which Knutson et al. [2012] interpret as evidence of a strong equatorial jet.

of planets with low albedo include TrES-2b, Kepler-5b and Kepler-6b. However, recent results have uncovered planets that buck this trend, for example Kepler-7b for which $A_g = 0.35 \pm 0.02$ in the Kepler bandpass, with high altitude silicate clouds being a possible explanation [Demory et al., 2011].

The level of day-side thermal emission for hot Jupiter planets is not only given by the opacity sources (which set the albedo), but also the efficiency with which the irradiating energy is transported from the (permanent) day-side to the night-side of the planet. Theoretically this contrast is often described in terms of two competing processes: re-radiation of the incident stellar energy, which supports the day-night contrast, and redistribution of the energy (by advection) to the planet's night-side, which suppresses the contrast [Fortney et al., 2008].

Information on this efficiency can be derived from thermal phase curves by measuring the planet's day-night flux contrast. Large contrasts suggest a poor heat redistribution efficiency and vice versa. Examples of large contrasts have been found for WASP-18b [Maxted et al., 2013] and WASP-12b [Cowan et al., 2012], with the value for WASP-12b being particularly extreme at $\Delta T_{4.5\mu\text{m}} \simeq 1900\ \text{K}$. More modest contrasts have been measured for other planets, including in multi-band phase curve

measurements for the planet HD 189733b [see Figure 1.13; Knutson et al., 2012]. The flux maxima and minima in the HD 189733b phase curves were found to be significantly offset from their expectations ($\phi = 0.5$ and $\phi = 0$, respectively) consistent with the advection of the longitudinal temperature structure by a super-rotating equatorial jet (supporting the inference of significant advection from the day-night contrast). These phase curve results show that, even amongst these handful of planets, there is diversity in the global atmospheric properties of hot Jupiters - another example of the diversity being seen in the exoplanet population.

A greater body of observations exists for infra-red secondary eclipse detections than exists for optical secondary eclipse and thermal phase curve observations. The number of hot Jupiter planets that have had at least one such detection currently⁸ stands at over 40. Many of these have been taken with *Spitzer's* IRAC instrument in various combinations of the photometric bands this instrument offers (centred on 3.6, 4.5, 5.8 and 8.0 μm). Making use of this larger sample, Cowan & Agol [2011] carried out a statistical analysis of possible albedo and redistribution efficiencies of 24 planets. From secondary eclipse measurements, estimates of the day-side effective temperature (T_d) of each planet were derived and used to constrain the Bond albedo (A_B) and redistribution efficiency using the parameterisation:

$$T_d = T_0 (1 - A_B)^{1/4} \left(\frac{2}{3} - \frac{5}{12} \varepsilon \right)^{1/4}, \quad (1.11)$$

where $T_0 = \sqrt{\frac{R_*}{a}} T_{\text{eff}}$ is the equilibrium temperature of the planet's sub-stellar point (a measure of the irradiating flux), and T_{eff} is the stellar effective temperature. ε describes the extent to which heat is redistributed from the day- to night-side of the planet. It can take values $0 \leq \varepsilon \leq 1$, where $\varepsilon = 1$ describes the fully redistributed case (i.e. no day-night contrast - the planet emits isotropically), while $\varepsilon = 0$ is for no redistribution (i.e. instantaneous re-radiation).

For a given estimate of T_d/T_0 , A_B and ε are degenerate (equation 1.11), so unique estimates of these quantities cannot be made. However, from the ensemble of 24 planets Cowan & Agol [2011] note interesting trends in the data. They find that low Bond albedo values are favoured (typically $A_B < 0.35$), which is an independent confirmation of the conclusions from reflected light measurements. They also find that the most highly irradiated planets have uniformly high T_d/T_0 values, suggesting they have both low albedos and low redistribution efficiencies (see Figure 1.14). On the other hand, less highly irradiated planets show a variety of albedos and/or redistribution efficiencies. This trend is found to qualitatively agree with the fact

⁸<http://exoplanets.org>; accessed 01/09/14.




Figure 1.14: Day-side effective temperature estimates, normalised to the equilibrium temperature of the planet’s sub-stellar point (T_d/T_0), plotted against the irradiation level of the planet. The solid red line represents the maximum day-side effective temperature, where the planet absorbs all of the incident stellar flux and re-radiates it without redistribution. More weakly irradiated planets show a range of T_d/T_0 , while the most highly irradiated planets ($T_{\epsilon=0} > 2400$ K) tend to have high T_d/T_0 values, implying both low A_B and ϵ values. Figure from Cowan & Agol [2011].

that the time-scale for radiative processes on these planets is expected to be a stronger inverse function of temperature than the advective time-scales [Cowan & Agol, 2011]. Subsequent infra-red observations, including some of the phase curve studies highlighted earlier, have supported this trend.

1.3.2.3 Vertical temperature structures

Another example of the diversity in hot Jupiter atmospheres is the apparent dichotomy in their vertical temperature structures. Many of the planets appear to host temperature inversions in their upper atmospheres, whilst others do not. The first evidence for this split came from a pair of secondary eclipse studies for the planets HD 209458b [Knutson et al., 2008] and HD 189733b [Charbonneau et al., 2008]. Secondary eclipses were observed in the four IRAC bands and combined with previous measurements at other infra-red wavelengths. These were compared with atmospheric models by Burrows et al. [2007], which are self-consistent, plane-parallel models that assume solar abundances and chemical equilibrium. Very few parameters in the models are used to fit the data, namely simple prescriptions of the day-night energy redistribution efficiency and an unknown stratospheric absorber with a specified optical depth. Knutson et al. [2008] found that the spectral measurements for

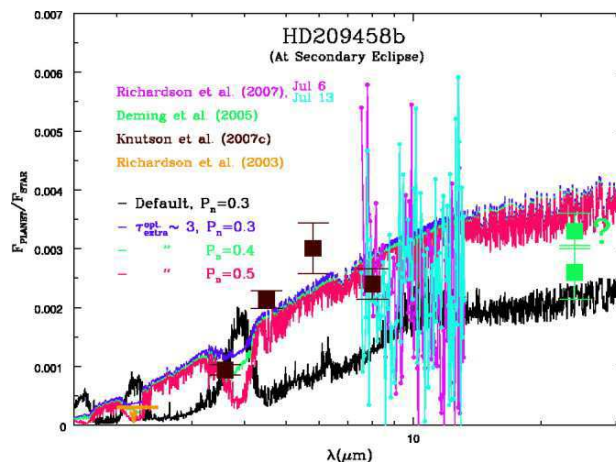


Figure 1.15: The planet-to-star flux ratio spectrum of HD 209458b presented by Burrows et al. [2007]. The data from Knutson et al. [2008] (shown as brown squares) are better fitted by models with temperature inversions (pink, green, blue lines) than those without (black line). The black model is dominated by H_2O opacity across these wavelengths. The peak at $3\ \mu\text{m}$ is due to a window in the H_2O opacity that samples flux in the lower, hotter atmosphere. In the temperature inversion models these features are inverted. H_2O is now seen in emission and the $3\ \mu\text{m}$ window now samples a cooler part of the atmosphere, giving a drop in flux.

HD 209458b were anti-correlated with the expected absorption features in models without a stratospheric absorber, as shown in Figure 1.15. The authors concluded that these features were in fact being seen in emission, rather than absorption, and that this was a result of an inverted temperature profile driven by an unknown stratospheric opacity. Conversely, Charbonneau et al. [2008] found the spectrum of HD 189733b in agreement with the absorption features, suggesting a temperature profile that decreases with altitude.

Fortney et al. [2008] attempted to understand the dichotomy in terms of correlations with other system properties. They proposed that hot Jupiters could fall into two classes, with the upper atmospheres of very hot planets being driven into emission due to the presence of gaseous titanium oxide (TiO) and vanadium oxide (VO). In less highly irradiated planets these gases would condense out of the atmosphere and so would not drive an inversion. Subsequent IRAC observations for other systems showed that, while there was agreement with this prediction in some cases [e.g. Knutson et al., 2009; Wheatley et al., 2010], there were clear discrepancies. Machalek et al. [2008] found that the ‘weakly’ irradiated planet XO-1b has an inversion, while Fressin et al. [2010] found that TrES-3b is highly irradiated but




Figure 1.16: Model fits to *Spitzer* IRAC data of XO-1b by Madhusudhan [2012]. While Machalek et al. [2008] originally concluded the planet to host a temperature inversion, Madhusudhan [2012] explain the data using non-inverted atmospheres but super-solar C/O ratios. Note that a distinction between the models presented here could be provided by accurate eclipse depth measurements in the J, H and K bands, which are available from the ground.

does not host an inversion.

An extension to this classification scheme was presented by Madhusudhan [2012], where planets are classified not only by the level of irradiation they receive, but also by the ratio of carbon-to-oxygen (C/O) in their atmospheres. Variations in both C/O and irradiation lead to variations in mixing ratios of the most spectroscopically active molecules (e.g. H_2O , CO , CH_4), significantly altering their thermal spectra. Oxygen-rich atmospheres ($\text{C/O} < 1$) are still expected to display the inverted/non-inverted dichotomy due to the presence of TiO and VO , as described by Fortney et al. [2008]. However, the lack of oxygen in carbon-rich planets ($\text{C/O} > 1$) precludes the presence of TiO/VO in the upper atmosphere and suppress the formation of temperature inversions for these planets, regardless of the irradiation level.

This scheme can be used to explain results that are discrepant under the classification of Fortney et al. [2008]. For example, the spectral features that seemed to support a temperature inversion for XO-1b [Machalek et al., 2008] are explained in the Madhusudhan [2012] scheme by a non-inverted atmosphere with a super-solar C/O ratio, as shown in Figure 1.16. However, many planets are not currently well defined within this scheme, even when *Spitzer* and near infra-red data are available. Additionally, the modelling technique used by Madhusudhan [2012] differs from that

of Burrows et al. [2007] and Fortney et al. [2008]. Rather than fully self-consistent models, Madhusudhan [2012] uses a ‘retrieval’ method, where a parameterised vertical temperature profile and the chemical abundances are free to vary in the model. This allows for a more detailed exploration of the potential composition and temperature structure of the planet, but with many free model parameters the fits are under constrained.

Knutson et al. [2010] highlighted another correlation, where the presence or absence of temperature inversions is seen to depend on the activity of the host star (measured using $\log R'_{\text{HK}}$). Planets around active stars are found to be non-inverted, while planets around quiet stars do display inversions, with an apparent cut-off at around $\log R'_{\text{HK}} = -4.9$, as shown in Figure 1.17. Knutson et al. [2010] propose a mechanism where an unknown absorber drives an atmospheric inversion except in the case of planets orbiting active stars, where the heightened UV flux breaks up the absorber, suppressing the inversion. In their description of the correlation, Knutson et al. [2010] introduced a useful empirical measure for the spectral slope across IRAC’s 3.6 μm and 4.5 μm bands (relative to a blackbody), which was found to correlate with interpretations of a temperature inversion from more detailed modelling. (I will refer to this empirical measure as ζ , following Anderson et al. 2011). A diagnostic such as this is useful for current atmosphere studies for a number of reasons. First, it provides a simple, model-independent way of characterising these sparsely sampled spectra which can be used to test correlations with other observables of the system. Second, the majority of secondary eclipse measurements come from *Spitzer*’s IRAC instrument and since 2009 the only available detectors onboard *Spitzer* have been the 3.6 μm and 4.5 μm IRAC channels. Therefore ζ will still be measurable for as long as *Spitzer* is in operation.

Though the correlation proposed by Knutson et al. [2010] is somewhat tentative, many of the subsequent results have been consistent with it [e.g. Anderson et al., 2011; Baskin et al., 2013]. However, some recent studies have found exceptions, including Stevenson et al. [2012] for HD 149026b and O’Rourke et al. [2014] for WASP-48b. In the latter case, WASP-48b was claimed not to host an inversion even though it orbits a quiet star. The authors note that the star is modestly metal-poor ($[\text{Fe}/\text{H}] = -0.12 \pm 0.12$) and that the observed lack of an inversion could be a result of a deficiency of the unknown absorber in the planet’s atmosphere. Observational correlations between the inverted/non-inverted nature of these planets and their metallicities have not yet been studied in the literature, but are explored in Chapter 4 of this thesis.

The nature of the absorber driving the inversions also remains uncertain. The

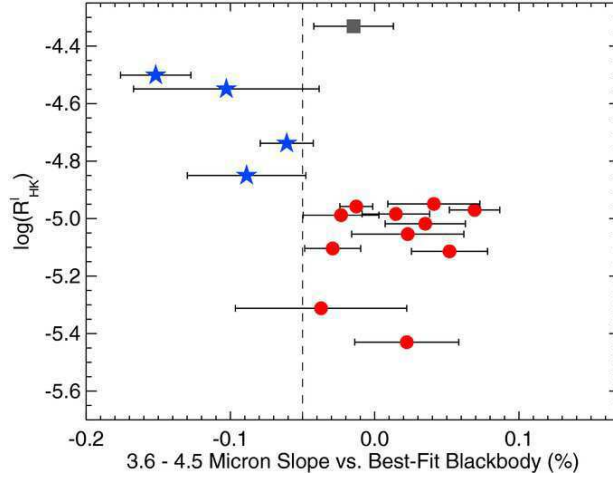


Figure 1.17: Plot from Knutson et al. [2010] showing the 3.6 – 4.5 μm slope of the planetary flux (ζ) against $\log R'_{\text{HK}}$ - a measure of the activity of the host star. Additionally points coloured red are planets which have been found to favour atmospheric models with inverted atmospheres, while those in blue do not. Thus, ζ provides a simple test of the nature of a planet’s vertical temperature structure. ζ is seen to correlate with the activity of the host, with an apparent cut-off between inverted and non-inverted atmospheres at around $\log R'_{\text{HK}} = -4.9$.

initial suggestion of TiO/VO from Fortney et al. [2008] is still a candidate, but there are arguments against this. Spiegel et al. [2009] highlighted a ‘cold trap’ mechanism expected to affect hot Jupiters, where condensation of Ti and V bearing species could remove the opacity from the planet’s photosphere. Observationally, Huitson et al. [2013] and Sing et al. [2013] rule out the presence for TiO in the transmission spectra of WASP-19b and WASP-12b (though note that the emission spectra of these planets suggest they do not host an inversion). Zahnle et al. [2009] suggested sulphur as an alternative, finding that inversions could result from absorption of UV and violet visible light by HS and S₂. However, they have difficulty explaining the inversion dichotomy, finding that sulphur-driven inversions would be common for a wide range of temperatures, irradiation levels and metallicities.

1.3.2.4 Compositions

In parallel to the work on temperature structures, there have been efforts to detect various molecules and atoms in the atmospheres of hot Jupiters. Although some early attempts at molecular detections remain controversial [e.g. detections of H₂O, CH₄, CO and CO₂ in the thermal and transmission spectra of HD 189733b, Swain




Figure 1.18: The optical transmission spectrum of HD 209458b as presented by Sing et al. [2008b]. The blue slope suggests Rayleigh scattering, while a sodium feature, with both a narrow core and broad wings, is also apparent. Rayleigh scattering is not as dominant here as it is for HD 189733b (Figure 1.19) and we are able to probe deep layers of the atmosphere where sodium is pressure broadened.

et al., 2008, 2009; Gibson et al., 2011], Grillmair et al. [2008] did find a tentative detection of H_2O in the near infra-red thermal spectrum of HD 189733b. More recent searches using the HST Wide Field Camera 3 (WFC3) have been successful, for example the detections of near infra-red H_2O signatures for WASP-19b [Huitson et al., 2013], HAT-P-1b [Wakeford et al., 2013], XO-1b and HD 209458b [Deming et al., 2013]. However these detections remain very difficult and often cannot reach the precision required to make clear detections [Ranjan et al., 2014].

Madhusudhan et al. [2011] used their retrieval method to extract compositional information from the thermal spectrum of WASP-12b across a range of infra-red wavelengths. They determined that the planet’s atmosphere was depleted in H_2O and CH_4 , explained by a carbon-to-oxygen ratio of $\text{C}/\text{O} \geq 1$. The latter result could have implications for formation mechanisms, since it is conceivable that the formation site of a giant planet, with respect to ice lines for various molecules (e.g. H_2O , CO_2) could impact this ratio [Madhusudhan et al., 2014]. However, more recent analyses of the system have revealed an M0 dwarf lying $1''$ from WASP-12. After accounting for this companion, the thermal spectrum is well described by a blackbody model, highlighting the insecure nature of some of these results.

Other work into atmospheric compositions has been carried out through detections of scattering opacities due to aerosols [e.g. WASP-12b, Sing et al., 2013] and the distinct opacity features of alkali metals [e.g. XO-2b, Sing et al. 2011; HAT-P-1b, Nikolov et al. 2014]. The well studied planets, HD 209458b and HD 189733b

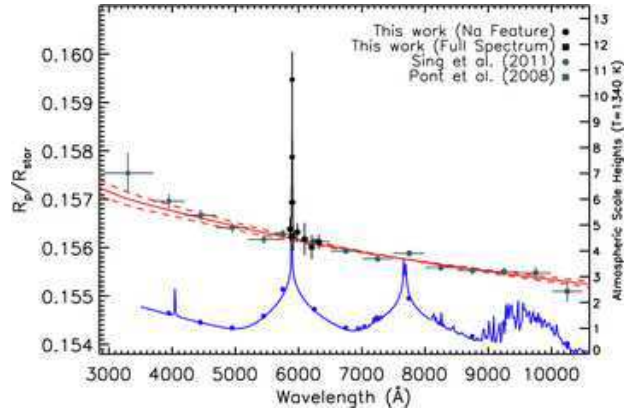


Figure 1.19: The optical transmission spectrum of HD 189733b as presented by Huitson et al. [2012]. The broad slope across the spectrum is consistent with Rayleigh scattering (red models) by a haze of condensates, while the narrow core of sodium is clearly visible. The lack of broad wings in the sodium feature (see blue model) is suggestive of a high altitude haze of condensates that blocks our view of the deeper regions of the atmosphere, where the sodium feature is subject to pressure broadening.

form an interesting pair of objects in terms of these opacities. HD 209458b has been found to have a sodium (Na) feature with broad wings, weak H₂O absorption and evidence for some Rayleigh scattering in the blue optical [see Figure 1.18; Sing et al., 2008a; Deming et al., 2013]. In contrast, the optical transmission spectrum of HD 189733b has been found to be dominated by a Rayleigh scattering signature, speculated to be due to an extended hazy/dusty atmosphere, with only the narrow cores of sodium and potassium as an additional feature [see Figure 1.19; Huitson et al., 2012; Pont et al., 2013]. In addition no evidence is found for a H₂O feature in the near infra-red [Gibson et al., 2012]. The features of these transmission spectra are consistent, since broad Na wings are a result of pressure-broadening in the lower atmosphere. HD 189733b’s lack of broad Na wings is consistent with a strong scattering opacity blocking our view of the lower atmosphere, while HD 209458b’s weaker scattering allows us to probe deeper, to the pressure-broadened regions. The different properties of these two spectra again hint at the potential diversity amongst exoplanet atmospheres. Pont et al. [2013] even suggest that the dichotomy described in Section 1.3.2.3 could be explained in terms of whether the atmosphere is dusty or not.

1.4 My work

In this thesis I have analysed secondary eclipse measurements, at both optical and infra-red wavelengths, of five planets found by the WASP transit survey: WASP-3b, WASP-21b, WASP-28b, WASP-33b and WASP-37b. These systems comprise bright stars ($V = 8.3\text{--}12.7$), hosting Jupiter-sized planets ($R_p = 1.07\text{--}1.50 R_J$) in tight orbits ($a = 0.026\text{--}0.052$ au). As a result, these planets are amenable to atmospheric characterisation through secondary eclipse measurements.

In Chapter 3, I analyse the infra-red secondary eclipses of WASP-3b, using the IRAC instrument onboard *Spitzer*. These observations were taken at $3.6\ \mu\text{m}$, $4.5\ \mu\text{m}$ and $8.0\ \mu\text{m}$ and can be used to constrain the vertical temperature structure and global energy properties of the planet. Additionally, WASP-3 is a moderately active star ($\log R'_{\text{HK}} = -4.87$) that probes the cut-off (at $\log R'_{\text{HK}} \sim -4.9$) in the proposed activity-inversion relation suggested by Knutson et al. [2010]. A conclusive detection of an inverted or non-inverted temperature structure for WASP-3b will be useful in characterising this cut-off.

In Chapter 4, I present infra-red secondary eclipse measurements of three low metallicity systems (WASP-21, WASP-28 and WASP-37). Again, these eclipses were observed using *Spitzer's* IRAC instrument, with measurements for each system taken at $3.6\ \mu\text{m}$ and $4.5\ \mu\text{m}$. The derived planetary flux estimates at these wavelengths are used to explore potential correlations between planetary spectral properties and host star metallicities.

In Chapter 5, I analyse ground-based secondary eclipse observations of WASP-33, using the ULTRACAM instrument on the William Herschel Telescope (WHT). WASP-33 is a δ Scuti pulsating star, with pulsation amplitudes being present at a similar level to the expected secondary eclipse signals. Therefore a large part of the analysis in this chapter concerns the removal of these pulsations. The secondary eclipse observations were made at both optical and near infra-red wavelengths, allowing for constraints to be placed on the reflected and thermal properties of the planet. In particular, the near infra-red z' band eclipse samples the spectral peak of the planet, providing a strong constraint on its effective temperature.

Chapter 2

Instrumentation and Methods

2.1 The *Spitzer Space Telescope* and the Infra-Red Array Camera (IRAC)

Chapters 3 and 4 of this thesis concern infra-red secondary eclipse observations of hot Jupiter exoplanets taken using the *Spitzer Space Telescope*. In this section I describe the telescope and the instrument used for these observations - the Infra-Red Array Camera (IRAC).

2.1.1 Telescope description

The *Spitzer Space Telescope*¹ is an infra-red observatory, launched by NASA on 2003 August 25 into an Earth-trailing heliocentric orbit. It comprises an 85 cm Ritchey-Chrétien mirror that focusses infra-red light onto one of three cryogenically cooled instruments. Together these instruments cover a wavelength range of 3.6 – 160 μm . They are:

The Infra-Red Array Camera (IRAC): Provides imaging at 3.6, 4.5, 5.8 and 8.0 μm .

The Infra-Red Spectrograph (IRS): Provides low-resolution, long-slit spectroscopy from 5.2 – 38 μm and high-resolution echelle spectroscopy from 9.9 – 37.2 μm . IRS can also provide imaging at 16 μm and 22 μm .

The Multiband Imaging Photometer for Spitzer (MIPS): Provides imaging at 24, 70 and 160 μm . MIPS can also provide low-resolution spectroscopy from 55 – 95 μm .

¹<http://irsa.ipac.caltech.edu/data/SPITZER/docs/spitzermission/missionoverview/spitzertelescopehandbook/>




Figure 2.1: *Spitzer's* Cryogenic Telescope Assembly, showing the telescope, Multiple Instrument Chamber, the liquid helium tank and the solar panel shield. Figure from the *Spitzer Space Telescope Handbook*.¹

The work in this thesis focuses on observations taken using the IRAC instrument. A more detailed description of this instrument is given in Section 2.1.2.

A diagram of *Spitzer's* cryogenic telescope assembly is given in Figure 2.1. This comprises the telescope itself (i.e. the primary and secondary mirrors), the cryostat (containing a tank of liquid helium and the Multiple Instrument Chamber) and an outer shell including a sun shield. The Multiple Instrument Chamber (MIC) contains the cold parts of *Spitzer's* three instruments. These are in contact with the helium bath at all times which, during *Spitzer's* cryogenic mission, was at a temperature of 1.24 K. Heat dissipation from the instruments created helium vapour, which was vented from the cryostat and used to cool the telescope to temperatures of 6 – 12 K. The cooling of the instruments and telescope was important because without it many of the instrument components would be unusable, due to the heightened thermal background.

This *cryogenic* phase of *Spitzer's* mission lasted from its launch until the liquid helium was exhausted in May 2009. From July 2009, *Spitzer* has been op-

erating its ‘warm mission’ where, without the coolant, the telescope operates at a temperature of around 28 K. In this mode, all the instruments except for the two shortest wavelength bands of IRAC (3.6 and 4.5 μm) are no longer usable. The two channels of IRAC operate with a higher thermal background and thus decreased sensitivity, but they are still capable of measuring eclipses of hot Jupiters.

A key aspect of the *Spitzer* mission is the orbit - an Earth-trailing heliocentric orbit (*Spitzer* is currently 1.3 au from Earth). This provides a cold, thermally stable environment, which contributed to the length of the cryogenic mission and has allowed *Spitzer* to carry out its warm operations. This orbit has other advantages too. For example, areas of sky are viewable for extended periods of time and not subject to the diurnal cycles of ground-based telescopes or the orbital cycles of telescopes in Earth-orbit (e.g. Hubble). The main restriction on *Spitzer’s* pointing is on the telescope boresight-sun angle which must be between 82.5 and 120°, meaning that targets are visible for at least 2 periods of about 40 days each year. Thus *Spitzer* is suitable for observations of exoplanet transits and secondary eclipses, which typically last around 6-8 hours (including baseline measurements) and even full phase-curve observations [e.g. Maxted et al., 2013].

Exoplanet time-series observations taken with *Spitzer’s* IRAC instrument typically use a staring technique (as opposed to dithered observations). This reduces systematic effects associated with array location dependent differences in measured flux [Harrington et al., 2007] by pointing and tracking a target to sub-pixel accuracies. During the cryogenic mission *Spitzer* was capable of pointing to a target with 0.5 pixel accuracy, while star trackers allow for tracking at ~ 0.2 pixel accuracy over the course of ~ 8 h. The precision is limited by the cycling of a heater which is believed to cause periodic variations in the relative pointing of the telescope boresight and the star tracker. The result is a wobble of the telescope’s pointing on the sky, with a period of ~ 1 h. In order to reduce systematic flux variations associated with this wobble (see Section 2.1.4) a new observing mode has been introduced that utilises *Spitzer’s* Pointing Calibration and Reference Sensor (PCRS) system to allow positioning of targets to a precision of 0.1 pixels.

2.1.2 IRAC description

The Infra-Red Array Camera [IRAC; Fazio et al., 2004] onboard *Spitzer* provides imaging at 3.6, 4.5, 5.8 and 8.0 μm . Following Fazio et al. [2004], I will refer to these as channels 1, 2, 3 and 4, respectively. Light from the telescope is focussed onto one of four detectors using a series of pickoff mirrors, lenses, dichroic beam splitters and filters (see Fig 2.2). Each detector measures light in a dedicated wavelength band




Figure 2.2: The optical layout of the IRAC instrument, from a side view (top image) and a top-down view (bottom image). These show the optical path of beams from the pickoff mirrors onto the IRAC InSb (channels 1 and 2) and Si:As (channels 3 and 4) detectors. The side view highlights that channels 1 and 3 image one field-of-view, while channels 2 and 4 image another field-of-view. Image taken from Fazio et al. [2004]

(see Fig. 2.3 for the response curves) and images one of two $5.2' \times 5.2'$ fields of view. One field is imaged by channels 1 and 3, while the other is imaged by channels 2 and 4. Data can be taken simultaneously by all four detectors, but since the two fields of view do not overlap, simultaneous observations of a particular target can only be made by the detector pairs that image the same field.

Each detector contains 256×256 pixels, giving a pixel scale of $1.2''$ in all four bands. Each can be operated in a ‘sub-array’ mode where only a sub-region of 32×32 pixels towards the edge of the full array are used. This mode is useful for taking shorter exposures (which are used to avoid saturation in bright stars) while not requiring large amounts of memory on the onboard computer.

The IRAC detectors are infra-red arrays. These share features with the charge-coupled devices (CCDs) used in optical astronomy, principally in that they are semi-conductor arrays comprised of p-n junctions that act as photodiodes [Dres-

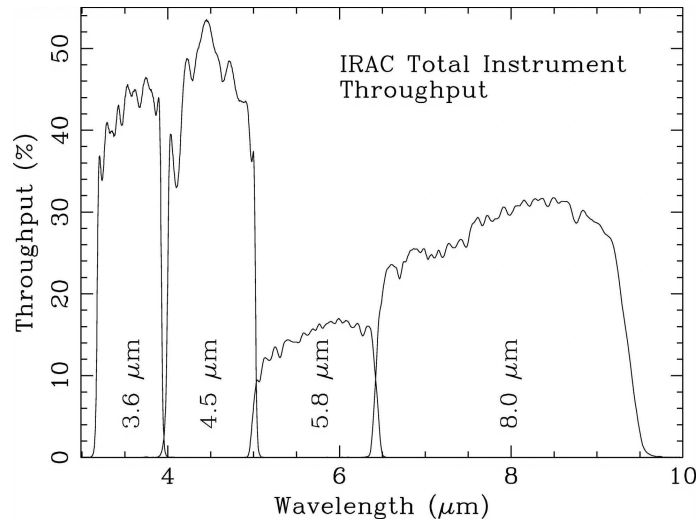


Figure 2.3: *Spitzer* IRAC transmission curves for channels 1–4. The bands are centred on 3.6, 4.5, 5.8 and 8.0 μm , respectively and have full-width at half maximum (FWHM) values of 0.7, 1.0, 1.4 and 2.8 μm . Image from Fazio et al. [2004].

sel, 2012]. Incoming infra-red photons create free electron-hole pairs near the p-n junction, where electrons are raised from the valence band to the conduction band and thus are free to move. The electric field set up by the p-n junction separates the free electron-hole pair, with the electrons being stored in the n-type semi-conductor. The amount of charge accumulated is therefore proportional to the number of photons illuminating the pixel.

In infra-red detectors the voltage change across the p-n junction associated with the charge accumulation is read out individually for each pixel by a dedicated readout amplifier. Pixel readouts are then sequentially connected to an output (multiplexing). For the IRAC detectors, a four-channel readout is used, with four columns being read out row-by-row, simultaneously.

The reading of pixel voltages is non-destructive i.e. charge is not removed from the pixel in order to be read out (as it is in CCDs). This affects how pixel values for a single exposure are measured. Fowler sampling is used, where successive voltage reads are made at the beginning of an exposure (pedestal levels) and the same number are made towards the end (signal levels; see Figure 2.4 for a schematic diagram of this technique). A single measurement for the pixel for a given exposure is determined as the average of the pedestal voltages subtracted from the average of the signal voltages [Fazio et al., 2004]. By averaging in this way, readout noise, which can be significant for infra-red arrays, can be reduced by a factor of \sqrt{N} (where N

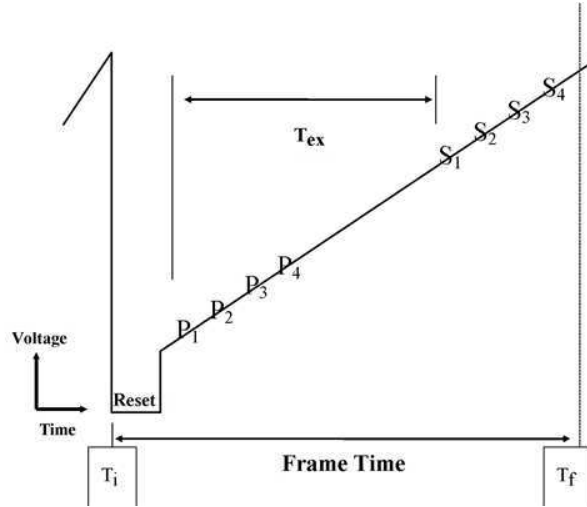


Figure 2.4: A schematic of how pixel count measurements are made for infra-red arrays like the IRAC detectors, using Fowler sampling. Successive voltage reads are made at the beginning (pedestal levels; p_i) and end (signal levels; s_i) of an exposure. The final value for the exposure is given by the average of the $s_i - p_i$ values. Effective exposure times are therefore the time difference between the corresponding p_i and s_i measurements, while the total frame time spans all of these measurements.

is the number of pedestal or signal reads).

A key choice for the IRAC detectors is the material used. The energy required to raise electrons from the valence band to the conduction band (the band gap energy) in the semi-conductor must be small enough so that infra-red photons can cause this change. Silicon, which is used in optical CCDs, has a band gap energy of ~ 1.1 eV which corresponds to a $\lambda_{\max} \sim 1.1 \mu\text{m}$. Infra-red photons would therefore not be detected in a silicon detector. For the 3.6 and 4.5 μm IRAC detectors, InSb is used - a silicon-like material that has a band gap energy of ~ 0.2 eV corresponding to a $\lambda_{\max} \sim 6.2 \mu\text{m}$. As shown in Figure 2.3, the bandpasses for these detectors are well matched to this cut-off. For the 5.8 and 8.0 μm detectors arsenic doped silicon (Si:As) is used. The band gap energy here is ~ 0.05 eV which has $\lambda_{\max} \sim 25 \mu\text{m}$ - again a suitable material for these bandpasses.

The respective sizes of these bandgap energies explains why the 3.6 and 4.5 μm detectors can operate during the warm mission, while the 5.8 and 8.0 μm detectors cannot. The longer wavelength detectors, with a smaller bandgap, are more susceptible to electrons being thermally excited into the conduction band, and at the temperatures of the warm mission (29 K) the resulting thermal noise at these wavelengths is too great for these detectors to be effective.

2.1.3 The Basic Calibrated Data (BCD) pipeline

For the *Spitzer* IRAC secondary eclipse observations I present in Chapters 3 and 4, the starting point for my analysis was the Basic Calibrated Data (BCD) frames provided by the Spitzer Heritage Archive (SHA).² The BCD pipeline³ is used to produce the standard data output frames from the IRAC instrument. It accounts for a number of well-understood systematics and produces frames with flux calibrated pixel values. Descriptions of the main steps involved in the BCD pipeline are detailed here.

Fowler normalisation: The successive pixel voltage reads in the Fowler sampling are summed, so these values must be divided by N (the number of pedestal or signal reads) to give the true pixel signal.

Lab dark subtraction: IRAC data show significant voltage offsets that are dependent on the time since the previous frame, the Fowler sampling (number and time between voltage reads) and the exposure time. For full array data, a library of lab darks, taken with different set-ups (e.g. time delays, exposure times) is interpolated creating a lab dark specific to the frame being calibrated. For sub-array data a single mean lab dark is applied (for a given detector), since sub-array lab dark data are sparse. Also, during the warm mission no lab darks are available so this step is not carried out.

Linearisation: The IRAC detectors are somewhat non-linear, with the effect becoming significant at around half full-well capacity. Pixel signals here are several percent lower than the prediction from the low signal linearity. This effect is accounted for by linearising each frame pixel-by-pixel, based on the pixel signal level, the frame time and a ‘linearity solution’. The linearity solution has been determined for each IRAC detector through ground and flight-based tests, for both the cryogenic and the warm phases of *Spitzer’s* operations.

Sky dark subtraction: Sky darks are used to account for changes to the dark and bias levels of the IRAC detectors during the mission. They are created by observing low zodiacal background regions (that are continually observable by *Spitzer*) at least twice per IRAC campaign (a 1-3 week period of IRAC observations). Sky darks are produced for a range of Fowler samplings and frame times and are applied accordingly to the science data. Before being applied to the (lab dark subtracted) science frames, the sky darks are themselves lab dark

²<http://irsa.ipac.caltech.edu/data/SPITZER/docs/spitzerdataarchives/>

³<http://irsa.ipac.caltech.edu/data/SPITZER/docs/irac/iracinstrumenthandbook>

subtracted. This is necessary because the science frames and sky darks will have, for example, different inter-frame delay times and the resulting voltage offsets need to be corrected separately. Since lab darks are not available in the warm mission, the sky darks are subtracted from the science frames without lab dark subtraction.

Flat-fielding: Infra-red arrays display pixel-to-pixel sensitivity variations, which can be mitigated by dividing data frames by a flat field frame. For IRAC, flat fields are created by observing regions of high zodiacal light and a low density of stars, to produce an even illumination. This is done twice per IRAC campaign. Dithering is used to ensure pixels are not constantly sampling light from stars or galaxies. The flats are calibrated in the same way as science frames, then averaged (with outlier rejection) and normalised to 1 before being applied. It has been found that the flat fields are unchanging, so a ‘super sky flat’ is used which combines years worth of data, giving a very high quality flat field (with pixel-to-pixel RMS values of 0.14 %, 0.09 %, 0.07 %, and 0.01 % in channels 1–4, respectively). Super sky flats are created using the same method, but separately, for the warm mission.

Flux conversion: The final BCD science data products are given in units of MJy/sr. In order to convert the science frames into these units regular observations of standard stars are made within each observing campaign, which act as flux calibrators.

The IRAC BCD data come in the form of FITS files, with one file per exposure, per channel. Before flux measurements (using aperture photometry) can be made several additional modifications are performed. These modifications are described on a case-by-case basis in Sections 3.3.1 (WASP-3), 4.3.1 (WASP-21), 4.4.1 (WASP-28) and 4.5.1 (WASP-37).

2.1.4 IRAC systematic features

The secondary eclipse light curves presented in Chapters 3 and 4 of this thesis were taken using a staring mode for IRAC, where the target star is centred on the same pixel throughout the observation. Observations typically last for around 8 h, to allow for out-of-eclipse baseline measurements (an eclipse typically lasts 2 – 3 h). Staring mitigates systematics caused by array location dependent effects which affect observations using the more standard dithering technique [Harrington et al., 2007]. However, significant systematic features still remain, namely the *intra-pixel*

sensitivity variation effect for channels 1 and 2 (InSb detectors) and the *detector ramp* effect for channels 3 and 4 (Si:As detectors).

Intra-pixel sensitivity variations

The intra-pixel sensitivity effect is a quasi-periodic variation in flux seen in the channel 1 and 2 detectors for staring data of bright sources. The variations have a typical amplitude of 1% on a time scale of 1 h. Figure 2.5 shows observations of the star-planet system TrES-4 by Knutson et al. [2009]. The effect is clearly seen and is stronger in channel 1 than in channel 2 - a feature that has been found generally. Charbonneau et al. [2005] found that these variations correlated with the centroid position of the point-spread function (PSF), which typically varies by 0.2 pixels across a secondary eclipse observation.

The physical origin of this effect is believed to be a variation of the sensitivity within each pixel, where the centre of a pixel is more sensitive than its edges. The observed PSF is therefore a combination of the true PSF and the intra-pixel sensitivity functions of the illuminated pixels. Since the point source full-width at half maximum (FWHM) for these detectors is 1.5 pixels, the PSFs are under-sampled and the central brightest pixel dominates the flux and the observed flux variation. The contribution to the flux variation from other pixels is also suppressed by the symmetry of the PSF. As the centre of the PSF moves from the centre of the brightest pixel to the edge, the total measured flux of the star decreases (and vice versa). The ~ 1 h variations in flux seen in light curves result from the telescope pointing variations described in Section 2.1.1.

In this thesis I corrected for this effect using a polynomial fit to the measured x and y positions of the PSF centre. This is a standard technique that has been used in many exoplanet eclipse studies [e.g. Charbonneau et al., 2005; Knutson et al., 2008; Todorov et al., 2012; Anderson et al., 2013]. In general the variations were modelled using subsets of the equation:

$$Flux = a_0 + a_x dx + a_y dy + a_{xx} dx^2 + a_{yy} dy^2 + a_t dt, \quad (2.1)$$

where $dx = x - \hat{x}$ and $dy = y - \hat{y}$ are the positions of the PSF centre relative to their weighted means, dt is the time since the first observation and a_0 , a_x , a_y , a_{xx} , a_{yy} and a_t are coefficients (determined as described in Section 2.4.2.3).

In some cases I tested additional terms, such as a cross term in position ($a_{xy} dx dy$) or a log temporal term ($a_t \ln(dt)$), although these were never adopted for the final analyses. The linear temporal term ($a_t dt$) has been found to be necessary

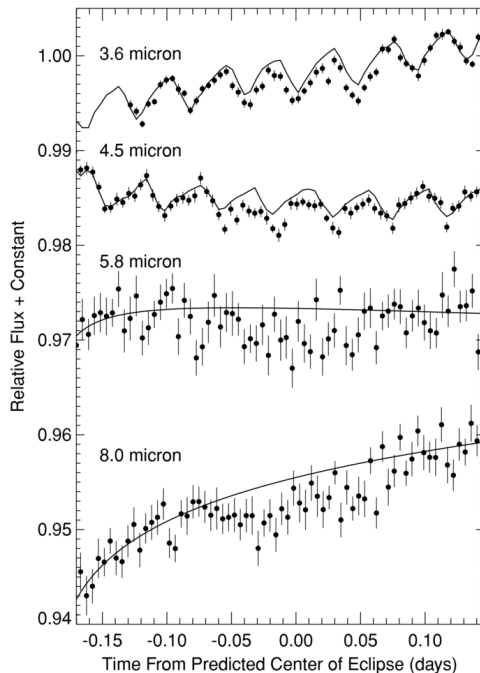


Figure 2.5: Raw *Spitzer* IRAC light curves of the secondary eclipse of TrES-4b by Knutson et al. [2009]. The quasi-periodic flux variations caused by the intra-pixel sensitivity effect is clearly seen in channels 1 and 2, while the detector ramp effect can be seen in channel 4. The slight decrease in flux in channel 3 is unexplained, but may be due a linear temporal trend in the weakly illuminated background pixels that continues after the bright source pixels have levelled off. A linear model was used for the intra-pixel sensitivity variations, while a quadratic log-time model was used for the ramps.

in some studies [e.g. Knutson et al., 2009], perhaps due to a weak detector ramp in these detectors (see below). Figure 2.5 shows a linear model in x and y applied to secondary eclipse data for the TrES-4, for channels 1 and 2. It can be seen that the intra-pixel sensitivity variations are well accounted for by such a model. The disagreements between the models and the data in the middle of the observations are the eclipse signals.

Alternative techniques to remove this systematic are used also in the literature. For example, Ballard et al. [2010] introduced a technique to create a point-by-point sensitivity map based on the flux measurements weighted by Gaussian functions in both spatial directions. Similarly, Stevenson et al. [2012] create a sensitivity map based on bi-linear interpolation (in space) of a grid of modelled flux ‘knots’, that span the range of x and y . The idea of these techniques is to account

for both large and small scale sensitivity variations. These techniques can improve the signal-to-noise on the residuals, but give consistent results when compared with polynomial fits [Stevenson et al., 2012; Blečić et al., 2013].

Detector ramp

The detector ramp effect is a systematic effect seen in *Spitzer* staring mode observations for the channel 3 and 4 detectors. It is seen as a temporal rise in measured flux from individual pixels that is dependent on its illumination history [Knutson et al., 2007]. The effect for highly illuminated pixels tends to level off after a few hours, while less bright pixels have a linear trend over the time-scale of a typical exoplanet eclipse observation. In many cases the target object is too faint for the ramps in individual pixels to be seen (changes in pixel fluxes due to telescope pointing variations dominate). However, the total flux from a star does show a clear ramp which is a sum of the individual pixel ramps (with the periodic flux variations being suppressed).

The physical explanation for this effect is believed to be that impurities in the Si:As detector material create charge traps which mask some fraction of the photoelectrons from the readout. In time, more and more charge traps are filled, meaning that fewer photo-electrons are lost from the readout i.e. the gain for the pixel increases. These charge traps also decay, releasing electrons which are then detected in the readout, though on a longer time-scale than the traps are filled. Eventually, the total decay rate of the traps will equal the rate at which the traps are filled and the pixel gain will remain constant [Agol et al., 2010].

In this thesis, I corrected for this effect by testing models that reproduce the qualitative behaviour of the ramp effect seen in the total source flux. These models take the form:

$$Flux = a_0 + a_t dt + a_{tt} dt^2, \quad (2.2)$$

$$Flux = a_0 + a_t \ln(dt + t_{\text{off}}) + a_{tt} \ln(dt + t_{\text{off}})^2, \quad (2.3)$$

$$Flux = a_0 + a_1 \exp(a_2 dt) + a_3 \exp(a_4 dt), \quad (2.4)$$

with subsets of these equations being tested. dt is the time since the first observation and a_0 , a_1 , a_2 , a_3 , a_4 , a_t , and a_{tt} are coefficients (determined as described in Section 2.4.2.3). The parameter t_{off} was used to ensure the argument of the logarithm did not go to 0.

The latter model is physically motivated, as the behaviour of individual pixels is expected to be exponential in nature, as described in the toy model of Agol

et al. [2010]. They found correcting for individual pixels was not possible due to the pointing variations, but that a double exponential modelled the total flux sufficiently. Presumably, one of the exponential terms accounts for the ramp effect from the highly illuminated pixels, while the other models the fainter pixels.

2.2 ULTRACAM on the William Herschel Telescope

In Chapter 5 I present ground-based secondary eclipse observations of WASP-33b. These were taken with the ULTRACAM instrument on the William Herschel Telescope, both of which are described here.

2.2.1 The William Herschel Telescope (WHT)

The William Herschel Telescope (WHT) at the Roque de los Muchachos Observatory on La Palma is a 4.2m alt-azimuth telescope used for studies at optical and near infra-red wavelengths. The site has a median seeing of around 0.7 arcsec with $\sim 75\%$ of nights being clear, making it one of the best astronomical sites in the world. The WHT has a Cassegrain configuration, with a parabolic primary mirror and a convex hyperbolic secondary mirror giving a focal ratio of 11 at the Cassegrain focus. Here, one of a number of instruments can be used, including ULTRACAM which is described in Section 2.2.2.

For this thesis I used the WHT for the purpose of precise secondary eclipse observations. Part of the observing strategy was to track the target accurately such that movements across pixels on the CCD cameras were minimised. Use of the telescope's tracking capabilities, along with an autoguider, allows tracking of a target to an accuracy of ~ 1 arcsec over several hours. This was suitable for my observations, although manual pointing corrections were still needed.

2.2.2 ULTRACAM description

ULTRACAM [ULTRA-fast CAMera; Dhillon et al., 2007] is a simultaneous three-colour photometer, mounted on the WHT, that is used for high speed imaging of faint astronomical objects. It is an extremely stable instrument, with no moving parts. Light from the telescope's Cassegrain focus is collimated before two dichroic beamsplitters separate the light into three beams. The beams contain light of wavelength 3200–3900 Å, 3900–5600 Å and 5600–11000 Å and each are focussed onto a separate detector allowing for simultaneous photometry in the three bands. (Figure 2.6 shows a schematic diagram of this set up). These bands are referred to as the




Figure 2.6: A ray trace diagram showing the optical path of light from the WHT focal plane to the three ULTRACAM cameras. The incoming light is collimated, then split into three beams by two dichroics, which are then focussed onto the blue, green and red ULTRACAM detectors.

blue, green and red channels, respectively [Dhillon et al., 2007]. Typical filters that are used include: SDSS u' (in the blue channel); SDSS g' (green channel); SDSS r' and z' (both red channel). For the work in Chapter 5, I used SDSS u' and z' , while in the green channel I used a ‘blue continuum’ filter. The response curves of these filters are shown in Figure 2.7.

The three detectors are frame transfer charge-coupled devices (CCDs) with an imaging area of 1024×1024 pixels at a scale of 0.3 arcsec/pixel (5.1×5.1 arcmin FOV). The detection of photons in these CCDs shares similarities with the infra-red detectors described in Section 2.1.2. Again, they are semi-conductor devices that utilise p-n junctions to detect photons via the photo-electric effect. One important difference is how the accumulated charge in each pixel is measured. In CCDs, charge is shifted from one pixel to the next by altering voltages on electrodes connected to each pixel. For a conventional CCD, charge in the 2D array is first shifted vertically by one pixel, so that one row enters a serial register. The charge in the serial register is then shifted horizontally multiple times so that the charge in each pixel can be read out by an output amplifier. Once the entire row is read, another vertical shift is applied so the next row enters the serial register. This process is repeated until the whole frame is read out.

The ULTRACAM frame transfer CCDs differ slightly from this in that charge




Figure 2.7: Response curves of standard filters used in ULTRACAM, along with the quantum efficiency curves of the three CCDs. The colours of the filter curves signifies which channel of ULTRACAM they are used in. These filters are: SDSS u' (blue line); Blue Continuum (green solid line); SDSS g' (green dashed); SDSS z' (red solid); SDSS r' (red dashed) and SDSS i' (red dotted). The solid lines are for the filters used in the WASP-33 observations in Chapter 5. The dashed grey line is the quantum efficiency for the blue and green CCDs, while the dotted grey line is the equivalent for the red CCD.

is first shifted to a storage area, which is an unilluminated part of the CCD. From here the frame read out occurs as described above, except that two halves of the chip are read out separately by different amplifiers (to increase read out speed). Shifting charge to the storage area occurs very quickly (~ 24 ms) and as soon as this is done the next exposure can begin. The minimum exposure time is therefore set by the time required to read the array, which is around 3 s for the full frame, but this can be reduced through windowing of the CCDs and binning of the pixels.

The use of frame transfer CCDs is crucial for the science goals of ULTRACAM i.e. high speed photometry. They are particularly well suited to the secondary eclipse observations of the bright system WASP-33 ($V = 8.3$, see Chapter 5) since the high duty cycle means the photon noise is as favourable as possible. Additionally, because ULTRACAM has no moving parts, it is a very stable and reliable instrument. This aids the robust removal of systematics (e.g. airmass trends and the pulsations of WASP-33), which is key to the detection of the very small eclipse signals ($< 0.1\%$) of the planet.

ULTRACAM data comes in the form of a single file for each exposure, which contains images for the three channels. Before aperture photometry is performed, these data are calibrated using standard techniques. For my observations of WASP-




Figure 2.8: A typical *Spitzer* IRAC channel 2 ($4.5\mu\text{m}$) image of WASP-3, from Chapter 3. The green source aperture has a radius of 3.0 pixels, while the sky annulus (red circles) spans a radius of 15–40 pixels. The purple apertures were used to mask the brightest stars in the background region. Aperture photometry was carried out using the ULTRACAM pipeline here.

33, this calibration is described in Section 5.3.1.

2.3 Aperture Photometry

Throughout this thesis I have used the technique of aperture photometry to extract measurements of stellar flux from both *Spitzer* IRAC and ULTRACAM images. Example images from these instruments are given in Figures 2.8 and 2.9. Both are 2D pixel arrays on which point-sources appear as roughly axis-symmetric, Gaussian-like point-spread functions (PSFs). To extract flux measurements from these sources, the centre of the PSF is first determined and a circular aperture of a specified radius is positioned. The flux contained within this source aperture is measured by performing a weighted sum of pixel values based on whether the pixel is entirely inside, partially inside or entirely outside of the source aperture. The resulting source flux is corrected for the contribution from the background, typically by assessing the average flux in an annulus (centred on the star). The size of this annulus is chosen such that it is: far enough from the star so that the PSF wings are negligible; not so far from the star that gradients in the background make the estimate inaccurate




Figure 2.9: A typical ULTRACAM red channel image of WASP-33 and a comparison star, from Chapter 5. The green source aperture has a radius of 25 pixels, while the sky annulus spans a radius 35–50 pixels. This is one of four windows used - the window here is of size 300×370 pixels. Windowing is useful for reducing the readout time of the detectors and thus increasing the duty cycle of the observations. The images are defocussed to ~ 4 arcsec to avoid saturation whilst still allowing reasonable window sizes that can image other stars that can be used as calibrators.

and large enough such that the uncertainty in the background estimate does not contribute significantly to the total error budget.

In the following sections the details of two aperture photometry packages used in this thesis to extract flux measurements of these stars are described.

2.3.1 The ULTRACAM pipeline

The ULTRACAM pipeline [Dhillon et al., 2007] is a fully featured photometry reduction package written in C++. It was created to analyse photometric measurements taken with the ULTRACAM instrument, but can also be used for data taken with a number of different instruments, including the *Spitzer* IRAC detectors. In this thesis I used the ULTRACAM pipeline in the analyses of the WASP-3 IRAC data in Chapter 3 and the WASP-33 ULTRACAM data in Chapter 5.

The pipeline is used to extract stellar flux estimates from a series of input frames in accordance with a set of user controllable parameters which define, for

example, the centroiding method, aperture sizes and error estimation method. The first stage in the aperture photometry is to define the initial positions of apertures using an interactive program (SETAPER). This is also used to define masked regions e.g. stars in the background annulus that could affect the background estimate. Once this is complete the REDUCE program is used to perform aperture photometry on a list of files for the apertures defined in SETAPER. For data taken using the ULTRACAM instrument (Chapter 5), flux estimates are extracted for the three channels simultaneously in the REDUCE program, while for the *Spitzer* IRAC data it is done for each channel individually.

Many aspects of the reduction are controlled by the user through an input parameter file. Below I describe some of the key parts of this reduction and the user controllable aspects of them.

Centroiding

For each defined aperture, centroiding is initially carried out using a Gaussian cross-correlation to the x and y marginal sums of the PSF. The marginal sums are calculated from a box with a specified width centred on the initial position defined by SETAPER, and the Gaussian profile has a fixed FWHM, again specified by the user.

Additionally, profile fitting of the PSF can be carried out, either using a 2D Gaussian or a 2D Moffat [1969] profile fit. The Gaussian profile fit has the form:

$$\text{Flux}(x, y) \propto \exp\left(-A(x - x_c)^2 - B(x - x_c)(y - y_c) - C(y - y_c)^2\right) \quad (2.5)$$

where x_c , y_c , A , B , C and the proportionality constants are free to vary. Here, x_c and y_c control the PSF centre while A , B and C control the profile widths in x and y and the orientation of the long-axis of the PSF. The user can restrict the profile to be axis-symmetric such that $A = C$ and $B = 0$. The fit is optimised using a Levenberg-Marquardt algorithm [Press et al., 1992] to minimise χ^2 on the fit.

For Moffat [1969] profile fitting, the PSF is modelled as:

$$\text{Flux}(x, y) \propto \frac{1}{[1 + A(x - x_c) + 2B(x - x_c)(y - y_c) + C(y - y_c)]^\beta} \quad (2.6)$$

where x_c , y_c , A , B , C , β and the proportionality constants are free to vary. Again a symmetric profile can be enforced by setting $A = C$ and $B = 0$ and the fit is optimised using a Levenberg-Marquardt χ^2 minimisation algorithm. A Moffat profile offers more freedom than a Gaussian profile and allows broader wings, which can often provide a better fit to PSFs affected by atmospheric seeing i.e. for ground-

based observations.

With the 2D fit complete, the new aperture centre is given by x_c and y_c .

Source flux estimation

To estimate source flux, a circular aperture of some specified radius is laid down at the previously estimated centroid position. A weighted sum of the PSF pixel values is then performed, with pixels entirely enclosed by the aperture having weight 1 and those entirely outside the aperture having weight 0. For pixels partially inside the aperture, linear tapering is used where:

$$weight = \begin{cases} 0 & r_{\text{source}} - r < -0.5 \\ r_{\text{source}} - r + 0.5 & -0.5 \leq r_{\text{source}} - r < 0.5 \\ 1 & r_{\text{source}} - r \geq 0.5 \end{cases}, \quad (2.7)$$

Here r_{source} is the source aperture radius and r is the distance from the aperture centre to the pixel centre.

The radius of the source aperture can either be set at a fixed value for all the images being reduced or it can be allowed to vary in proportion with the PSF's full-width at half maximum (FWHM) as determined from the 2D profile fit. Which option is used, and how large the aperture should be, is dictated by the signal-to-noise of the extracted flux.

Background estimation

After the flux inside the source aperture is measured it needs to be corrected for the background contribution, to give the true stellar flux estimate. This is done using an annulus centred on the measured centroid, with specified inner and outer radii. The background per pixel is estimated as either a clipped mean (with a specified clipping level) or a median of the pixel values within this annulus (after accounting for masked regions). This estimate is scaled by the total weight found for the source aperture and is subtracted from the source flux measurement.

As for the source aperture, the annulus radii can be fixed values or scaled to the FWHM determined in the 2D profile fitting. Their values are set to avoid the wings of the PSF, but to give as accurate a measure of the background as possible.

Error estimation

Errors on the stellar flux estimates can be found in one of two ways using the ULTRACAM pipeline. The first assumes that pixel variances are in accordance

with Poisson statistics ($\sigma^2 = \textit{pixel count}$, which has contributions from the dark and sky photons and astrophysical sources) and sums these variances over the source aperture, with an additional readout noise contribution. Along with the uncertainty on source aperture pixels, the uncertainty in the background estimate is also found, again using Poisson statistics and accounting for readout noise.

The second method uses a combination of Poisson errors for the estimated stellar flux and the pixel-to-pixel variance in the background, again summed over the source aperture. Using the pixel-to-pixel background variance accounts for uncertainties due to the dark and sky photons, as well as readout noise. Again the uncertainty on the background estimate is accounted for, in this case using the pixel-to-pixel background variance. This option is insensitive to whole frame offsets, which can typically affect data from the IRAC instruments due to poorly calibrated background estimates.

2.3.2 Aperture photometry with IDL

Aperture photometry using IDL scripts was carried out on the *Spitzer* IRAC data presented in Chapter 4. FITS files are passed to a customised program that first estimates the centroid of the target star’s PSF and then estimates the stellar flux and uncertainty.

Centroiding

Centroiding is carried out using a 2D Gaussian fit (GAUSS2DFIT⁴) to data in a box (of specified width) centred on a user specified pixel. The functional form of the 2D Gaussian is as given in equation 2.5, but with $B = 0$ i.e. the long-axis of the PSF is restricted to lie along either the x or y axes. The fit is optimised using a Levenberg-Marquardt χ^2 minimisation algorithm [Press et al., 1992].

Source flux estimation

Source flux estimates are made using APER⁵, which lays down an aperture of fixed radius (specified by the user) at the centroid position determined in the 2D Gaussian fit. As in the ULTRACAM pipeline, a weighted sum of the PSF pixel values is carried out, with fully enclosed pixels having weight 1 and pixels entirely outside the aperture having weight 0. For pixels partially inside the aperture the intersection

⁴<http://www.exelisvis.com/docs/GAUSS2DFIT.html>

⁵<http://www.exelisvis.com/docs/APER.html>

is computed exactly, with the weight given as the area of the pixel covered by the aperture.

Background estimation

The background estimates are also handled by APER and are determined using a clipped mean of pixel values in an annulus centred on the stellar centroid position. The source flux estimate is then corrected for this background contribution, giving an estimate for the stellar flux. As in the ULTRACAM pipeline the annulus radii and clipping level are defined by the user.

Uncertainty estimation

Photometric uncertainties are found using a combination of Poisson errors in the stellar counts and pixel-to-pixel variations in the background annulus. Specifically:

$$\sigma_{\text{phot}} = \sqrt{F_{\star} + n_{\text{source}} \left(1 + \frac{n_{\text{source}}}{n_{\text{bkg}}} \right) \sigma_{\text{bkg}}^2} \quad (2.8)$$

where F_{\star} is the stellar flux estimate, n_{source} is the source aperture area, n_{bkg} is the background annulus area and σ_{bkg} is standard deviation of pixel values in the background annulus (post-clipping). Note that F_{\star} and σ_{bkg} are in units of electrons here.

2.4 System modelling for *Spitzer* IRAC data

The work on *Spitzer* secondary eclipses presented in Chapters 3 and 4 of this thesis involved extracting information about exoplanetary systems from a combination of photometric light curves, radial velocity data and other constraints on system parameters. In this section I describe the information that can be obtained and the techniques used to do this in the Markov Chain Monte Carlo (MCMC) code of Collier Cameron et al. [2007a] - MCMCTRANSIT

I note here that the modelling of the ULTRACAM light curves of WASP-33 involves modelling of the host star's δ Scuti pulsations, and so this is described separately in Sections 5.3.5 and 5.3.6.

2.4.1 The system model

MCMCTRANSIT uses a set of parameters from which photometric lightcurves (transits and secondary eclipses) and radial velocities are modelled, and system parameters

Table 2.1: Definitions of the proposal parameters used in MCMCTRANSIT.

Parameter	Description
T_0	Transit epoch
P	Orbital period
δ	Transit depth
T_{14}	Transit duration
b	Impact parameter
K_1	Stellar reflex velocity
T_{eff}	Stellar effective temperature
[Fe/H]	Stellar metallicity
$\sqrt{e} \cos \omega$	Combinations of the orbital eccentricity (e) and argument of periastron (ω)
$\sqrt{e} \sin \omega$	
$\Delta F = \frac{F_p}{F_\star}$	Day-side planet-to-star flux ratio ^a

^aThere could be multiple instances of this variable, depending on how many *Spitzer* channels are being analysed.

are derived. For the work on *Spitzer* data this parameter set is defined in Table 2.1.

From this set of parameters, the following properties of the system can be determined directly (see also Table 2.2 for definitions of derived quantities).

$$\frac{R_p}{R_\star} = \sqrt{\delta} \quad (2.9)$$

$$\frac{R_\star}{a} = \frac{\pi T_{14}}{P} \frac{1}{\sqrt{(1 + \frac{R_p}{R_\star})^2 - b^2}} \frac{1 + e \sin \omega}{\sqrt{1 - e^2}} \quad (2.10)$$

$$\rho_\star = \frac{3\pi}{G} \left(\frac{a}{R_\star} \right)^3 \frac{1}{P^2} \quad (2.11)$$

$$e = (\sqrt{e} \cos \omega)^2 + (\sqrt{e} \sin \omega)^2 \quad (2.12)$$

$$\omega = \arctan \left(\frac{\sqrt{e} \sin \omega}{\sqrt{e} \cos \omega} \right) \quad (2.13)$$

$$i = \arccos \left(\frac{R_\star}{a} \frac{b}{1 - e \cos E_t(e, \omega)} \right) \quad (2.14)$$

These properties are then used to determine the apparent separation of the star and planet centres, as seen from Earth, as a function of the orbital phase (ϕ) of the

Table 2.2: Definitions of parameters derived from parameters in Table 2.1.

Derived Parameter	Description
R_p	Planet radius
R_\star	Stellar radius
M_p	Planet mass
M_\star	Stellar mass
ρ_\star	Stellar density
a	Semi-major axis
i	Orbital inclination
ϕ	Orbital phase
$V_r(\phi)$	Stellar radial velocity
$z(\phi)$	Apparent separation of star and planet centres (normalised to R_p)
$E(\phi)$	Eccentric anomaly
E_t	Eccentric anomaly at mid-transit
$\theta(\phi)$	True anomaly

planet. In general

$$z(\phi) = \frac{a}{R_\star} (1 - e \cos E(\phi)) \sqrt{1 - \sin^2 i \cos^2 \left(\theta(\phi) + \omega + \frac{\pi}{2} \right)}. \quad (2.15)$$

For a circular orbit this simplifies to

$$z(\phi) = \frac{a}{R_\star} \sqrt{1 - \sin^2 i \cos^2 (2\pi\phi)}. \quad (2.16)$$

Here, the orbital phase (ϕ) is a function of the time of interest (t), the transit epoch (i.e. the mid-transit time; T_0) and the period (P), defined as:

$$t = T_0 + nP + \phi P, \quad (2.17)$$

where n is the number of full orbits the planet has undergone between T_0 and t .

Transit and secondary eclipse lightcurves are then modelled by finding the ratio of the obscured system flux to the unobscured stellar flux. This is found as a function of $z(\phi)$ and $p = R_p/R_\star = \sqrt{\delta}$, and for secondary eclipses, ΔF . Figure 2.10 shows the geometry used for this modelling. Transit light curves are modelled as described in Section 5 of Mandel & Agol [2002], using the small planet approximation

($p < 0.1$) and ignoring the flux contribution from the planet. Analytically, the system flux as a function of orbital phase is modelled as:

$$\frac{F(\phi)}{F_\star} = 1 - \lambda(\phi) \quad (2.18)$$

$$\lambda = \begin{cases} 0 & 1 + p < z \\ \frac{1}{\pi} \left[p^2 \arccos\left(\frac{z-1}{p}\right) - (z-1)\sqrt{p^2 - (z-1)^2} \right] & 1 - p < z \leq 1 + p \\ p^2 & z \leq 1 - p \end{cases} \quad (2.19)$$

Here, λ and z are functions of ϕ . Realistic modelling of transits also requires accounting for the limb-darkening of the star [e.g. Claret, 2000], however in this thesis transit light curves are not fitted to data (see Section 2.4.2.3) and so this will not be discussed further. Note that, neglecting limb-darkening, the difference between out-of-transit and in-transit flux is $p^2 = \frac{R_p^2}{R_\star^2}$ i.e. the depth of the transit is a measure of the planet-to-star radius ratio.

Secondary eclipse light curves are modelled as described in Wheatley et al. [2010], by finding the visible fraction of the planet, $\eta(\phi)$ and scaling this by the day-side planet-to-star flux ratio, ΔF :

$$\frac{F(\phi)}{F_\star} = 1 + \eta(\phi) \Delta F \quad (2.20)$$

$$\eta = \begin{cases} 1 & 1 + p < z \\ \frac{1}{\pi} \left[\beta - \cos \beta \sin \beta + \frac{(\alpha - \cos \alpha \sin \alpha)}{\pi^2} \right] & 1 - p < z \leq 1 + p \\ 0 & z \leq 1 - p \end{cases} \quad (2.21)$$

where $\alpha = \arccos\left(\frac{1-p^2+z^2}{2z}\right)$ and $\beta = \arccos\left(\frac{1-p^2-z^2}{2pz}\right)$ (also, see Figure 2.10). Separate secondary eclipse models using separate ΔF values, but the same $\eta(\phi)$ function, are used when multiple IRAC datasets are being analysed simultaneously. Note that the difference between the out-of-eclipse and in-eclipse flux is $\Delta F = \frac{F_p}{F_\star}$ i.e. the depth of the eclipse gives the day-side planet-to-star flux ratio in the given IRAC band.

If, in addition, there are radial velocity (V_{rad}) measurements, these can be modelled as a function of phase using

$$V_{\text{rad}}(\phi) = K_1 (e \cos(\omega) + \cos(\theta(\phi) + \omega)). \quad (2.22)$$




Figure 2.10: The geometry used to model the visible fraction of a planet, $\eta(\phi)$, during secondary eclipse. The angles α and β used in equation 2.20 are shown here. Note that the planet is shown in front of the star here only for clarity.

For a circular orbit this simplifies to

$$V_{\text{rad}}(\phi) = K_1 \cos\left(2\pi\phi - \frac{\pi}{2}\right) \quad (2.23)$$

In order to find unscaled system parameters (R_p , R_\star , M_p , M_\star and a) an additional constraint is required. This is provided by a stellar mass estimate, parameterised by the stellar density (equation 2.11), along with stellar effective temperature and metallicity estimates (obtained from spectra) as described by Enoch et al. [2010]. Using this mass estimate along with Kepler's third law (assuming $M_\star \gg M_p$) sets the scale of the system through

$$a = \left(\frac{GM_\star P^2}{4\pi^2}\right)^{\frac{1}{3}} \quad (2.24)$$

From here the stellar and planetary radii can be found from equations 2.9 and 2.10. With radial velocity measurements, the mass of the planet can also be found using the stellar reflex velocity:

$$M_p = \frac{K_1}{\frac{2\pi a \sin i}{P} - K_1} M_\star \quad (2.25)$$

Hence, from the parameters in Table 2.1 one can obtain a set of detailed system




Figure 2.11: 2D histograms of a typical MCMC chain for the analysis of an exoplanet system. Left panel: distribution for e and ω shows a highly correlated structure. Much of the space spanned by allowable e and ω values does not produce good fits, so exploration of this parameter space will not be efficient. Right panel: distribution for $\sqrt{e} \cos \omega$ and $\sqrt{e} \sin \omega$ shows a weaker correlation, so exploration of this parameter space will be more efficient.

parameters (see Table 2.2).⁶

The parameters in Table 2.1 were chosen for the purposes of the Markov chain Monte Carlo (MCMC) technique used for parameter optimisation (Section 2.4.2). These parameters show little to no correlation with each other [Collier Cameron et al., 2007a], ensuring the efficient running of the MCMC code.

As an example, the choice of $\sqrt{e} \cos \omega$ and $\sqrt{e} \sin \omega$ may seem like a curious one - why not use e and ω ? The reason is that e and ω are highly correlated, as shown in Fig. 2.11. The MCMC would run inefficiently as much of the space spanned by allowable e and ω values would not produce good fits. A better choice is to use $e \cos \omega$ and $e \sin \omega$ [for small e ; Ford, 2006], as these parameters are nearly orthogonal. However, this choice imposes an implicit prior on the eccentricity that is linear in e , leading to overestimates of its value. $\sqrt{e} \cos \omega$ and $\sqrt{e} \sin \omega$ are chosen because they mitigate this problem - restoring a uniform prior on e - whilst remaining nearly orthogonal [Anderson et al., 2011].

⁶There is a mistake in equation 2.25 - a factor of $\frac{1}{\sqrt{1-e^2}}$ is missing from the first term of the denominator. This reduces the M_p value by $\sim 2\%$ for the largest values of e I have found in Chapters 3 and 4, which is at the level of the uncertainties I found on M_p .

2.4.2 Parameter and uncertainty estimation

2.4.2.1 Figure of merit for parameter optimisation

The fitting of the light curve and radial velocity models to data is carried out using Bayesian methods. Under this scheme one is interested in characterising the posterior distribution - the probability distribution of models given a set of data.

For a particular model, the posterior probability is found by combining the likelihood of obtaining the dataset given the model, and the prior probability of the model:

$$p(\text{model}|\text{data}) \propto p(\text{data}|\text{model})p(\text{model}) \quad (2.26)$$

Under the condition that the data uncertainties represent a Gaussian distribution and that their values are accurate [Ford, 2005], the likelihood is given as:

$$p(\text{data}|\text{model}) \propto \exp\left(-\frac{\chi^2}{2}\right) \quad (2.27)$$

$$\chi^2 = \sum_i \frac{(\text{data}_i - \text{model}_i)^2}{\sigma_i^2} \quad (2.28)$$

χ^2 is found from the model fits to all the inputted data. In this thesis these data are secondary eclipse and radial velocity datasets.

The prior probability, $p(\text{model})$ is used to account for prior knowledge of the system. In this work they are applied as non-uniform probability distributions of system model parameters and/or the derived parameters. Specifically, parameters are restricted using Gaussian distributions:

$$p(\text{model}) = \exp\left(-\sum_{j=1}^m \frac{(p_j - p_{j,0})^2}{2\sigma_j^2}\right), \quad (2.29)$$

where $p_{j,0}$ is the best estimate for the parameter and σ_j is the error. These values are based on estimates from previous studies. Details of the priors used for the IRAC observations, and specifically their use in constraining transit parameters, can be found in Section 2.4.2.3 and in Chapters 3 and 4. For parameters where no prior is explicitly given, a uniform prior is implicitly used. This is equivalent to there being no prior knowledge of the parameter.

The posterior probability distribution is therefore given by:

$$p(\text{model}|\text{data}) \propto \exp\left(-\frac{\chi^2}{2} - \sum_{j=1}^m \frac{(p_j - p_{j,0})^2}{2\sigma_j^2}\right) \quad (2.30)$$

Model optimisation and uncertainty estimation is carried out by exploring the parameter space around the maximum of the posterior probability distribution. This is equivalent to the minimum of a figure-of-merit, Q , given by:

$$Q = \chi^2 + \sum_{j=1}^m \frac{(p_j - p_{j,0})^2}{\sigma_j^2} \quad (2.31)$$

2.4.2.2 Markov Chain Monte Carlo

Exploration of model fits to data in the MCMCTRANSIT code is carried out using a Markov Chain Monte Carlo (MCMC) technique. As described in Section 2.4.1, models are defined by a set of proposal parameters (see Table 2.1). The MCMC algorithm is used to estimate both the optimal model parameters and their uncertainties. One of the main advantages of MCMC is that it properly accounts for the inter-dependencies of model parameters in the parameter uncertainty estimates.

In an MCMC algorithm a cloud of points is created in the model parameter space, with each point corresponding to a particular instance of the model. The number of points in a region of this parameter space is proportional to the posterior probability in that region. This is achieved in MCMCTRANSIT using the Metropolis-Hastings algorithm, which is as follows:

1. Find a model that is reasonably close to the global minimum in Q (as given in equation 2.31). Use the model, M_0 , as the starting point for the MCMC ($i = 0$, where i is the i^{th} step in the chain).
2. Generate a proposal model, $M_{i+1,\text{proposal}}$. Each parameter, p , of the model is altered according to $p_{i+1,\text{proposal}} = p_i + s_p G(0, 1) f_i$. Here s_p is an estimate of the uncertainty in p ; $G(0, 1)$ is a Gaussian random number (mean of 0, variance of 1); f_i is a variable scale factor used to ensure the acceptance rate of proposals is $\sim 25\%$.
3. Calculate $Q_{i+1,\text{proposal}}$ for $M_{i+1,\text{proposal}}$.
4. Apply Metropolis-Hastings rules. If:

$$\begin{aligned} Q_{i+1,\text{proposal}} < Q_i, & \quad \text{accept model } M_{i+1,\text{proposal}} \\ Q_{i+1,\text{proposal}} \geq Q_i, & \quad \text{accept model } M_{i+1,\text{proposal}} \text{ with probability} \\ & \quad \exp(-(Q_{i+1,\text{proposal}} - Q_i)/2). \end{aligned}$$

5. If accepted set $M_{i+1} = M_{i+1,\text{proposal}}$. Then return to step 2 with $i = i + 1$

6. If rejected return to step 2 (i does not increase).
7. End the chain when i equals a specified number (typically $10^4 - 10^5$).

These rules ensure that the cloud of points created in the model parameter space will converge to the desired probability distribution [Ford, 2005].

Information on individual parameters can then be obtained by finding the marginal distribution of the parameter (i.e. by producing its histogram). Uncertainties for the parameter are estimated using the 1σ confidence intervals of its marginal distribution, while the optimal value can be estimated from the median of this distribution.

2.4.2.3 Optimal parameter and uncertainty estimation

In the usual implementation of MCMCTRANSIT, transit light curve data are provided and fitted using the model in equation 2.18 [e.g. Collier Cameron et al., 2007b; Pollacco et al., 2008]. In the implementation used in this thesis the transit fitting was bypassed through the use of priors. Constraints that are usually imposed on the parameters entering $\lambda(\phi)$ by transit data were instead applied using priors on transit light curve model parameters from previous studies. These studies are often quite detailed in their assessments of transit light curves. Consequently I saw the use of priors based on their results as a better option than another analysis of the available transit data. The use of priors still allowed for the dependencies of secondary eclipse parameters on transit light curve parameters to be explored.

For WASP-3 (Chapter 3), transit light curve priors were applied by constraining the derived parameters $\frac{R_p}{a}$, $\frac{R_\star}{a}$ and i (see Section 3.3.4). These parameters were constrained because they were used for the light curve model of Southworth [2011], which was the follow-up study used for this object. In Chapter 4 for WASP-21, $\frac{R_p}{a} + \frac{R_\star}{a}$, $\frac{R_p}{R_\star}$ and i were used to provide the transit light curve constraints [Ciceri et al., 2013, also see Section 4.3.3], while for WASP-28 and WASP-37 δ , b and T_{14} were used [Anderson et al., 2014; Simpson et al., 2011, also see Sections 4.4.3 and 4.5.3]. For each of these objects, priors were also placed on the ephemeris parameters T_0 and P (which affect the transit model through ϕ), and on T_{eff} and $[\text{Fe}/\text{H}]$ (which set the value for M_\star). These priors enter the Q statistic in accordance with equation 2.31. For example, the prior constraint on i for WASP-3 is determined as $\frac{(i-i_0)^2}{\sigma_{i_0}^2}$, where $i_0 \pm \sigma_{i_0}$ is the inclination estimate from Southworth [2011].

Secondary eclipse light curves were fitted to the flux data resulting from the aperture photometry analyses (using the ULTRACAM pipeline for WASP-3 and the IDL reduction for WASP-21, 28 and 37). The units of these data were electrons i.e.

they are not scaled or normalised to the flux of a comparison star. There were two stages to the secondary eclipse fitting. First, the secondary eclipse model (equation 2.20) was divided out of the measured fluxes to give an estimate for the stellar flux:

$$F_{\star} = \frac{F_{\text{measured}}}{1 + \eta(\phi)\Delta F} \quad (2.32)$$

Then, F_{\star} was modelled using a subset of one of the equations 2.1–2.4. For example, for IRAC channel 1 data, one could choose to model the intra-pixel sensitivity variations using:

$$F_{\star,\text{model}} = \hat{F}_{\star} + a_0 + a_x dx + a_y dy \quad (2.33)$$

The optimal coefficient values can be determined uniquely (in terms of minimising χ^2) using a singular-value decomposition (SVD) technique [Press et al., 1992]. This is possible due to the linear nature of the equations used for detrending. A slight complication is that t_{off} , a_2 and a_4 (from equations 2.3 and 2.4) cannot be evaluated using this method, so these were added as MCMC proposal parameters when in use.

Radial velocity data were modelled using equation 2.22. The combination of radial velocity and secondary eclipse data provides important constraints on e and ω . In the case of the secondary eclipses, the timing of the centre of the eclipse constrains the quantity $e \cos \omega$ through [Charbonneau et al., 2005]:

$$\Delta\phi \equiv \phi_{\text{E}} - 0.5 \simeq \frac{2}{\pi} e \cos \omega, \quad (2.34)$$

where ϕ_{E} is the observed mid-eclipse phase and 0.5 is the expectation from a circular orbit.

The starting point for the MCMC proposal parameters came either from an analysis of the SuperWASP light curves used to discover WASP objects (for WASP-3) or they were derived from the priors (for WASP-21, 28 and 37). A circular orbit was also assumed for the first model ($\sqrt{e} \cos \omega = \sqrt{e} \cos \omega = 0$) and the eclipse depths were started from 0. The starting point for K_1 was found by initially modelling the radial velocity data using a circular orbit.

These values were each perturbed by $5G(0,1)s_p$ at the beginning, to ensure areas of parameter space other than the expected solution were explored. From this starting point, a ‘burn-in’ phase occurred, to allow the optimal region of parameter space to be found. At least 2000 (accepted) steps were used for this, with the burn-in ending once the χ^2 for an accepted step was greater than the median χ^2 of all the preceding steps. This ensured that the chain had found the region of parameter space in which the optimal solution lies.

After the burn-in phase, the photometric uncertainties on the secondary eclipse data were scaled such that the reduced χ^2 for each set equalled 1. This was done to ensure that realistic uncertainties on the proposal parameters were obtained. The values of s_p were also re-evaluated by producing a chain of 1000 values and determining the standard deviations of the marginal distributions for each parameter.

A production chain of between 10^4 and 10^5 successful jumps was then made in order to fully explore the parameter space around the optimal solution, to account for the dependencies of the proposal parameters on each other and to make an honest assessment of their uncertainties [Anderson et al., 2011]. Optimal parameter estimates (for both proposal and derived parameters) were taken as the median of the marginal distribution for that parameter, while the uncertainties were taken as the 15.9 and 84.1 percentiles (to give 1σ uncertainties).

2.5 The Bayesian Information Criterion (BIC)

In order to correct for the systematic effects associated with the *Spitzer* IRAC detector, I used functions based on subsets of the equations 2.1 (for IRAC channels 1 and 2) and 2.2–2.4 (for channels 3 and 4) to model the stellar flux, as described in Section 2.4.2.3. In order to distinguish the relative quality of fits to a given dataset, using different models, χ^2 cannot be used, as adding more model parameters will at worst give the same χ^2 and would lead to over-fitting.

In this thesis, I use the Bayesian Information Criterion [BIC; Schwarz, 1978] to provide the distinction between different models. The BIC does this by weighting the likelihood of different model fits by number of free parameters. It is given by:

$$\text{BIC} = \chi^2 + k \ln N, \quad (2.35)$$

where k is the number of free model parameters and N is the number of data points being fitted. The BIC prefers more simple models (with lower k), unless there is a significant improvement in χ^2 with additional model parameters [Anderson et al., 2011].

For two given models fitting the same dataset, the relative BIC values provide an approximation to the Bayes factor, which is the ratio of the likelihoods for the two fits. This approximation is valid for large N - the smallest dataset where the BIC was used in this thesis had $N = 1500$, so the approximation should be valid. Kass & Raftery [1995] provide a useful qualitative summary of the evidence against the higher BIC value, as shown in Table 2.3.

Table 2.3: Qualitative summary of evidence provided by ΔBIC values.

ΔBIC	Bayes factor	Evidence against higher BIC
0–2	0.37–1	Weak
2–6	0.05–0.37	Positive
6–10	0.007–0.05	Strong
> 10	< 0.007	Very Strong

2.6 Temperature estimation

2.6.1 Brightness temperatures

As described in Section 2.4.1, the depth of an exoplanetary secondary eclipse gives the day-side planet-to-star flux ratio (in the wavelength band observed). This can be used to estimate the brightness temperature, T_b , of the planet. Throughout this thesis brightness temperatures have been estimated using the following method: First the photon flux of a blackbody was integrated over the response function for the photometric band being used, then divided by the stellar flux using the models of Kurucz [1993]⁷ integrated over the same response function. To obtain a flux ratio, this was scaled by the square of the planet-to-star radius ratio (δ). The temperature of the blackbody, T_b , was iterated until the above quantity matched the relevant eclipse depth.

The models of Kurucz [1993] are parameterised by the effective temperature, metallicity and surface gravity of the host star. For a given stellar effective temperature the closest matching Kurucz [1993] models in temperature were interpolated (since models are only given in steps of $\Delta T_{\text{eff}} = 250\text{ K}$) to get the relevant stellar flux model. The closest matching models in terms of metallicity and surface gravity were used, but not interpolated as it was found to have a negligible effect on the resulting brightness temperatures.

The errors on T_b were found using a Monte-Carlo technique, where 10^5 fake data sets were created based on the eclipse depth measurement and its error, assuming Gaussian errors. δ was also perturbed using its error. The above procedure was applied to each of these fake data sets, and the error on the brightness temperature was found as the 1σ confidence interval of the resulting distribution. Additionally, the best fitting brightness temperature was found using Kurucz [1993] models with effective temperatures altered by $\pm 1\sigma$. The difference between the resulting brightness temperature estimates was halved and added in quadrature to the errors from

⁷<http://www.stsci.edu/hst/observatory/crds/k93models.html>

the Monte-Carlo simulations. This was done to account for the uncertainty in the stellar temperature.

2.6.2 Effective temperatures, albedos and heat redistribution

Obtaining a representative temperature for a planet can be useful for placing simple limits on its reflectance and heat redistribution properties. A temperature estimate could perhaps be found by averaging the measured brightness temperatures across a range of wavelengths, or by determining a best-fitting blackbody model for the planet, utilising the method described above. In this thesis I have followed the method of Cowan & Agol [2011], who use an interpolation method to estimate the day-side effective temperature of the planet through:

$$\frac{\sigma}{\pi} T_d^4 = \int_0^\infty B_\lambda(T(\lambda)) d\lambda, \quad (2.36)$$

where $T(\lambda)$ is given by a piecewise model of the planetary brightness temperatures, derived as described in Section 2.6.1. The model is linear between the wavelength bands used and constant longward (shortward) of the longest (shortest) wavelength band. I assessed the errors on T_d using a Monte-Carlo method, by perturbing the brightness temperature estimates according to their errors, reassessing T_d each time and assessing the 1σ confidence interval of the resulting distribution. An additional systematic error component was added in quadrature, following Cowan & Agol [2011], to account for the fact that small numbers of brightness temperature estimates (which are sensitive to spectral features) are being used estimate the effective temperature of the planet.

Cowan & Agol [2011] then use this effective temperature estimate to constrain the Bond albedo and heat redistribution properties of the planet using:

$$T_d = T_0 (1 - A_B)^{1/4} \left(\frac{2}{3} - \frac{5}{12} \varepsilon \right)^{1/4} \quad (2.37)$$

where

$$T_0 = \sqrt{\frac{R_\star}{a}} T_{\text{eff}}. \quad (2.38)$$

T_0 values and uncertainties were found using the marginal distributions of R_\star , a and T_{eff} outputted from the final MCMC runs for each system. ε describes the extent to which heat is redistributed from the day- to night-side of the planet. It can take values $0 \leq \varepsilon \leq 1$, where $\varepsilon = 1$ describes the fully redistributed case (i.e. no day-night contrast - the planet emits isotropically), while $\varepsilon = 0$ is for no redistribution.

Cowan & Agol [2011] found some interesting trends in the population of exoplanets they analysed in this way (see Section 1.3.2.2). For example, A_B values are found to be restricted to low values (< 0.35), the most highly irradiated planets are found to have systematically high T_d/T_0 values and more moderately irradiated planets show a variety of T_d/T_0 values. The above analysis was used to place the planets studied in this thesis into the context of these trends.

Given an estimate of T_d/T_0 , only a swathe through $A_B - \varepsilon$ space can be constrained due to the degeneracy between these quantities (see equation 2.37). To show this more clearly a 2D probability distribution, based on the measured T_d/T_0 value, can be found. Assessing this 2D PDF can highlight which regions of $A_B - \varepsilon$ space are allowed by the data. Throughout this thesis this was carried out using equation 7 of Cowan & Agol [2011]:

$$\text{PDF}(A_B, \varepsilon) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[-\frac{\left(\frac{T_{d,\text{measured}}}{T_0} - \frac{T_d}{T_0}(A_B, \varepsilon) \right)^2}{2\sigma^2} \right] \quad (2.39)$$

where $\frac{T_{d,\text{measured}}}{T_0}$ is the measured value, σ is the uncertainty on this and $\frac{T_d}{T_0}(A_B, \varepsilon)$ is determined from equation 2.37 for a given A_B and ε . Equation 2.39 was evaluated over allowable A_B and ε values to create the 2D PDF. Marginalised 1D PDFs for the two quantities were also found by integrating over the other quantity.

Additionally, the parameterisation given in equation 2.37 was used to quote representative equilibrium temperatures for planets throughout this thesis (with specified A_B and ε values).

Chapter 3

The thermal emission of the exoplanet WASP-3b

3.1 Introduction

WASP-3b was the third planet discovered by the WASP project [Pollacco et al., 2008] and is an $M_p = 1.8 - 2.0 M_J$, $R_p = 1.3 - 1.4 R_J$ planet orbiting a late F-type star with a 1.8 d period and a separation of $a = 0.03$ au [Southworth, 2011; Maciejewski et al., 2013]. Numerous follow-up studies of the system have been made, including: space-based transit photometry [Christiansen et al., 2011]; Rossiter-Maclaughlin measurements, which have revealed a prograde orbit for the planet [Simpson et al., 2010; Miller et al., 2010] and transit timing variation measurements which initially hinted at the possibility of another planet in the system, but this scenario has subsequently been ruled out [Maciejewski et al., 2010; Montalto et al., 2012; Maciejewski et al., 2013].

The planet receives intense irradiation from its host and this makes it an excellent target for studies of its thermal emission, with an equilibrium temperature ($A_B = 0$, $\varepsilon = 1$) of 1990 K. A ground-based secondary eclipse observation in the K_s band has already been carried out by Zhao et al. [2012]. The brightness temperature derived from the eclipse measurement was around 2700 K - larger than the maximum expected from simple equilibrium temperature calculations (e.g. equation 2.37 with $A_B = 0$ and $\varepsilon = 0$). The *Spitzer* IRAC secondary eclipse observations of WASP-3b presented in this chapter reveal the planet's thermal emission over a wider range of infra-red wavelengths, giving more insight into the planet's energy output and vertical temperature structure.

With regard to the activity-inversion correlation suggested by Knutson et al.

[2010] (see Section 1.3.2.3), the activity measure for WASP-3 [$\log R'_{\text{HK}} = -4.872$, Knutson et al., 2010] puts it interestingly near the expected cut-off between inverted and non-inverted atmospheres ($\log R'_{\text{HK}} \sim -4.9$). A conclusive determination of the nature of WASP-3b's inversion will be useful in characterising this cut-off. Equilibrium temperature estimates also suggest WASP-3b could be part of the group of planets with low albedo and heat redistribution efficiencies highlighted by Cowan & Agol [2011] (Section 1.3.2.2). Accurate brightness temperature estimates across the IRAC bands will allow a test of this hypothesis.

The remainder of this chapter is ordered as follows: in Section 3.2 the secondary eclipse observations taken using *Spitzer* are described; the analysis carried out on these observations, including light curve model fitting, is described in Section 3.3; in Section 3.4 the results of this analysis are presented and in Section 3.5 these results are discussed in the context described above. Finally, these findings are summarised in Section 3.6.

3.2 Observations

Secondary eclipse observations for WASP-3 were taken using the IRAC instrument onboard the *Spitzer Space Telescope*. Both of these are described in Section 2.1. One set of observations were taken at the predicted time of secondary eclipse on 18 September 2008, during the cryogenic mission. Simultaneous exposures of 12 s (effective integration time of 10.4 s) were taken in full-array mode at 4.5 μm (channel 2) and 8.0 μm (channel 4) over a period of 5.8 hours, providing 1559 images in each channel (program ID 50759, PI P. Wheatley).

Secondary eclipse measurements were also made at 3.6 μm (channel 1) on 26 October 2009, during the warm mission. Observations were taken using IRAC's sub-array mode with a frame time of 2 s (1.92 s effective exposure), over 7.7 hours, providing 13670 individual images (program ID 60021, PI H. Knutson). The combination of the 2 s frame time and sub-array mode provided the best sensitivity whilst avoiding the non-linear regime of the detector. The short exposures were also useful in the removal of systematics from the data.

3.3 Analysis

3.3.1 Basic Calibrated Data

To analyse these data I used Basic Calibrated Data (BCD) frames produced by the IRAC Level 1 pipeline (version S18.18.0), downloaded from the *Spitzer Heritage*




Figure 3.1: Typical *Spitzer*/IRAC channel 1 ($3.6\ \mu\text{m}$) image of WASP-3. The green source aperture has a radius of radius of 2.5 pixels. The red aperture is the inner radius of the sky annulus - the outer sky boundary was set by the limits of the frame. The purple aperture was used to mask the other star in the image from background estimates.

Archive (SHA).¹ As described in Section 2.1.3, these frames are flux calibrated, with dark and flat-field corrections applied. Typical images from channels 1, 2 and 4 are shown in Figures 3.1, 3.2 and 3.3, respectively. Before obtaining the desired flux measurements from these frames, I applied additional modifications, which are described below for each channel.

3.3.1.1 Channel 1 ($3.6\ \mu\text{m}$) Sub-array Warm Mission Data

IRAC sub-array data are provided in the form of data cubes, each containing $64\ 32 \times 32$ pixel frames. 215 data cubes were obtained from the channel 1 observations. I elected to split each cube into individual images after assessing the autocorrelation of WASP-3's image centroid. The image centroid autocorrelation functions are shown in Figure 3.4 (along with the autocorrelation function for the flux), where two clear features are evident. First, there are long timescale (~ 1 hour) variations corresponding to *Spitzer's* pointing wobble.² This variation is also apparent in the raw photometry for channel 1 (upper left plot of Figure 3.12). Second, sharp peaks

¹<http://irsa.ipac.caltech.edu/data/SPITZER/docs/spitzerdataarchives/>

²ssc.spitzer.caltech.edu/warmmission/news/21oct2010memo.pdf




Figure 3.2: Typical *Spitzer*/IRAC channel 2 ($4.5\ \mu\text{m}$) image of WASP-3. The green source aperture has a radius of 3.0 pixels. The sky annulus (in red) spans a radius from 15-40 pixels. The purple apertures were used to mask the brightest stars in the background region.

in the autocorrelation functions are seen at small offsets. The width of these peaks imply that short timescale excursions of the image centroid are present and typically last for ~ 20 s. Given that the duration of the data cubes were ~ 120 s, I felt that summing the frames (to produce one image per data cube) would not sufficiently sample this short-timescale variation.

I used timing information found in the BCD FITS file headers to determine the Heliocentric Modified Julian Date (HMJD) times for each individual frame at mid-exposure. I converted the times from UTC to the terrestrial timing standard (TT) by adding 66.184 s to the HMJD(UTC) times and also converted from HMJD to HJD. This was done for consistency with the ephemeris used to constrain the parameter fits (see Section 3.3.4).

In order to estimate errors on the flux values I converted pixel values from MJy/sr (as given for the BCD frames) to electron units by multiplying by the gain and effective exposure time, then dividing by a flux conversion factor. These values were all taken from the FITS file headers. An estimate for the zodiacal background flux in the subtracted sky dark was also added, using the SKYDRKZB FITS header value [Anderson et al., 2011].

I found a trend in background values repeated in each data cube, similar to




Figure 3.3: Typical *Spitzer*/IRAC channel 4 ($8.0\ \mu\text{m}$) image of WASP-3. The green source aperture has a radius of 3.0 pixels for this channel. Other apertures are as in Figure 3.2.

that seen in other IRAC data [e.g. Harrington et al., 2007; Agol et al., 2010; Deming et al., 2011; Todorov et al., 2012, 2013] - see Figure 3.5. The reason for this trend is not clear. The 1st and 58th frame within each cube are clearly the largest outliers. The same trend is no longer seen for background subtracted source fluxes, although the 1st frame in each cube still remains an outlier. Because of this I chose to remove the first frame in each data cube for the remainder of the analysis. 215 (1.6%) of the channel 1 frames were removed here.

3.3.1.2 Channel 2 ($4.5\ \mu\text{m}$) Cryogenic Mission Data

Pixel values in channel 2 were converted from MJy/sr to electrons in the same way as described in channel 1. FITS header timing information was again used to determine the HMJD(UTC) times at mid-exposure for each frame. These times were then converted to HJD(TT), with 65.184s being added to convert from the UTC to the TT timing system.

I found that a small number of the channel 2 frames were affected by the ‘column pull-down’ effect.³ This causes a change in the intensity of pixels in the same column as very bright sources such as saturated stars and cosmic ray hits. An example of a frame affected by column pull-down is shown in Figure 3.6. There are

³ssc.spitzer.caltech.edu/warmmission/news/21oct2010memo.pdf.




Figure 3.4: Autocorrelation of the flux and x and y centroid position of WASP-3 in the channel 1 sub-array images. The central 6000 images were used for the comparison time-series and I did not assess the autocorrelation where no shift was applied. The long timescale (~ 1 hour) variations are due to *Spitzer's* pointing wobble, which is also evident in the extracted flux values - see Figure 3.12. The width of the central peak (see inset, zoomed plots) suggests that shorter timescale excursions of the image centroid typically last ~ 20 s.

no saturated stars in the frames, but cosmic ray hits do cause temporary changes to column intensity values. In addition to the ‘pull-down’ I also found a rise in the intensity of column values in frames directly after those affected by cosmic ray hits. I removed frames that showed signs of these effects by creating time series of median column values and rejecting frames where any column was more the 10σ away from the median filtered time series for that column (using a window width of 20 observations). I applied this only to columns coincident with the source aperture, trusting that for the background estimate the clipped mean procedure (see Section 3.3.2) would account for this effect. For the optimal source aperture size of 3 pixels (see Section 3.3.7), I removed 3 frames as a result of this effect.

Post-science blank sky observations were carried out to check for warm pixels




Figure 3.5: *Upper panel:* Channel 1 clipped mean background values for a given sub-array image within each data cube are plotted. For each sub-array image, the average value across the 215 cubes is given here. A clear trend is seen with background values decreasing asymptotically. The 1st and 58th frames stand out as being discrepant from this trend, as has been found in other channel 1 warm mission data.

Lower panel: As above, except here I plot the background subtracted source flux. Here only the 1st frame is an outlier (at 14σ) and so the first frame in each data cube is rejected. The 58th frame is 2.5σ from the median of the 64 values plotted. I elected to use these frames, since the background subtraction adequately accounts for the effect seen in the upper panel.

that could compromise the analysis of WASP-3. I found no evidence for such warm pixels.

3.3.1.3 Channel 4 ($8.0\ \mu\text{m}$) Cryogenic Mission Data

Electron pixel values and mid-exposure HJD(TT) times for channel 4 were calculated as for channel 2. Again, a small number of frames were affected by detector effects - this time by ‘row pull-up’, a result of electronic banding⁴ caused by cosmic ray hits. An example of a frame affected by this effect is shown in Figure 3.7. I removed affected frames in an equivalent way to that described for the ‘column pull-down’

⁴ssc.spitzer.caltech.edu/warmmission/news/21oct2010memo.pdf




Figure 3.6: Typical *Spitzer*/IRAC channel 2 ($4.5\ \mu\text{m}$) image of WASP-3 showing column pull-down. The effect here has been caused by a cosmic ray hit.

effect in channel 2, but this time by creating a time series of median row values. For the optimal source aperture size of 2.5 pixels, 4 frames were removed in this way.

As for channel 2, post-science blank sky observations were carried out. Again, no evidence of warm pixels that would affect the analysis of WASP-3 was found.

3.3.2 Aperture Photometry

3.3.2.1 Source Aperture

Aperture photometry was performed on the remaining frames using the ULTRACAM pipeline software [Dhillon et al., 2007], as described in Section 2.3.1. Centring of the source apertures was carried out using a cross-correlation of a 1D Gaussian profile (FWHM ~ 1.6 pixels) with the marginal sums of the source flux in x and y . A box of 9×9 pixels centred on the brightest pixel was used for the cross-correlation. Flux values from WASP-3 were estimated using a range of fixed source aperture radii, from 2.0–4.0 pixels in steps of 0.5 pixels and additionally 5 and 6 pixel radii. I selected the most suitable aperture size as the one which minimised the normalised residual RMS, once systematic trends and the eclipse model had been removed. The final values adopted were 2.5, 3.0 and 2.5 pixels in channels 1, 2 and 4 respectively (see Section 3.3.7). Binned versions of the extracted fluxes using these source apertures can be seen in the top row Figure 3.12.




Figure 3.7: Typical *Spitzer*/IRAC channel 4 (8.0 μm) image of WASP-3 showing row pull-up, resulting from a cosmic ray hit.

3.3.2.2 Estimating the Background

Background estimates were obtained using a 4σ clipped mean applied to pixel values in a region surrounding the source aperture, on a frame-by-frame basis. Semi-static pixel masks ('p_masks'), provided along with the BCD frames, were applied to each frame and visible stars were also masked from the background estimates. For channel 1, the whole sub-array frame was used, excluding a circular aperture of 6 pixel radius, centred on the source (Figure 3.1). For channels 2 and 4 an annulus of 15–40 pixels was used, with the inner boundary set to avoid significant contributions from the wings of the target PSF (Figures 3.2 and 3.3). The outer radii in these channels were set such that the background estimate was not made too far from WASP-3, but that there were no significant gains (in terms of the photometric errors) in using larger radii.

I tested different methods for estimating the background, including altering the sigma-clipping level (I tried 3, 4 and 5σ) and using the median. The biggest changes compared to my adopted reduction (using a 4σ clipped mean) were seen when using the median, but as shown in Figure 3.8 even these changes are negligible.




Figure 3.8: Figure showing difference in photometry between using my chosen background estimation method (4σ clipped mean) and a median. For comparison the eclipse depths are 150^{+30}_{-20} , 650 ± 30 and 160^{+40}_{-30} electrons for channels 1, 2 and 4 respectively. Offsets in flux will have a negligible effect on the eclipse depth, but changes to these flux differences could impact the results. However, the largest variation, seen in channel 2, is still at only 0.3σ of the final eclipse depth. Hence the differences between different background estimation methods are insignificant.

Table 3.1: Comparison of variances derived from the background regions for a typical image in channels 1, 2 and 4. The second column gives the background noise estimate using the ‘photon’ option in the ULTRACAM pipeline, while the third column gives the estimate from the ‘variance’ option. The latter is significantly larger across the three channels. The final column gives the median variance found in subregions of the background (6×6 pixel boxes in channel 1, 8×8 pixel boxes in channels 2 and 4). In each channel this variance makes up a significant fraction of the total background variance, suggesting that the latter is dominated by pixel-to-pixel variations rather than larger scale background structure.

Channel	σ_{bkg}^2 (photon)	σ_{bkg}^2 (variance)	σ^2 (median from subregions)
1	136	566	435
2	73	233	203
4	606	1007	1068

3.3.2.3 Photometric errors

I estimated uncertainties in the measured flux values using a combination of the Poisson errors in the source counts and the pixel-to-pixel variance in the background region (see Section 2.3.1). I used the variance in the background, rather than the average background count level, because the former was found to be significantly larger.

I found that the background variance was dominated by pixel-to-pixel variations, as opposed to larger scale structure, by comparing the background variance (from the aperture photometry) to the typical variance from subregions of the background. For channel 1 the subregions were boxes of 6×6 pixels placed in clean areas of the background (e.g. without contamination from stars or cosmic ray hits). For channels 2 and 4 I used boxes of 8×8 pixels. The results, given in Table 3.1, show the variances within the subregions dominate the variance of the whole background region.

The underestimation of this pixel-to-pixel background variation when using the photon noise prediction could have resulted from a poorly estimated zodiacal contribution, which was added to the frames as described in Section 3.3.1. The variance option for error estimation is insensitive to this poor calibration, further motivating its use.

3.3.2.4 Centroiding

Centring of the source apertures was important for the analysis, particularly in the removal of systematic effects in channels 1 and 2 (see Section 3.3.3.1). In addition to the standard reduction where Gaussian cross-correlation was used for centring, I also centred source apertures using Gaussian and Moffat profile fitting (see Section 2.3.1). In each case symmetric and non-symmetric profile fits were used. I made flux estimates using the optimal source aperture sizes (2.5, 3.0 and 2.5 pixels in channels 1, 2 and 4 respectively, as selected in Section 3.3.7) and tested the effects of using the different centring methods on the final results in Section 3.3.8.

3.3.2.5 Frame rejection

Frames were rejected in each channel based on flux measurements and centroid positions. Any frame with a flux, or x or y position more than 4σ away from the median of the surrounding 150 (channel 1) or 20 (channels 2 and 4) frames was rejected. The total number of frames rejected in this procedure was 37 (0.3%) in channel 1, 24 (1.5%) in channel 2 and 27 (1.7%) in channel 4. A break down of

Table 3.2: Number of frames rejected as outliers based on flux and position measurements.

Channel	Flux	x -position	y -position	Total
1 (3.6 μm)	22	12	3	37 (0.3%)
2 (4.5 μm)	24	0	0	24 (1.5%)
4 (8.0 μm)	26	1	0	27 (1.7%)

these numbers is given in Table 3.2.

In addition to this, I removed the first 4 data cubes for channel 1 since the x and y positions of WASP-3 showed that the telescope was still settling during this time. I also removed the first 10 minutes of data from channel 4, based on the analysis in Section 3.3.6.

3.3.3 IRAC systematics

3.3.3.1 Channels 1 and 2

Each of the secondary eclipse data sets displayed the systematic effects associated with the IRAC detectors, as described in Section 2.1.4. Both the warm channel 1 and cryogenic channel 2 data are affected by the intra-pixel sensitivity effect. In Figures 3.9 and 3.10 I plot the flux extracted from the aperture photometry against x and y position and also against time. Figure 3.9 shows that for channel 1, there is a clear correlation between the flux and x position, while the correlation with y is not so clear as the PSF crosses a boundary between two pixels in y , where the intra-pixel sensitivity is at a minimum. The trend of flux in time shows the typical $\sim 1\%$ variations on a ~ 1 hr timescale. A clearer representation of this can be found in the top left plot of Figure 3.12, which shows the binned flux against time. Figure 3.10 shows that in channel 2 the correlation between flux and position is not as strong, though there is still a clear trend with the y position. Note that the eclipse signal can be seen in the raw flux against time for channel 2.

To model this effect I tested subsets of the function:

$$F_{\star, \text{model}} = \hat{F}_{\star} + a_0 + a_x dx + a_y dy + a_{xy} dx dy + a_{xx} dx^2 + a_{yy} dy^2 + a_t dt, \quad (3.1)$$

where \hat{F}_{\star} is the weighted mean of the stellar flux and $dx = x - \hat{x}$ and $dy = y - \hat{y}$ are the positions of the PSF centre relative to their weighted means (also see equation 2.1). I used un-smoothed positions to allow sampling of the short timescale jitter




Figure 3.9: A plot of the channel 1 raw flux against dx and dy position and also against time. $dx = x - \hat{x}$ and $dy = y - \hat{y}$ are the positions of the PSF centre relative to their weighted means. There is a clear correlation between the flux and x position, while in y the PSF crosses a pixel boundary, where the minimum of the intra-pixel sensitivity function is sampled. A slight u-shape can be seen in the middle panel however. The trend of flux in time shows the typical $\sim 1\%$ variations on a ~ 1 hr timescale.

seen in Figure 3.4. dt is the time since the first observation. I also tested functions using a logarithmic time term, $a_t \ln(dt + t_{\text{off}})$, in place of the linear time term. The parameter t_{off} was used to ensure the argument of the logarithm did not go to 0.

I select between these models in Section 3.3.6. For the final analysis in Section 3.3.10, the channel 1 data was detrended using the function:

$$F_{\star, \text{model}} = \hat{F}_{\star} + a_0 + a_x dx + a_y dy + a_{xx} dx^2 + a_{yy} dy^2, \quad (3.2)$$

and for channel 2 I used:

$$F_{\star, \text{model}} = \hat{F}_{\star} + a_0 + a_x dx + a_y dy + a_{xx} dx^2 + a_{yy} dy^2 + a_t dt. \quad (3.3)$$

3.3.3.2 Channel 4

The channel 4 data was affected by the ramp effect, as described in Section 2.1.4. Again, this effect can clearly be seen in the raw flux in Figure 3.12. I tried detrending the channel 4 data using the following functions:




Figure 3.10: The equivalent plot to Figure 3.9, but here for the channel 2 data. The correlation between flux and the positions is not as strong as for channel 1, though there is still a clear trend with the y position. Note that the eclipse signal can be seen in the raw flux against time for channel 2.

$$F_{\star,\text{model}} = \hat{F}_{\star} + a_0 + a_t dt + a_{tt} dt^2, \quad (3.4)$$

$$F_{\star,\text{model}} = \hat{F}_{\star} + a_0 + a_t \ln(dt + t_{\text{off}}) + a_{tt} \ln(dt + t_{\text{off}})^2, \quad (3.5)$$

$$F_{\star,\text{model}} = \hat{F}_{\star} + a_0 + a_1 \exp(a_2 dt) + a_3 \exp(a_4 dt), \quad (3.6)$$

where \hat{F}_{\star} , dt and t_{off} are as described previously. I also tried each of these functions without their final term (i.e. $a_{tt} = 0$ for equations 3.4 and 3.5, and $a_3 = 0$ for equation 3.6). Again, I select between these models in Section 3.3.6, where the single exponential function was chosen for my final analysis:

$$F_{\star,\text{model}} = \hat{F}_{\star} + a_0 + a_1 \exp(a_2 dt). \quad (3.7)$$

3.3.4 Markov Chain Monte Carlo Parameter Fitting

Using the MCMCTRANSIT code of Collier Cameron et al. [2007a] as described in Section 2.4, I determined secondary eclipse light curve solutions for the IRAC observations of WASP-3. Briefly, from a set of proposal parameters (see Table 2.1)

chosen at each step in the MCMC, secondary eclipse light curves in each channel were modelled by evaluating the visible fraction of the planetary disk scaled by the day-side planet-to-star flux ratio, for the given channel. The planetary flux contribution was then removed from the measured flux to give the stellar flux, using:

$$F_{\star} = \frac{F_{\text{measured}}}{1 + \eta(\phi)\Delta F} \quad (3.8)$$

The stellar flux, which contained the systematic effects described in Section 3.3.3, was then modelled using subsets of equation 3.1 or equations 3.4-3.6. The detrending function coefficients were obtained at each step in the MCMC using singular value decomposition [Press et al., 1992]. The terms t_{off} , a_2 and a_4 cannot be evaluated using this method and so they were introduced as proposal parameters when they were in use. t_{off} was also subject to a Gaussian prior to prevent its value from drifting to more than an hour.

Constraints that are usually imposed from fitting to the transit light curves were here imposed using priors based on the light curve model parameters ($\frac{R_{\star}}{a}$, $\frac{R_{\text{p}}}{a}$, i) of Southworth [2011] and transit ephemeris (T_0 , P) from Maciejewski et al. [2010]. Southworth [2011] provide a detailed study of WASP-3 light curves, including both high quality ground-based and space-based data. The transit ephemeris from Maciejewski et al. [2010] is based on a transit timing variation study, which naturally sets an accurate ephemeris. Given that these studies had already been carried out on the WASP-3 system, I saw no reason to carry out another analysis of the transit light curves. The use of priors based on the results from these studies still allowed the dependencies of the eclipse parameters on parameters that would be constrained by transit data to be explored. Additional priors on T_{eff} and $[\text{Fe}/\text{H}]$, which were used to estimate the stellar mass, were taken from Pollacco et al. [2008]. In summary, the priors used were:

$$\frac{R_{\star}}{a} = 0.1994 \pm 0.0032$$

$$\frac{R_{\text{p}}}{a} = 0.02125 \pm 0.00041$$

$$i = 83.72 \pm 0.39^{\circ}$$

$$T_0 = 2454605.56000 \pm 0.00011 \text{ HJD(TT)}$$

$$P = 1.8468355 \pm 0.0000007 \text{ d}$$

$$T_{\text{eff}} = 6400 \pm 100 \text{ K}$$

$$[\text{Fe}/\text{H}] = 0.0 \pm 0.2$$

I note that more recent ephemerides have been published [e.g. Sada et al., 2012; Nascimbeni et al., 2013]. However, the uncertainties on the mid-eclipse phases I have found are dominated by the data and the differences between ephemerides have a negligible effect (within 0.1σ).

I fitted to the out-of-transit radial velocity data from Simpson et al. [2010] and Tripathi et al. [2010], since this puts constraints on the eccentricity and argument of periastron (ω) of the system. As described in Section 2.4.2.3, the timing of the secondary eclipse also constrains these parameters through the relation [Charbonneau et al., 2005]:

$$\Delta\phi \equiv \phi_E - 0.5 \simeq \frac{2}{\pi}e \cos \omega, \quad (3.9)$$

where ϕ_E is the observed mid-eclipse phase and 0.5 is the expectation from a circular orbit.

In the following sections, the error bars on the IRAC photometry were scaled so that the reduced χ^2 in each channel was 1, unless otherwise stated. Typical scaling values were: 1.06 in channel 1; 1.01 in channel 2 and 1.06 in channel 4, implying fits to the IRAC data were good, even without the rescaling. Additionally, the errors on the radial velocity data were scaled so that the reduced χ^2 on individual sets was also 1. This was done to ensure that realistic uncertainties on the proposal parameters were obtained from the MCMC.

In the following sections, I select the optimised detrending functions and source aperture radius in each channel, amongst other explorations of the data (Sections 3.3.5–3.3.9). In all of these tests the MCMC code was run using 10^4 jumps, with optimal parameter estimates taken as the median of the resulting marginal distributions for each parameter, and the uncertainties taken as the 15.9 and 84.1 percentiles.

3.3.5 Light travel time corrections

I accounted for the light travel time across the WASP-3 system for each of the data sets. For a circular orbit, this effect causes the secondary eclipse to appear later than the $\phi = 0.5$ prediction by $\frac{2a}{cP}$. Here a is the radius of the orbit and c is speed of light. This delay occurs because the epoch of the orbit is given for mid-transit. Similarly, $\phi = 0$ for the radial velocity measurements will occur $\frac{a}{cP}$ later than mid-transit. This can be understood by considering the radial velocity signature that would be measured for the stellar light blocked by the planet at mid-transit. It would have the radial velocity signature of the star at a time $\frac{a}{c}$ before mid-transit.

To account for these effects I subtracted 31.4s from the secondary eclipse timings and 15.7s from the radial velocity timings. These values were estimated using a value for a from an initial MCMC run (without these corrections) and assuming a circular orbit. I found that my value for a was not affected by the correction, and that the light-travel time correction did not change significantly for eccentric orbits allowed by the final results.

I did not correct for the fact that some times were given as Heliocentric Julian Dates (HJD), while others were given as Barycentric Julian Dates (BJD). The differences here are typically ~ 1 s, much smaller than the final errors on the central eclipse time, and so will not have significantly impacted the results.

3.3.6 Detrending model selection

I assessed the relative quality of the fits given by the detrending functions described in Sections 3.3.3.1 and 3.3.3.2 using the Bayesian Information Criterion [BIC, Schwarz, 1978], as described in Section 2.5.

Initially I performed MCMC analyses on individual *Spitzer* datasets to find the best 2 or 3 detrending functions in each case (i.e. with the lowest BIC value). For these tests I used datasets where the source flux had been extracted using aperture sizes of 2.5, 3.0 and 2.5 pixels in channels 1, 2 and 4 respectively. The results of these tests are given in Tables 3.3, 3.4 and 3.5. Note that in all of these comparisons of detrending functions, secondary eclipse errors were not rescaled to ensure BIC values were not affected by differences in rescaling between different models.

The best detrending functions, highlighted in these tables, were then tested in various combinations, using combined MCMC fits to all the IRAC datasets. Combined fits were used because I wanted to ensure parameters that were simultaneously constrained by all the datasets (e.g. $\Delta\phi$) were not affecting my conclusions about the detrending functions. In addition, motivated by systematic features in the channel 1 residuals (see Section 3.3.9) that could have been biasing parameters like $\Delta\phi$, I tested combinations of detrending functions using combined MCMC fits to the channel 2 and 4 data only. My final choice for the channel 1 detrending function came from the combined fits to all the IRAC data, while the final choices for channels 2 and 4 came from the combined fits to the channel 2 and 4 data.

For channel 1 I found the best detrending function from the simultaneous fits to all IRAC data was:

$$F_{\star,\text{model}} = \hat{F}_{\star} + a_0 + a_x dx + a_y dy + a_{xx} dx^2 + a_{yy} dy^2. \quad (3.10)$$

Table 3.3: Eclipse depths and BIC values for fit to the channel 1 data only, using a variety of detrending models. Fits including quadratic terms are clearly favoured by the BIC and the eclipse depths derived from these fits are consistent with the final result of $0.209^{+0.040}_{-0.028}\%$. The best two models from these individual fits were used in the subsequent tests of simultaneous fits to data in all channels. The worst three detrending models given here show markedly poorer BIC values. This is a result of the lack of spatial detrending terms for these models.

Channel 1 detrending model ($-\hat{F}_*$)	Eclipse depth (%)	BIC = $\chi^2 + k \ln N$
$a_0 + a_x dx + a_y dy + a_{xx} dx^2 + a_{yy} dy^2$	0.210 ± 0.009	14969.7
$a_0 + a_x dx + a_y dy + a_{xx} dx^2 + a_{yy} dy^2 + a_t dt$	0.212 ± 0.009	14977.1
$a_0 + a_x dx + a_y dy + a_{xy} dx dy + a_{xx} dx^2 + a_{yy} dy^2$	0.212 ± 0.009	14978.5
$a_0 + a_x dx + a_y dy + a_{xx} dx^2 + a_{yy} dy^2 + a_t \ln(dt + t_{\text{off}})$	0.212 ± 0.009	14984.9
$a_0 + a_x dx + a_y dy$	0.244 ± 0.008	16012.6
$a_0 + a_x dx + a_y dy + a_t dt$	0.246 ± 0.008	16028.2
$a_0 + a_x dx + a_y dy + a_t \ln(dt + t_{\text{off}})$	0.242 ± 0.008	16040.9
$a_0 + a_t \ln(dt + t_{\text{off}})$	0.179 ± 0.008	24218.0
$a_0 + a_t dt$	0.131 ± 0.008	24331.1
a_0	0.072 ± 0.008	25899.3

Table 3.4: As Table 3.3, but for channel 2. Again detrending models with quadratic terms are favoured. The best three models here have eclipse depths consistent with the final result of $0.282 \pm 0.012\%$. As in channel 1, the best two models from these individual fits were used in the subsequent tests of simultaneous fits to data in channels 2 and 4. The models with no spatial detrending terms are not as markedly poor here as they were for channel 1 because the intra-pixel sensitivity effect is not as strong in channel 2.

Channel 2 detrending model ($-\hat{F}_*$)	Eclipse depth (%)	BIC = $\chi^2 + k \ln N$
$a_0 + a_x dx + a_y dy + a_{xx} dx^2 + a_{yy} dy^2 + a_t dt$	0.282 ± 0.012	1637.7
$a_0 + a_x dx + a_y dy + a_{xx} dx^2 + a_{yy} dy^2$	0.279 ± 0.012	1639.1
$a_0 + a_x dx + a_y dy + a_{xy} dx dy + a_{xx} dx^2 + a_{yy} dy^2 + a_t dt$	0.281 ± 0.012	1645.0
$a_0 + a_x dx + a_y dy + a_{xx} dx^2 + a_{yy} dy^2 + a_t \ln(dt + t_{\text{off}})$	0.271 ± 0.012	1647.4
$a_0 + a_x dx + a_y dy$	0.286 ± 0.012	1661.5
$a_0 + a_x dx + a_y dy + a_t dt$	0.289 ± 0.012	1667.5
$a_0 + a_x dx + a_y dy + a_t \ln(dt + t_{\text{off}})$	0.279 ± 0.012	1677.2
$a_0 + a_t dt$	0.277 ± 0.011	1815.9
$a_0 + a_t \ln(dt + t_{\text{off}})$	0.254 ± 0.011	1827.8
a_0	0.266 ± 0.012	1860.9

Table 3.5: As Tables 3.3 and 3.4, but for channel 4 data. The best three detrending functions were chosen for the subsequent tests of simultaneous fits to data in all channels. The difference in eclipse depth for these three detrending functions is discussed in the text. The very poor performance of the worst model in this table highlights the necessity of including temporal terms in the detrending model.

Channel 4 detrending model ($-\hat{F}_*$)	Eclipse depth (%)	BIC = $\chi^2 + k \ln N$
$a_0 + a_1 \exp(a_2 dt)$	0.371 ± 0.046	1794.2
$a_0 + a_t \ln(dt + t_{\text{off}})$	0.332 ± 0.034	1795.4
$a_0 + a_1 \exp(a_2 dt) + a_3 \exp(a_4 dt)$	0.339 ± 0.050	1798.5
$a_0 + a_t dt + a_{tt} dt^2$	0.328 ± 0.048	1801.4
$a_0 + a_x dx + a_y dy + a_1 \exp(a_2 dt)$	0.366 ± 0.047	1803.2
$a_0 + a_t \ln(dt + t_{\text{off}}) + a_{tt} \ln(dt + t_{\text{off}})^2$	0.316 ± 0.035	1807.9
$a_0 + a_t dt$	0.107 ± 0.030	1827.3
a_0	0.827 ± 0.028	2119.2

Table 3.6: Eclipse depths and BIC values for channel 4, from the simultaneous fits to the channel 2 and 4 data. The first 10 minutes of channel 4 data were removed for these fits as these were the source of the inconsistencies in the channel 4 eclipse depths seen in Table 3.5. The channel 2 data was detrended using equation 3.3, while the detrending for channel 4 is shown in this table. The eclipse depths found were all consistent with each other, with the single exponential function being favoured by the BIC, as in the individual fits.

Detrending model ($-\hat{F}_\star$)	Eclipse depth (%)	BIC = $\chi^2 + k \ln N$
$a_0 + a_1 \exp(a_2 dt)$	0.329 ± 0.052	3329.2
$a_0 + a_t \ln(dt + t_{\text{off}})$	0.333 ± 0.032	3338.5
$a_0 + a_1 \exp(a_2 dt) + a_3 \exp(a_4 dt)$	0.326 ± 0.053	3343.6

This is the same as was found for the individual fits. Table 3.3 shows that for the individual fits, there was little change in the eclipse depth for the best fitting detrending functions. This was also the case for the simultaneous fits. Equation 3.10 was therefore chosen as the detrending function for the final analysis of the channel 1 data.

In channel 2 the results followed a similar pattern. The best detrending function from the simultaneous fits (both including and excluding the channel 1 data) was:

$$F_{\star, \text{model}} = \hat{F}_\star + a_0 + a_x dx + a_y dy + a_{xx} dx^2 + a_{yy} dy^2 + a_t dt. \quad (3.11)$$

as was found for the individual fits. Again there was little variation in the eclipse depths found from both the individual (see Table 3.4) and simultaneous fits. Equation 3.11 was therefore chosen as the detrending function for the final analysis of the channel 2 data.

In channel 4, as in channel 2, there was consistency between the best fit to both the individual and combined datasets. The best detrending function was found to be the single exponential in time:

$$F_{\star, \text{model}} = \hat{F}_\star + a_0 + a_1 \exp(a_2 dt). \quad (3.12)$$

As shown in Table 3.5, this function was only marginally favoured over the linear logarithmic time (equation 3.5 with $a_{tt} = 0$) and double exponential (equation 3.6) functions. In addition, the eclipse depth from the single exponential function only agrees with the other two at 1σ . I found this difference stemmed from data at the beginning of the time-series, where the single exponential function did not

give a good fit. With the first 10 minutes of data removed, these three models gave consistent channel 4 eclipse depths, with the single exponential again being favoured by the BIC (see Table 3.6). As a result, for the final analysis in Section 3.3.10 the channel 4 data was detrended using the single exponential function with the first 10 minutes of data removed.

3.3.7 Source aperture size

With the detrending functions chosen, I then carried out combined MCMC analyses on light curves extracted using source aperture radii between 2.0–6.0 pixels, as described in Section 3.3.2. The optimal radii were found by assessing where the RMS of the normalised residuals from the fitted light curve were minimised. Values of 2.5, 3.0 and 2.5 pixels were found for channels 1, 2 and 4 respectively. Figure 3.11 highlights these choices. I plot the signal-to-noise ratios of flux measurements from a typical image, the RMS of the normalised residuals and the eclipse depths as a function of source aperture radius for the three channels. Minima in the RMS values coincide well with maxima in the typical signal-to-noise values, as expected. In all channels the eclipse depths remain well within the MCMC uncertainties, except for the larger radii in channel 1, which differ from those at smaller radii at the 2σ level. However, there is a significant increase in the RMS of the normalised residuals, along with an associated drop in typical signal-to-noise, at larger radii due to the inclusion of more background photons.

For channels 1 and 4, Figure 3.11 shows a discrepancy between the estimated and actual RMS values at the chosen radii. The trends in these discrepancies with radius suggest there is a noise floor at smaller radii, where some additional source of noise (not accounted for in the flux error estimation) becomes significant.

Figure 3.12 shows binned time-series for the three IRAC channels, using the chosen source apertures. The top row shows the raw flux light curves with the full secondary eclipse model, including the detrending functions chosen in Section 3.3.6. The middle row shows the flux with the detrending function removed and normalised to the flux of the star, along with the eclipse model. The model residuals are shown in the bottom row.

3.3.8 Centroiding method

In Section 3.3.2 I extracted flux from each of the *Spitzer* datasets using centroiding methods different to the cross-correlation technique used in the standard reduction. I used Gaussian and Moffat profile fitting and in each case symmetric and non-




Figure 3.11: Columns from left to right are for channels 1, 2 and 4 respectively. The top row shows signal-to-noise ratios of flux measurements from a typical image of WASP-3, at a range of source aperture radii. The middle row gives normalised residual RMS values on the light curves detrended using the functions found in Section 3.3.6. Also shown is the expectation from the signal-to-noise values (blue line). A discrepancy between these values, for a given source aperture radius, suggests the presence of additional noise sources. The bottom row gives the eclipse depth values, with MCMC errors. The grey dashed lines show the final quoted error for the optimal eclipse depth, taken as either the error from the MCMC or the prayer bead analysis, whichever was largest. The vertical dotted lines show the adopted radii, chosen to be where the normalised residual RMS value was minimised.

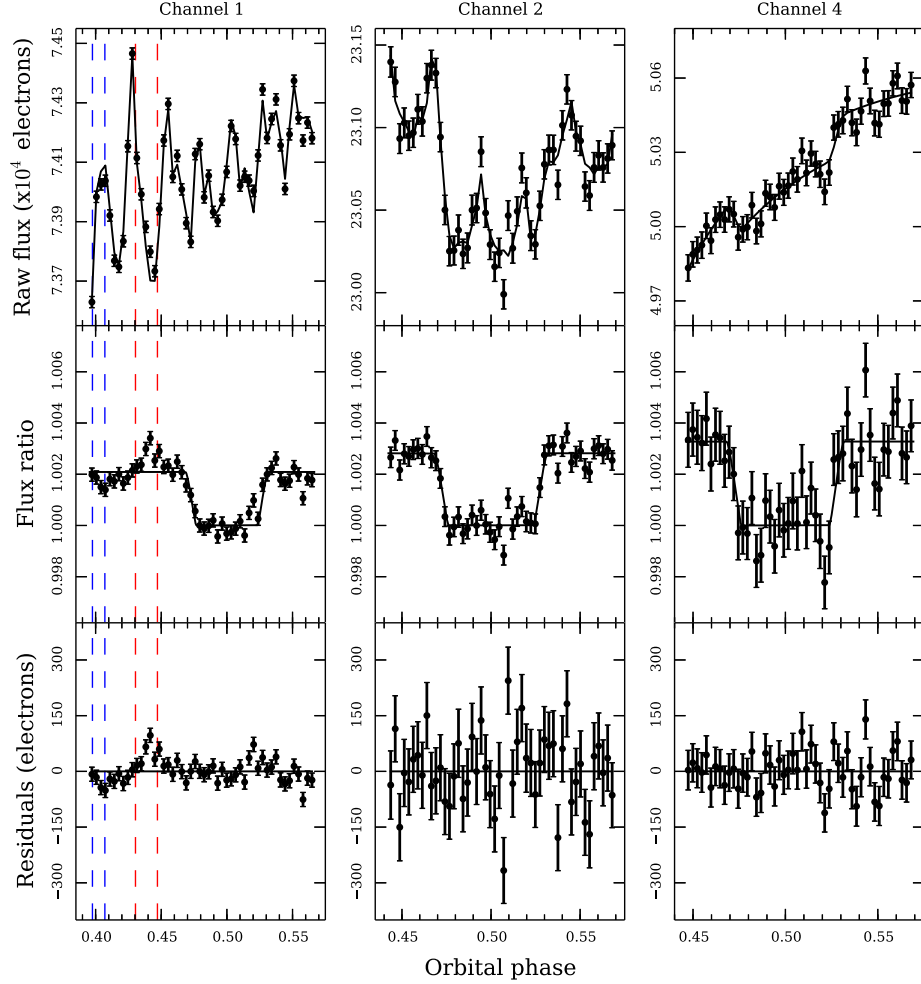


Figure 3.12: Columns from left to right are for channels 1, 2 and 4 respectively. The top row shows the (binned) raw flux light curves for the optimal source aperture radii along with the full secondary eclipse model, which includes the detrending functions found in Section 3.3.6. The middle row shows the fluxes with the detrending functions removed and normalised to the flux of the star. Residuals are shown in the bottom row. The IRAC systematic effects are clearly seen in the raw data. These effects are generally well corrected for using the detrending functions given in equations 3.10, 3.11; and 3.12, however there appear to be remaining features in the channel 1 data, which are discussed in the text. The eclipse depths, as displayed in the middle row, are 0.209%, 0.282% and 0.328% in the three channels respectively. The blue and red vertical lines for the channel 1 data correspond to the phase ranges of the positions highlighted in the same colours in the centroid position plot (Figure 3.14). See Section 3.3.9 for a discussion of this.

Table 3.7: Eclipse depths and normalised residual RMS values for different centroiding methods. The channel 1 results were obtained from simultaneous fits to all the data, while the channel 2 and 4 results were obtained from simultaneous fits to the channel 2 and 4 data only. The bottom row gives the final results for comparison. The biggest discrepancy in eclipse depth is $< 0.5\sigma$ in channel 2 and other differences are much smaller than this. There is also very little difference between RMS values. The cross-correlation technique used in my reduction gives consistently good RMS values across the 3 channels, so there is no motivation to use an alternate centroiding method.

Centroiding method	Channel 1		Channel 2		Channel 4	
	ED (%)	RMS ($\times 10^{-3}$)	ED (%)	RMS ($\times 10^{-3}$)	ED (%)	RMS ($\times 10^{-3}$)
Gaussian	0.210	4.217	0.282	2.134	0.335	5.622
Gaussian (non sym.)	0.215	4.221	0.278	2.143	0.334	5.628
Moffat	0.213	4.189	0.277	2.137	0.332	5.651
Moffat (non sym.)	0.208	4.264	0.277	2.136	0.334	5.664
Cross-correlation	$0.209^{+0.040}_{-0.028}$	4.190	0.282 ± 0.012	2.136	$0.328^{+0.086}_{-0.055}$	5.623

symmetric profile fits were used. The optimal source aperture sizes found in Section 3.3.7 were used here.

For each centroiding method, two MCMC analyses was carried out as described in Section 3.3.4: one using data from all 3 channels and another using data from channels 2 and 4. The results for channel 1 were taken from the first analysis, while the results for channels 2 and 4 were taken from the second analysis. The resulting eclipse depths and normalised residual RMS values are given in Table 3.7. These results show little variation in eclipse depths with centroiding method - the biggest variation is at $< 0.5\sigma$ and is typically much smaller than this. The RMS values also show very little variation, with the largest fractional difference in RMS at $< 2\%$.

Across the three *Spitzer* bands the cross-correlation centroiding method had consistently good RMS values and eclipse depths that were in agreement with the alternative centroiding methods. Therefore, there was no motivation to alter this aspect of the reduction.

3.3.9 Channel 1 residual systematics

Figure 3.12 shows that for channel 1, systematic features remain in the residuals after modelling using equation 3.10. A rise and fall of around 0.15% can be seen from $\phi = 0.40$ -0.44 and another, smaller bump is present around $\phi = 0.52$. I looked into a number of potential sources of these features, which are discussed below.

3.3.9.1 Detrending function

Given that alternative detrending functions provided reasonably good fits to the channel 1 data (see Table 3.3), I tested whether the residual features seen in Figure 3.12 were suppressed for any of these. Visual inspection of the residuals for the best four detrending functions in Table 3.3 (which have reasonable BIC values) indicated that the residual systematic features were still present and with similar amplitudes.

3.3.9.2 Effect of the source aperture radius

Given the variation of eclipse depth with source aperture radius in channel 1 (see Figure 3.11), I considered the possibility that the residual systematic noise could be reduced with larger source aperture radii. In Figure 3.13 I show normalised residual RMS values as a function of bin width for different radii in each of the three channels used. In channel 1, red noise is clearly seen as the RMS values increase with respect to the white noise expectation for large bin sizes. Crucially this feature is present at large radii ($r = 5.0 - 6.0$ pixels) as well as at the chosen radius of 2.5 pixels. Thus there is no motivation for using the larger aperture.

3.3.9.3 Centroid position

At the beginning of the channel 1 observations there was a large movement of the image centroid in y , covering a distance of ~ 0.6 pixels in ~ 4 hours (Figure 3.14). This was where the most clear residual systematic was also present, so I suspect this motion played a significant role. However, as shown in Figures 3.12 and 3.14, the extremes of a slope in the residuals ($\phi = 0.40 - 0.44$) occurred while the centroid positions covered very similar parts of the pixel. This shows that a simple detrending function in space and time cannot correct these features.

3.3.9.4 Centroiding method

Using the reductions for the different centroiding methods (see Section 3.3.8), I tested if my use of the cross-correlation technique was the source of the noise. Visual inspection of the channel 1 light curves for the various centroiding methods revealed that the general nature of the residual features remained (e.g. there was always a slope in the residuals at the beginning of the time series) and that changes to their amplitudes were minimal.

Figure 3.15 is an equivalent plot to the top panel of Figure 3.13, except here I plot for different centroiding methods (rather than source aperture radii). It shows that while there may be a slight reduction in the red noise for different centroiding




Figure 3.13: Binned normalised residual RMS values as a function of bin width. From top to bottom, plots are for channels 1, 2 and 4. The different line colours represent different source aperture radii. Black lines are for the optimal radii (2.5, 3.0 and 2.5 pixels respectively), while the red, blue and green lines are for 4.0, 5.0 and 6.0 pixel radii respectively. The dotted lines show the $1/\sqrt{n}$ expectation, where n is the number of data points per bin. This expectation is fixed to the unbinned RMS value for the optimal source aperture radius. Red noise is clearly present in the channel 1 residuals across the range of aperture sizes and does not decrease significantly at large radii. Since the signal-to-noise and normalised residual RMS values degrade at these larger radii, there is no motivation in using a larger aperture. Channels 2 and 4 show well behaved trends, following the $1/\sqrt{n}$ expectation.

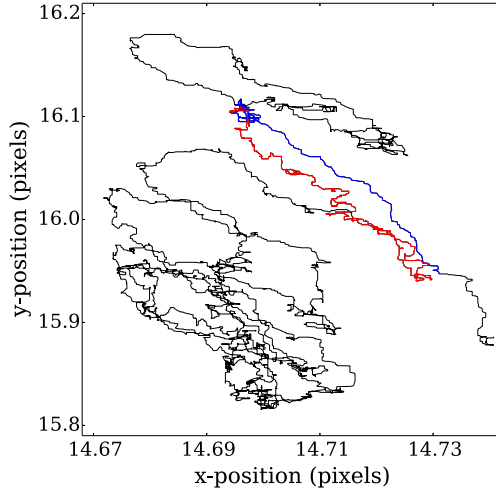


Figure 3.14: Median filtered centroid positions for the channel 1 data set (with a filter width of 200 frames). At the start of the observations the centroid position is $x \sim 14.74$, $y \sim 15.9$. Note the significantly larger scale in y compared to x . Throughout the first ~ 4 hours of observations, the image moves across a large range in y , before settling down for the final ~ 3 hours. The blue and red highlighted regions correspond to those marked on Figure 3.12, and are discussed in Section 3.3.9.3.

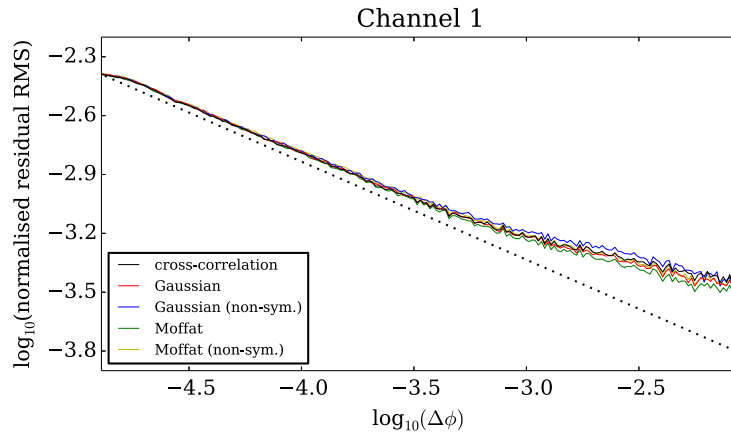


Figure 3.15: An equivalent plot to the top panel of 3.13, except here I vary the centroiding method rather than the source aperture size. This plot shows that while the red noise may be slightly suppressed when using some alternate centroiding techniques, it is still clearly present and so my choice of centroiding technique is not the root of the noise.




Figure 3.16: Channel 1 flux residuals plotted against the FWHM (using Gaussian profile fitting centroiding method). The residual and FWHM values were binned (in time) at the same level as in Figure 3.12 before being plotted. There is no clear trend in the residuals with FWHM, suggesting that variations in the FWHM are not causing the residual systematic variations in channel 1.

methods, it is not a significant change. The red noise is certainly still present when using the different centroiding methods and therefore I have no reason to alter this aspect of the reduction.

3.3.9.5 Trend with FWHM

Another potential source of the residual systematic trends could be variations in the FWHM of the image. The flux measurements could be susceptible to this due to my use of a fixed aperture size (rather than one which varies with FWHM). To test this I looked for trends between the channel 1 residuals and the measured FWHM. This was done for the case of Gaussian profile fitting, since the FWHM is not measured in the cross-correlation centroiding. In Fig 3.16 I show the light curve residuals plotted against the FWHM, binned to the same level as Figure 3.12. No obvious trends are seen here, so it is unlikely that variations in the FWHM are causing the systematic features.

3.3.9.6 Stellar activity

The systematic features seen in the channel 1 residuals could be an astrophysical effect, rather than being detector related. One possibility is variability in the stellar

flux due to flares. As a simple test of this I looked at Kepler light curves of F-type stars with similar spin periods to WASP-3 [Mathur et al., 2013] and found photometric variability with amplitudes and timescales comparable to the channel 1 residual systematic features. It is therefore plausible that they result from stellar variability rather than a detector related systematic.

3.3.10 Final analysis and prayer bead errors

With the detrending functions and source aperture sizes chosen, I ran two final MCMC analyses: one including all the data and another including only the channel 2 and 4 data. The former was used to set the results for the channel 1 eclipse depth, while results for all other parameters were found from the latter. The motivation here was to derive the main set of results from the two well behaved data sets (channels 2 and 4).

These were run using 10^5 production run steps, to ensure a thorough exploration of the parameter space. The optimal parameter values were taken as the medians of the resulting marginal distributions for each parameter, which can be found in Tables 3.8 and 3.9.

To estimate the extent to which red noise affects the eclipse depth and timing measurements, I performed a ‘prayer bead’ analysis [Gillon et al., 2007]. Light curve residuals were subtracted from the raw data and then added back in with a cyclic offset in phase. I used 20 different offsets, equally spaced across the phases sampled, and I applied these shifts to each of the light curves. The resulting data sets were run through the MCMC procedure as described previously, though using 10^4 jumps here. From the resulting distribution I assessed the 1σ errors as the 68% confidence interval centred on the value obtained from the original MCMC fit. The final uncertainty values in Tables 3.8 and 3.9 were taken as the larger of either this error or that given by the MCMC analysis on the unmodified data.

The prayer bead analysis assumes that the residual red noise pattern was just as likely to occur at different phases relative to the time series, and finds the resulting variations in parameter values. For example, it assumes that the slope seen in the residuals of channel 1 at the beginning of the observations could have occurred at any time during the observations. Given that the slope occurred when the PSF centre was undergoing quite a large change in position (see Figure 3.14), this assumption is probably too conservative. However, given that there was no simple correlation between the residual pattern and position, this analysis was used to give a fair reflection of the uncertainty.

The only parameters for which prayer bead errors were larger than the




Figure 3.17: Eclipse depth values, with MCMC errors, for the fake datasets created in the prayer bead analysis by adding residuals back into the best fitting model, with a given offset. The dotted horizontal lines give the final adopted errors. The top 3 panels are for the eclipse depths in channels 1, 2 and 4 respectively, and the bottom panel is for $e \cos \omega$. Channel 1 is clearly the most affected by these shifts, as evidenced by the correlated noise signatures seen in Figures 3.12 and 3.13. I adopt the prayer bead errors here since it is clearly more representative than the errors derived from the MCMC code. The values for the channels 2 and 4 eclipse depths were derived from the second prayer bead analysis (which excluded the channel 1 data). The residuals here are much cleaner and the variation in eclipse depths is comparable to the MCMC error bars. For $e \cos \omega$ I plot values from both prayer bead analyses (those from the second analysis are shown in black), showing an improvement in the distribution when the channel 1 data are excluded. However, the positive error bar is still significantly affected by the residual permutations.

MCMC errors were the channel 1 and 4 eclipse depths and the parameters associated with the eclipse timings ($\sqrt{e} \cos \omega$, $e \cos \omega$, ϕ_E). Figure 3.17 shows the results of the prayer bead MCMC runs for the eclipse depths and for $e \cos \omega$. As expected, the channel 1 eclipse depth was the most affected, with its errors increasing by a factor of 3–4 above the MCMC errors.

Since the eclipse timing ($e \cos \omega$) is constrained by each channel simultaneously, I felt it appropriate to re-run the prayer-bead analysis with the channel 1 data removed. This was done because systematics associated with that channel could have an adverse effect on the determinations of $e \cos \omega$, as well as the channel 2 and 4 eclipse depths. The resulting distribution for $e \cos \omega$ did show an improvement, though the positive error was still larger than the MCMC error. There was little change to the channel 2 and 4 eclipse depth distributions after dropping the channel 1 data. The channel 2 prayer bead variation was smaller than the MCMC errors, and for channel 4 only the positive error was set by the prayer bead analysis.

3.4 Results

The results from my MCMC analyses of the *Spitzer* IRAC secondary eclipse data for WASP-3b can be found in Tables 3.8 and 3.9. They are given for the optimal source aperture radii, with errors accounting for the results of the prayer bead analysis. The first column of results are for the fits to the channel 1, 2 and 4 data, while the second column is for the fits excluding the channel 1 data. I adopt the parameters in the second column as the final set. The exceptions to this are the eclipse depth and brightness temperature for channel 1, which are taken from the first column.

I find eclipse depths of $0.209^{+0.040}_{-0.028}$ %, 0.282 ± 0.012 % and $0.328^{+0.086}_{-0.055}$ % in channels 1, 2 and 4 respectively. These eclipse depths are equivalent to day-side planet-to-star flux ratios. Using simple blackbody estimates I expect the planetary thermal emission to dominate the planetary reflected light across these wavelengths. For example, using a very conservative value of 0.5 for the Bond albedo, the ratio of thermally emitted photons to reflected photons from the planet is 20–40 across the IRAC channels. I therefore interpret the secondary eclipse measurements as resulting from the thermal emission of the planet.

I used the measured eclipse depths to estimate day-side brightness temperatures for each channel, using the technique described in Section 2.6.1. I find values of $T_{3.6\mu\text{m}} = 2280^{+210}_{-150}$ K, $T_{4.5\mu\text{m}} = 2400 \pm 80$ K and $T_{8.0\mu\text{m}} = 2210^{+390}_{-250}$ K.

The constraint on $e \cos \omega$ from the eclipse timings and radial velocity data is $-0.0006^{+0.0010}_{-0.0006}$, with the positive error resulting from the prayer bead analysis and

the negative error from the MCMC. This result suggests the timing of the eclipse is not significantly different from the expectation for a circular orbit. This implies the eccentricity of the system can only be large if $\omega \simeq 90^\circ$. For example, for $e > 0.03$, ω is restricted to $87.5^\circ < |\omega| < 94.2^\circ$ at 3σ . A measurement of $e \cos \omega = 0.0070 \pm 0.0032$ by Zhao et al. [2012] gives a marginal 2σ agreement with my result.

Since the prayer bead analysis does not affect the eccentricity estimates, I use the MCMC distribution for the constraints, which gives a result of $e = 0.003_{-0.002}^{+0.013}$. The 3σ upper limit on the eccentricity is 0.049.

Table 3.8: Fitted parameters of the WASP-3 system derived from the MCMC analysis of *Spitzer* secondary eclipse data (channels 1, 2 and 4), using priors to constrain many properties of the system. Epoch and period priors are from Maciejewski et al. [2010], T_{eff} and $[\text{Fe}/\text{H}]$ are from Pollacco et al. [2008], and others are from Southworth [2011]. Parameter values in the column ‘Channels 1, 2 & 4’ are from the simultaneous fit to all the *Spitzer* data, whereas those in the column ‘Channels 2 & 4’ are from the fit to these two channels (see Section 3.3.10).

Parameter	Symbol	Prior	Channels 1, 2 & 4	Channels 2 & 4 ^a	Unit
Fitted parameters:					
Mid-transit time ^b	T_0	$2454605.56000 \pm 0.00011$	$2454640.64993 \pm 0.00011$	$2454640.64993 \pm 0.00011$	d
Orbital period	P	1.8468355 ± 0.0000007	1.8468355 ± 0.0000007	1.8468355 ± 0.0000007	d
Stellar RV amplitude	K_1	6400 ± 100	0.284 ± 0.008	0.284 ± 0.008	km s^{-1}
Stellar temperature	T_{eff}	0.0 ± 0.2	6400 ± 100	6400 ± 100	K
Stellar metallicity	$[\text{Fe}/\text{H}]$		0.0 ± 0.2	0.0 ± 0.2	
Planet/star area ratio	δ		0.0112 ± 0.0005	0.0111 ± 0.0005	
Primary transit duration	T_{14}		$0.1143_{-0.0016}^{+0.0013}$	$0.1155_{-0.0019}^{+0.0016}$	
Impact parameter	b		0.527 ± 0.019	0.511 ± 0.020	
	$\sqrt{e} \cos \omega$		$-0.009_{-0.025}^{+0.022}$	$-0.010_{-0.017}^{+0.017}$	
	$\sqrt{e} \sin \omega$		$0.019_{-0.006}^{+0.005}$	$0.021_{-0.065}^{+0.033}$	
Channel 1 (3.6 μm) eclipse depth ^c	$\Delta F_{3.6\mu\text{m}}$		0.282 ± 0.012	0.282 ± 0.012	%
Channel 2 (4.5 μm) eclipse depth	$\Delta F_{4.5\mu\text{m}}$		$0.323_{-0.081}^{+0.081}$	$0.328_{-0.086}^{+0.086}$	%
Channel 4 (8.0 μm) eclipse depth ^c	$\Delta F_{8.0\mu\text{m}}$		$0.323_{-0.054}^{+0.081}$	$0.328_{-0.055}^{+0.086}$	%

^aFinal values are taken from this column, except for the eclipse depth for channel 1.

^bHJD (TT).

^cFor these parameters the prayer bead analysis led to an increase in the error estimates.

Table 3.9: As Table 3.8, but for parameters derived from the fitted MCMC proposal parameters.

Parameter	Symbol	Prior	Channels 1, 2 & 4	Channels 2 & 4 ^a	Unit
Derived parameters:					
Orbital separation	a		0.0315 ± 0.0003	0.0315 ± 0.0003	AU
Orbital inclination	i	83.72 ± 0.39	83.91 ± 0.28	84.06 ± 0.29	$^{\circ}$
b	$e \cos \omega$		$-0.0005^{+0.0017}_{-0.0005}$	$-0.0006^{+0.0010}_{-0.0006}$	
	$e \sin \omega$		$0.001^{+0.011}_{-0.002}$	$0.001^{+0.014}_{-0.003}$	
Orbital eccentricity	e		$0.003^{+0.011}_{-0.002}$	$0.003^{+0.013}_{-0.002}$	
Mid-eclipse phase b	ϕ_E		$< 0.045 (3\sigma)$	$< 0.049 (3\sigma)$	d
Secondary eclipse duration	T_{58}		$0.4997^{+0.0011}_{-0.0003}$	$0.4996^{+0.0007}_{-0.0004}$	M_{\odot}
Stellar mass	M_{\star}		0.1147 ± 0.0011	0.1161 ± 0.0013	R_{\odot}
Stellar radius	R_{\star}		1.23 ± 0.03	1.23 ± 0.03	
Scaled stellar radius	R_{\star}/a	0.1994 ± 0.0032	1.36 ± 0.02	1.36 ± 0.02	
Stellar density	ρ_{\star}		0.2004 ± 0.0025	0.2011 ± 0.0026	ρ_{\odot}
Stellar surface gravity	$\log g_{\star}$		$0.488^{+0.018}_{-0.017}$	$0.482^{+0.019}_{-0.018}$	(egs)
Planet mass	M_p		4.260 ± 0.011	4.256 ± 0.011	M_J
Planet radius	R_p		1.98 ± 0.07	1.98 ± 0.06	R_J
Scaled planetary radius	R_p/a	0.02125 ± 0.00041	1.40 ± 0.03	1.40 ± 0.03	
Planet density	ρ_p		0.02122 ± 0.00042	0.02120 ± 0.00042	ρ_J
Planet surface gravity	$\log g_p$		0.72 ± 0.05	0.72 ± 0.05	(egs)
Channel 1 brightness temperature	$T_{3.6\mu m}$		3.36 ± 0.02	3.36 ± 0.02	K
Channel 2 brightness temperature	$T_{4.5\mu m}$		2280^{+210}_{-150}	2400 ± 80	K
Channel 4 brightness temperature	$T_{8.0\mu m}$		2390 ± 80	2210^{+390}_{-250}	K

^aFinal values are taken from this column, except for the brightness temperature for channel 1.

^bFor these parameters the prayer bead analysis led to an increase in the error estimates.

3.5 Discussion

3.5.1 Albedo and energy redistribution

Two key properties in our understanding of exoplanetary atmospheres are the fraction of the incident stellar radiation that is reflected by the planet (the Bond albedo) and the efficiency with which the planet redistributes the absorbed energy around the planet. These properties strongly influence the planetary day-side flux and hence the expected secondary eclipse depths. While thermal emission measurements alone cannot break the degeneracy that exists between these two properties (an increase in the Bond albedo is indistinguishable from an increase in the efficiency of heat redistribution), they can still provide useful constraints.

To do this I followed the methodology of Cowan & Agol [2011], using the technique outlined in Section 2.6.2 to estimate the day-side effective temperature of the planet. The value found was $T_d = 2280_{-150}^{+200}$ K, based on the brightness temperature estimates for my *Spitzer* IRAC data. Constraints on the Bond albedo and heat redistribution can then be made through the simple parameterisation given in equation 2.37.

Using the system parameters in Tables 3.8 and 3.9, along with equation 2.38, I find $T_0 = 2870 \pm 50$ K, giving $T_d/T_0 = 0.79_{-0.06}^{+0.07}$. Cowan & Agol [2011] have highlighted an emerging trend for highly irradiated planets (with high T_0 values) to have low albedo and energy redistribution efficiency values (i.e. high T_d/T_0 values). With $T_0 = 2870 \pm 50$ K and $T_d/T_0 = 0.79_{-0.06}^{+0.07}$, WASP-3b is another case of such a planet. I note here that this result is based on my *Spitzer* data only. Including a K_s brightness temperature based on the results of Zhao et al. [2012] increases this value to $T_d/T_0 = 0.94 \pm 0.05$.

In Figure 3.18 I plot a 2D probability distribution found for my T_d/T_0 value, determined using the method described in Section 2.6.2. The degeneracy between A_B and ε is clear - a low albedo with a moderate ε reproduces T_d/T_0 equally as well as a moderate albedo with a low ε . However, it can be seen that for WASP-3 high values of both A_B and ε are excluded, and while a large range of ε value are possible, A_B is restricted to lower values (e.g. $A_B < 0.34$ at 1σ).

3.5.2 Atmospheric modelling

In Figure 3.19 I compare my eclipse depths with the 1D plane-parallel models of Fortney et al. [2008]. These models assume solar metallicity, with abundances from Lodders [2003], and are cloud-free. Atmospheric compositions are found under the assumption of chemical equilibrium throughout the atmosphere. Only 2 model inputs




Figure 3.18: A probability distribution found using my T_d/T_0 value with equation 2.39, assessed over a grid of allowable albedo and ε values. Darker regions are of higher probability than lighter regions. The resulting distribution shows the ambiguity between the albedo and heat redistribution parameters, where a low albedo with a moderate ε reproduces T_d/T_0 equally as well as a moderate albedo with a low ε . Also shown are the marginalised PDFs for the two parameters.

are varied. First, in order to simulate varying degrees of redistribution of incident stellar energy around the planet, the irradiating flux at the top of the atmosphere is weighted by a geometric factor (f). Models were produced for $f = 1/4$ to simulate isotropic planetary radiation, $f = 1/2$ for even emission over the dayside only, and $f = 2/3$ to simulate instantaneous re-radiation of flux by the planet (e.g. Burrows et al., 2008; Hansen, 2008; note also that $f = 1/4$ is equivalent to $\varepsilon = 1$ and $f = 2/3$ is equivalent to $\varepsilon = 0$ in equation 2.37). Second, models were produced either with TiO and VO present in the atmosphere with their equilibrium abundances (‘TiO’ in Figure 3.19), or with these molecules removed at $P < 10$ bars (‘no TiO’). This is motivated by the apparent inverted/non-inverted dichotomy amongst hot Jupiter exoplanets. For the ‘TiO’ case, regardless of the choice of f , the atmospheric temperatures of WASP-3b always lead to significant amounts of gaseous phase TiO/VO in the upper atmosphere, which drives temperature inversions. Conversely the ‘no TiO’ cases lack inversions. The differing temperature structures of these models are highlighted in Figure 3.20, which also shows the relative contributions of different

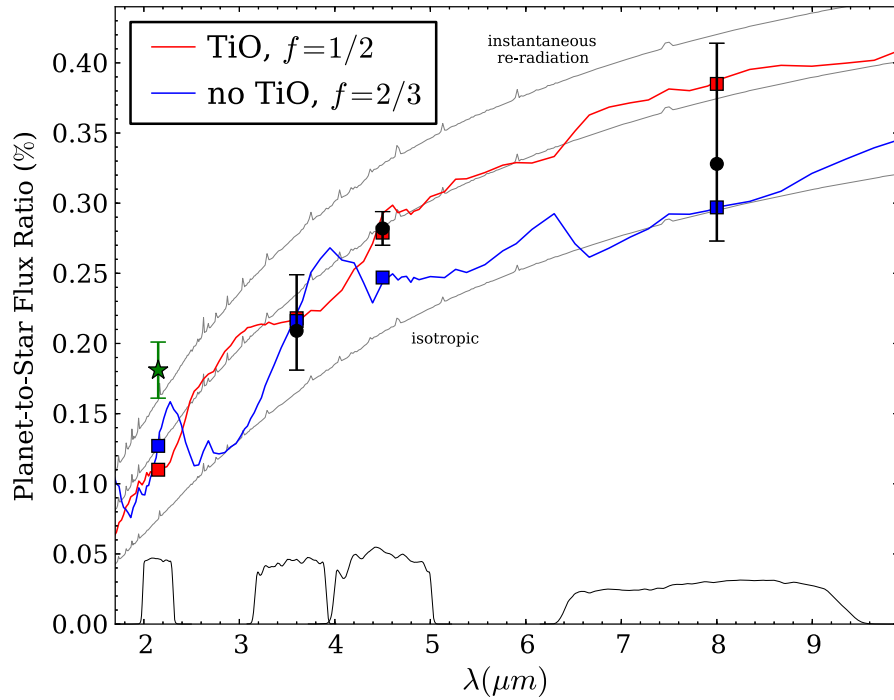


Figure 3.19: *Spitzer*/IRAC eclipse depths are shown along with the K_s band eclipse depth (green star) found by Zhao et al. [2012]. The red curve is the Fortney et al. [2008] model for the case where TiO and VO are present in their equilibrium abundances. The blue curve is the same, but with TiO and VO removed at $P < 10$ bar. The passband-integrated model predictions are given as squares in the same colours. Comparisons of these models to the data are given in the text. The grey curves use planetary blackbody models with Kurucz [1993] stellar models for WASP-3. The top and bottom curves give the extremes of the expected planetary dayside temperature, assuming a Bond albedo of 0. The lower curve is for isotropic re-radiation ($\varepsilon = 1$ in equation 2.37) and the upper curve is for instantaneous reradiation ($\varepsilon = 0$). For reference the upper grey curve has a temperature of 2590 K, and the lower has a temperature of 2030 K. The middle grey curve is the best fitting blackbody to the *Spitzer* bands, with a temperature of 2390 K. IRAC response curves for channels 1, 2 and 4 are also shown with arbitrary units, along with the K_s band response curve.

layers in the modelled atmospheres to the flux in each of the *Spitzer* bands.

For both ‘TiO’ and ‘no TiO’ models I find that inefficient heat redistribution is heavily favoured (Figure 3.19), in agreement with the analysis in Section 3.5.1. For the ‘TiO’ models the data are best reproduced in the $f = 1/2$ ($\varepsilon = 2/5$) case, whereas for the ‘no TiO’ models the extreme $f = 2/3$ ($\varepsilon = 0$) case is favoured.

The passband integrated model predictions for the ‘TiO’ model provide good agreement to the measurements in all three *Spitzer* bands. The ‘no TiO’ model reproduces the channel 1 and 4 eclipse depths well, however the channel 2 model prediction is discrepant, at 3σ . The fact that this is the most robust of my measurements, and that the ‘no TiO’ model uses the maximum reasonable value for f [Hansen, 2008] argues strongly in favour of the ‘TiO’, inverted atmosphere case.

I also compare these models to the K_s band measurement from Zhao et al. [2012], and find that neither reproduces the very high brightness temperature found. A similar situation has been found for the system HAT-P-1 [de Mooij et al., 2011], though with planetary brightness temperatures typically ~ 700 K lower than WASP-3b. de Mooij et al. [2011] highlight that for HAT-P-1b making the lower atmosphere hotter would account for their K_s measurement, since H_2O opacity windows across this band result in sampling flux from these deeper regions. However, for their models such a change would also increase the channel 1 eclipse depth prediction to be inconsistent with the measurement from Todorov et al. [2010]. Comparing the 2 models in Figure 3.19, the ‘no TiO’ case provides a better fit in the K_s band, due to its relatively hotter lower atmosphere, but it still underestimates the eclipse depth by $\sim 3\sigma$. This suggests the presence or absence of TiO and VO is not driving this discrepancy. As a result I still favour the inverted atmosphere conclusion based on the *Spitzer* measurements.

The distinction in the data between the two vertical temperature structures is driven mainly by the $4.5\ \mu\text{m}$ eclipse depth. Its relatively high value compared to the $3.6\ \mu\text{m}$ measurement favours the presence of H_2O and CO being seen in emission (a result of a temperature inversion). This fact has been used in the empirical measure suggested by Knutson et al. [2010]. Following Anderson et al. [2011] I define this measure (ζ) as the gradient of the measurements at 3.6 and $4.5\ \mu\text{m}$ i.e. $(\Delta F_{4.5\ \mu\text{m}} - \Delta F_{3.6\ \mu\text{m}})/0.9\ \mu\text{m}$, minus the corresponding gradient of the blackbody that is the best fit to the two measurements. Planets giving smaller ζ values tend not to show inverted atmospheres, whereas those with inversions tend to have larger ζ values, with a cut-off at around $\zeta \sim -0.05\ \% \mu\text{m}^{-1}$. For WASP-3b I find $\zeta = 0.027 \pm 0.046\ \% \mu\text{m}^{-1}$, placing it in the group of planets with inverted atmospheres.




Figure 3.20: Left panel: Atmospheric pressure-temperature profiles for the models presented in Figure 3.19 (using the same colour scheme). Significantly different temperature structures can be seen, with an inversion occurring for the red (‘TiO’, $f = 1/2$) model at $P \sim 10^{-1}$ bar. Right panel: Normalised contribution functions for my three *Spitzer*/IRAC bandpasses. Solid lines are for channel 1 ($3.6 \mu\text{m}$), dashed lines for channel 2 ($4.5 \mu\text{m}$) and dotted lines are for channel 4 ($8.0 \mu\text{m}$). At wavelengths where we see less deep into the atmosphere (i.e. $8 \mu\text{m}$) the inverted atmosphere is brighter. This difference is what drives the difference in the models plotted in Figure 3.19.

3.5.3 Stellar activity correlation

In Figure 3.21 I reassess the possible correlation between ζ and the chromospheric activity of the host star, as suggested by Knutson et al. [2010]. Values are plotted for the current set of planets with $\log R'_{\text{HK}}$ and ζ values in the literature. The $\log R'_{\text{HK}}$ index is used to quantify the chromospheric activity of stars. It is a measure of the strength of the line cores of Ca II H & K, corrected for the bolometric flux of the star. These line cores sample the stellar chromospheric emission and therefore provide a measure of the chromospheric activity.

A simple measure of the correlation between $\log R'_{\text{HK}}$ and ζ is provided by the Pearson correlation coefficient (r), which varies between 1 (for a perfect correlation) and -1 (for perfect anti-correlation). Without the inclusion of WASP-3, $r = -0.61$. With WASP-3 included $r = -0.59$, meaning that my result weakens the strength of the correlation very slightly. I also compare the quality of fits of some simple functional relationships between $\log R'_{\text{HK}}$ and ζ . To start, I calculate the χ^2 value of a simple weighted average of all the ζ values (see blue line in Figure 3.21). For the 24 systems analysed the χ^2 value here is 133, suggesting a poor fit. This is not




Figure 3.21: Stellar activity measure $\log R'_{\text{HK}}$ plotted against the empirical 3.6–4.5 μm slope measurement, ζ . The labelled systems, WASP-3 (red star), WASP-4 (black triangle) and HAT-P-6 (black square), whilst having very similar host star activities, seem to display quite different vertical temperature structures. The $\log R'_{\text{HK}}$ errorbar for WASP-3 reflects the range of activity values found for this object by Montalto et al. [2012]. I also plot three models of the proposed correlation between $\log R'_{\text{HK}}$ and ζ : a weighted average to all ζ values (blue line), a linear model (red line) and a step function comprised of the weighted average either side of a cut-off (green line). These are discussed in Section 3.5.3. The data here are taken mainly from Mahtani et al. [2013], with additional values from Désert et al. [2011a,b]; Todorov et al. [2012]; Smith et al. [2012]; Anderson et al. [2013].

surprising - given the dichotomy seen in these atmospheres, one would not expect all ζ values to be consistent. I also calculated χ^2 for two simple correlations: a linear slope and a step function (red and green lines in Figure 3.21 respectively). χ^2 values here were 73 and 87, respectively. While these are a clear improvement over the weighted average, they are still not good fits, suggesting the trends in ζ are more complex than simple functions of the host stars activity. Again this is not surprising - there may be many other factors that affect a planet's atmospheric structure and/or ζ values, for example its metallicity (see Chapter 4) or its C/O ratio.

A correlation in the updated dataset still looks plausible, although there is interesting structure around the proposed cut-off between inverted and non-inverted

atmospheres at $\log R'_{\text{HK}} \sim -4.9$. With $\log R'_{\text{HK}} = -4.872$ [Knutson et al., 2010], the WASP-3 system is part of an interesting group of systems that probe the cut-off. The other systems are WASP-4 [Beerer et al., 2011] and HAT-P-6 [Todorov et al., 2012], as highlighted in Figure 3.21. These systems span a range of ζ near the proposed cut-off, with HAT-P-6b being found to support a weak inversion, and WASP-4b found to support either a weak inversion, or no inversion at all [Beerer et al., 2011]. Though the errors on ζ are large, it is still intriguing that the inversion/non-inversion cut-off with activity may not be as clear as initial data suggested. For WASP-3, I have also plotted the range of $\log R'_{\text{HK}}$ values found by Montalto et al. [2012] over a time-span of 3 years. These values span the potential cut-off, so monitoring of ζ and $\log R'_{\text{HK}}$ (for WASP-4 and HAT-P-6 as well as WASP-3) could be very informative, particularly if there is a strong cut-off.

I note here that $\log R'_{\text{HK}}$ may not be well calibrated for WASP-3 due to its spectral type [late-F, Knutson et al., 2010]. For stars of early spectral type, small amounts of emission in the H and K line cores can be difficult to detect due to the higher continuum flux. As such, I expect that any inaccuracy in the calibration of $\log R'_{\text{HK}}$ for WASP-3 would tend to underestimate the true value. A higher activity for WASP-3 would make for an even more intriguing comparison with WASP-4 and HAT-P-6, given the proposed activity-inversion trend.

3.6 Conclusion

I have presented *Spitzer* observations of the WASP-3 system in 3 infra-red bands, centred on 3.6, 4.5 and 8.0 μm and have detected significant secondary eclipse signals from the planet WASP-3b in each band. The planet-to-star flux ratios derived from the eclipse depths in these bands were found to be $0.209_{-0.028}^{+0.040}$ %, 0.282 ± 0.012 % and $0.328_{-0.055}^{+0.086}$ %, respectively, corresponding to brightness temperatures of $T_{3.6\mu\text{m}} = 2280_{-150}^{+210}$ K, $T_{4.5\mu\text{m}} = 2400 \pm 80$ K and $T_{8.0\mu\text{m}} = 2210_{-250}^{+390}$ K.

I found that the strength of thermal emission from WASP-3b suggests the planet absorbs the incident stellar radiation efficiently, and/or redistributes this energy to the night side of the planet inefficiently. The latter point is also heavily favoured when comparing the eclipse depths to the models of Fortney et al. [2008]. These models also favour the presence of a temperature inversion in the atmosphere.

The WASP-3 system probes the cut-off of the suggested correlation between host star activity and the planetary 3.6–4.5 μm spectral slope, and is found to have a slope consistent with other planets with temperature inversions. Two other planets with similar host star $\log R'_{\text{HK}}$ values, WASP-4b and HAT-P-6b, have been found

to have weakly/non-inverted atmospheres. Since these objects appear to be on the cusp of the activity-inversion cut-off it would be interesting to carry out monitoring of their $\log R'_{\text{HK}}$ and $3.6\ \mu\text{m}$ and $4.5\ \mu\text{m}$ eclipse depths. A detection of an anti-correlation between $\log R'_{\text{HK}}$ and the presence of an inversion would provide strong support for the destruction of a high altitude absorber by UV irradiation, as proposed by Knutson et al. [2010].

Chapter 4

Testing the effects of metallicity on the temperature structure of exoplanet atmospheres

4.1 Introduction

With the increasing number of thermal emission detections for extra-solar planets we are now in a position to start characterising and catagorising properties of a reasonably large sample of these objects. One prominent line of enquiry is the characterisation of the apparent dichotomy amongst thermal spectra for hot Jupiters, where some atmospheres host temperature inversions whilst others do not. There have been numerous attempts to shed light on this dichotomy by looking for correlations between inversions and other system properties (as discussed in Chapters 1 and 3).

In this chapter, I explore potential correlations with another important property - the stellar metallicity. This property has already been found to be important in other areas of exoplanet research, for example in the correlation between giant planet frequency and host star metallicity [Udry & Santos, 2007; Mayor et al., 2011], which provides support for the core accretion theory of planet formation. Bulk planetary properties are also expected to correlate with metallicity (e.g. higher metallicity systems would promote planets with larger cores) and some evidence for this has been found [Laughlin et al., 2011; Enoch et al., 2012]. In the context of planetary atmospheres, a planet's emission spectrum depends on the atmospheric temperature structure, which is controlled by opacity sources, which in turn are dependent on the metallicity of the planet. With regard to work on the temperature inversions

observed for many hot Jupiters, a connection to the planet’s metallicity is certainly plausible [Knutson et al., 2009; Zahnle et al., 2009]. For example, if a planet is deficient in metals, perhaps inversions are suppressed on these planets due to a deficiency of the unknown, inversion-driving absorber.

An interesting case in this context is that of WASP-48b [O’Rourke et al., 2014], which does not appear to host an inversion despite being highly irradiated and orbiting a quiet star. This provides a counter example to the suggestions by Fortney et al. [2008] and Knutson et al. [2010]. The authors note that the host is marginally metal-poor, with $[\text{Fe}/\text{H}] = -0.12 \pm 0.12$ and suggest this could be the reason behind the discrepancy. Whilst this case is intriguing, a correlation between inversions and metallicities has not yet been systematically explored through observations. Such a study could provide invaluable information, particularly since the nature of the absorber is still unknown.

Figure 4.1 presents a simple test of such a correlation from existing measurements with *Spitzer’s* IRAC instrument. A measurement of the host star’s metallicity ($[\text{Fe}/\text{H}]$) is compared to ζ - the empirical $3.6 - 4.5 \mu\text{m}$ slope suggested by Knutson et al. [2009] and defined, following Anderson et al. [2011], in Section 3.5.2. A major assumption here is that the metallicity of the planet will correlate with the metallicity of the host star. It is by no means certain that hot Jupiter planets will have the same metallicities as their host stars. Indeed Jupiter and Saturn have super-solar abundances. However, even if a correlation between planetary and stellar metallicity is uncertain, the metallicity of the host star is still the best proxy we have for the planetary composition. The test is simplistic, but matches the nature of our thermal spectral measurements for many of these objects. Often we only have eclipse measurements from the 3.6 and $4.5 \mu\text{m}$ channels of IRAC. Even so, there is a hint of a trend in Figure 4.1 (with a Pearson correlation coefficient of $r = 0.27$) with planets around metal-poor stars having lower ζ values. This is in qualitative agreement with the expectation of an inversion-driving absorber whose presence is dependent on metallicity. A good way to test this hypothesis - that planets around metal poor stars will not show signs of temperature inversions - is to measure ζ for very metal-poor systems. To date, studies of metal poor systems have been rare because there is a bias in the hot Jupiter population towards metal-rich systems [Mayor et al., 2011]. However, with an increasing number of planets being found, we are beginning to push to a wider range of metallicities and a small number of low metallicity systems (with potentially detectable secondary eclipses) have been discovered. In this chapter, therefore, I assess the thermal emission of three metal poor systems through *Spitzer* IRAC secondary eclipse observations.




Figure 4.1: A plot of ζ (a measure of the planetary 3.6–4.5 μm spectral slope) against host star metallicity, $[\text{Fe}/\text{H}]$, for planets with published values in the literature. The grey solid line is a best fit linear model, which accounts for the errors in both ζ and $[\text{Fe}/\text{H}]$, and hints at a positive correlation between the two quantities. The Pearson correlation coefficient between $[\text{Fe}/\text{H}]$ and ζ here is $r = 0.27$. To test this potential correlation I studied WASP-21, WASP-28 and WASP-37 which are low metallicity systems, with $[\text{Fe}/\text{H}]$ values between the dashed blue lines. Note that the original $[\text{Fe}/\text{H}]$ estimates lay between the dashed grey lines. Data here are taken mainly from Mahtani et al. [2013] and Mortier et al. [2013], with additional values from Désert et al. [2011a,b]; Todorov et al. [2012]; Smith et al. [2012]; Anderson et al. [2013].

The systems chosen for this study were WASP-21, WASP-28 and WASP-37. The metallicities for these systems are marked in Figure 4.1 and brief descriptions are provided below. Note that, since the *Spitzer* observations were proposed, the metallicity estimates for WASP-21 and WASP-28 have been revised and are now less extreme, as highlighted in Figure 4.1.

WASP-21

WASP-21b was discovered by Bouchy et al. [2010] as a Saturn-mass planet with $M_p = 0.30 \pm 0.01 M_J$ and $R_p = 1.07 \pm 0.06 R_J$, orbiting a metal-poor ($[\text{Fe}/\text{H}] = -0.46 \pm 0.11$) thick-disk host in a 4.3 d ($a = 0.05$ au) orbit. The host and planet masses were revised down by Barros et al. [2011] who noted that the host star

was slightly evolved. Further photometric follow-up by Ciceri et al. [2013] gave $M_p = 0.28 \pm 0.02 M_J$ and $R_p = 1.16 \pm 0.05 R_J$. A study of the properties of numerous planet host stars has also altered the metallicity for this system slightly, with the best estimate now at $[\text{Fe}/\text{H}] = -0.22 \pm 0.04$ [Mortier et al., 2013].

WASP-28

Anderson et al. [2014] discovered WASP-28b, an $M_p = 0.91 \pm 0.04 M_J$, $R_p = 1.21 \pm 0.04 R_J$ planet in a 3.4 d ($a = 0.04$ au) orbit around its metal-poor host ($[\text{Fe}/\text{H}] = -0.29 \pm 0.10$). It has recently been followed-up with the *Kepler* K2 mission, with consistent results to the Anderson et al. [2014] study being found. As for WASP-21, Mortier et al. [2013] present a revised value for the host's metallicity, putting it at $[\text{Fe}/\text{H}] = -0.12 \pm 0.03$.

WASP-37

WASP-37b was discovered by Simpson et al. [2011] and is a $M_p = 1.80 \pm 0.17 M_J$, $R_p = 1.16 \pm 0.07 R_J$ planet in a 3.6 d ($a = 0.04$ au) orbit around a metal-poor G2V star, with $[\text{Fe}/\text{H}] = -0.40 \pm 0.12$. No further follow-up work has been carried out on this planet nor on the host star.

I order the remainder of this chapter as follows: in Section 4.2 I describe the *Spitzer* IRAC secondary eclipse observations of the three systems studied; in Sections 4.3, 4.4 and 4.5 I present the analyses of these observations for WASP-21, WASP-28 and WASP-37, respectively; in Section 4.6 I summarise the results from these analyses and in Section 4.7 I discuss the implications of the results. Conclusions are given in Section 4.8.

4.2 Observations

Since July 2009 *Spitzer* has been operating in a warm mode, after its liquid helium cryogen was exhausted. During the warm mission only the IRAC instrument has been available (Section 2.1.2), with only the 3.6 and 4.5 μm bands (here referred to as channels 1 and 2) operating. Observations of WASP-21, WASP-28 and WASP-37 were taken at the predicted time of secondary eclipse in channels 1 and 2. Since simultaneous observations of the same field cannot be taken in both channels, separate secondary eclipse events were observed for each channel resulting in 6 sets of observations.

Table 4.1: Details of the secondary eclipse observations taken for WASP-21, WASP-28 and WASP-37, including the start (T_{start}), predicted mid-eclipse (T_{mid}) and end times (T_{end}). T_{exp} is the effective exposure time of the observation (in seconds), while N_{obs} is the total number of exposures that were taken.

Target	Channel	Date	T_{start}	T_{mid}	T_{end}	T_{exp}	N_{obs}
WASP-21	1	2012 Aug 23	06:46	11:43	14:51	1.92	14464
WASP-21	2	2012 Aug 27	14:44	19:27	22:49	1.92	14464
WASP-28	1	2012 Aug 31/Sep 01	21:26	02:07	05:14	4.4	3960
WASP-28	2	2012 Sep 07/08	17:14	21:45	01:02	4.4	3960
WASP-37	1	2012 April 09	14:23	18:47	22:06	4.4	3910
WASP-37	2	2012 May 01	01:22	05:56	09:05	4.4	3910

Each observation set made use of the new IRAC Peak-up mode (PCRS), which allows the target to be positioned to an accuracy of 0.1 pixels. The targets were placed near the centre of a pixel which has been determined to minimise the effects of the intra-pixel sensitivity (see Section 2.1.4).

The details of these observations are given in Table 4.1. Each set of observations lasted around 8 hrs. Observations start and end times were chosen so as to allow enough time to determine the out of eclipse baseline level with sufficient accuracy. Also, the times were chosen such that more time was spent observing the baseline before the eclipse than after it, to minimise the impact of any settling effects of the telescope. The eclipse timing prediction was based on the assumption of a circular orbit in each case.

The exposure times in each case were chosen such that the targets were well below the level where saturation and non-linearity effects of the frames become significant¹. For WASP-21, IRAC’s subarray mode was used in both channels to allow for an appropriately short exposure time, while not using up too much memory on the on-board computer. For WASP-28 and WASP-37 the full array mode was used.

4.3 Analysis of WASP-21

To analyse the observations of WASP-21 described in Section 4.2 and Table 4.1, I used the Basic Calibrated Data (BCD) frames downloaded from the SHA.² The production of these BCD frames from the raw data is described in Section 2.1.3. The specific version of the BCD pipeline used here was 19.1.0.

¹<http://ssc.spitzer.caltech.edu/warmmission/propkit/som/>

²<http://irsa.ipac.caltech.edu/data/SPITZER/docs/spitzerdataarchives/>

4.3.1 WASP-21 pre-aperture photometry

Before I applied aperture photometry to the BCD frames, a number of additional modifications and checks were made to minimise the possibility of systematics in the data. These are described below, with the same analyses being applied to channels 1 and 2 unless otherwise stated.

First, since the data in both channels were taken in sub-array mode, they came in the form of data cubes - each containing 64 32×32 pixel images (with a pixel scale of $1.2''/\text{pixel}$). In each channel there were 226 cubes for the whole observation. These were split into 14464 individual 32×32 images, to allow for sampling of short time scale variations, similar to those found in the WASP-3 sub-array data (Section 3.3.1.1). New FITS headers were created for each individual image, based on the header of the data cube from which it came. The only header value that was changed was 'HMJD_OBS' which was altered to reflect the time at mid-exposure of the individual image.

Pixel values in each frame were then converted from MJy/sr to electrons by multiplying by the gain and effective exposure time and dividing by a flux conversion factor. An estimate for the zodiacal background flux in the subtracted sky dark was also added to the data. Each of these values were found in the FITS header of each frame.

Mask files outputted from the BCD pipeline were also applied. These masks flag such things as saturated or non-linear pixels, cosmic ray hits and column pull-down. In both channels individual masks for each individual image were supplied, allowing temporarily discrepant pixels to be masked. Masking was done by changing the relevant pixel value for the frame to a value of NaN. The IDL aperture photometry software used in this chapter treats NaN values as masked points.

Next, a background trend repeated within each data cube was assessed. This trend was very similar to the one found for WASP-3 (Section 3.3.1.1). Figure 4.2 shows this trend for channel 1, along with the corresponding background subtracted source fluxes. In the top panel it can be seen that the 1st and 58th images within each data cube are the largest outliers in term of their background flux. In the background subtracted source fluxes (bottom panel) the trend is no longer seen, but the 1st image is still a clear outlier and the 58th image is a marginal outlier. Due to this, I chose to remove the 1st and 58th image from each cube. Figure 4.3 shows the equivalent plot as Figure 4.2 but for channel 2. Following a similar reasoning to channel 1, the first image in each cube was removed for channel 2.

Masking of cosmic rays was then carried out. To do this I created individual pixel light curves that were sigma-clipped at 6σ , to ensure only true outliers were




Figure 4.2: *Upper panel:* Channel 1 clipped mean background values for a given sub-array image within each data cube for WASP-21. For each sub-array image, the average value across the 215 cubes is given here. An asymptotic trend in the background values is seen, very similar to that found for the warm mission sub-array data for WASP-3 (Figure 3.5). The 1st and 58th images stand out as being discrepant from this trend.

Lower panel: As above, except here I show the background subtracted source flux. Here, the 1st image is a clear outlier, while the 58th image is still slightly discrepant. As a result I decided to remove the 1st and 58th images from each cube.

masked. The effect of the variable sub-array backgrounds was accounted for by first subtracting the median frame value from each pixel on a frame-by-frame basis. Flux variations in pixels due to the wobble of the telescope (see Section 2.1.1) were accounted for by creating a windowed, median smoothed light curve for each pixel and subtracting this away from the unsmoothed data. The size of the window was chosen based on the time-scale of the flux variations in the brightest pixel. σ values were calculated using $1.48 \times$ the median absolute deviation [MAD; Ruppert, 2011] of the smoothed-subtracted light curves. I used the MAD, rather than the standard deviation, because it is more robust to the effects of extreme outliers, like cosmic ray hits. Again, pixels were masked here using NaN values.

Masking was also applied to noisy pixels in the background region that were found to have abnormally high standard deviations (post-cosmic ray masking). In




Figure 4.3: Equivalent plot to Figure 4.2, but here for the channel 2 data of WASP-21. The background values here show a weaker repeated trend than for channel 1, with values rising for the first 15 sub-array images, before settling down (and perhaps decreasing slowly). As for channel 1, the 1st and 58th images are clear outliers. The discrepancy in the 58th image appears to be well corrected for in the background subtracted source flux, but the 1st image is still an outlier. Therefore, the 1st image is removed from each data cube.

total, 7 pixels in channel 1 and 11 pixels in channel 2 were masked in this way.

Columns showing signs of pull-down/up were also masked. These are columns for which the intensity is altered by the presence of a very bright source or a cosmic ray hit. In contrast to the method used for WASP-3’s channel 2 data (see Section 3.3.1.2), this was done using a manual approach. As for the cosmic ray masking, the variable sub-array background was accounted for by subtracting the median frame value from each pixel. Time series of each of the 32 columns in each frame were found using median column values and frames containing a column with a noticeably large or small value were then checked manually for signs of pull-down or pull-up. If present, the columns were masked. In channel 1, 3 frames were affected while in channel 2, 10 frames were affected.

The final stage of the pre-aperture photometry analysis was to apply manual masks. These were used to mask detector systematics and other stars in the field-of-

view that could have affected the background estimate in the aperture photometry analysis (Section 4.3.2). In both channels a star around 15 pixels ($18''$) from WASP-21 was masked. Also, in channel 1 the top left of the frame was systematically brighter than the rest of the background (due to a positive residual bias pattern³) and so this region was masked. In channel 2, the residual bias pattern was not present, but one column was permanently pulled down and so it was masked. Note that none of the source apertures used in the aperture photometry sampled these masked regions.

4.3.2 WASP-21 aperture photometry

In each channel, aperture photometry was performed on each frame using the IDL software described in Section 2.3.2. Before extracting the data to be used in the MCMC parameter fitting, the effects of choosing different inputs to the IDL software were tested, to ensure appropriate choices were made. The inputs tested were: the size of the box used for centroiding; the background annulus inner and outer radii and the sigma clipping level of the background clipped mean estimate.

Tests of the box sizes used for centroiding were then carried out, using sizes of 5×5 , 7×7 and 9×9 pixels (with $1.2''/\text{pixel}$), each centred on the PSF. Boxes of 7×7 pixels were chosen for both channels. The 5×5 pixel box gave a noisier distribution in the measured x and y positions of the PSF and while the 9×9 pixel box gave marginally tighter x and y distributions, it was more susceptible to centroiding failures in the IDL software (perhaps due to the presence of more masked points in the larger box).

The inner background radius for both channels was tested using values from $r_{\text{in}} = 5\text{--}15$ pixels and in both channels $r_{\text{in}} = 12$ pixels was selected. This value for the inner annulus avoided contributions from the PSF wings while maximising the number of background pixels to minimise the photometric error contribution from the uncertainty in the background estimate (the $\frac{n_{\text{source}}}{n_{\text{bkg}}} \sigma_{\text{bkg}}^2$ term in equation 2.8).

The outer background radius was tested using values from $r_{\text{in}} = 15\text{--}25$ pixels and in both channels $r_{\text{out}} = 25$ pixels was selected. This value for r_{out} includes all the pixels in the (32×32 pixel) sub-array images - again minimising the uncertainty on the background estimate.

The sigma clipping level for the background clipped mean estimate was then tested. For both channels a clipping level of 3σ was chosen. This was selected by comparing the background and uncertainty values for reductions using 3, 4 and 5σ clipping. In both channels it was found that for the 3σ case the background

³<http://irsa.ipac.caltech.edu/data/SPITZER/docs/irac/warmfeatures/>




Figure 4.4: The WASP-21 channel 1 raw fluxes against dx and dy position and also against time. $dx = x - \hat{x}$ and $dy = y - \hat{y}$ are the positions of the PSF centre relative to their weighted means. There is a clear correlation between the flux and y position, but no obvious correlation with the x position. The effect of this correlation can be seen in the panel on the right-hand side - as the image of WASP-21 moves with *Spitzer's* pointing wobble it samples different parts of the intra-pixel sensitivity function, giving rise to a $\sim 1\%$ variation in flux.

uncertainty contributed least to the total flux error budget (as one would expect given σ_{bkg} term in σ_{phot}), but that the measured background levels were not altered to any significant degree (compared with the 4 and 5σ cases).

Using the inputs selected above, I ran the IDL aperture photometry software on the data for each channel using source aperture radii (r_{source}) from 2–6 pixels, in steps of 0.5 pixels. Values extracted from each frame were: the HMJD(UTC) time at mid-exposure; the background subtracted source flux; the photometric error and the x and y centroid position from the Gaussian fit (as described in Section 2.3.2).

I then applied sigma-clipping to these extracted data. This was done by rejecting any frame with a flux or x or y value more than $k\sigma$ from the median of the surrounding 250 frames in channel 1 and 300 frames in channel 2. I tested k values of 3, 4 and 5. $k = 5$ was found to be too relaxed, but there was very little difference in the time-series produced by the 3 and 4σ clipping. As a result I decided to use the more conservative 4σ clip for both channels.

Two additional corrections were made to the extracted data sets. First, there appeared to be ramps in the background values in both channels, where the back-




Figure 4.5: An equivalent plot to Figure 4.4, for the WASP-21 channel 2 data. No clear correlations are seen in the raw flux with position and there are no signs of quasi-periodic variations in the flux. However, the intra-pixel variations in this channel are typically below 0.5% while the RMS on the raw fluxes here is 0.7%. Indeed, evidence is found for a correlation in the detrending functions tests in Section 4.3.4.

ground flux was systematically low at the beginning of the observations, before settling. As a result of this the first data cube was removed for channel 1 and the first two data cubes for channel 2 were removed. The final datasets used (for the optimal source aperture radii of $r_{\text{source}} = 2.5$ pixels - see Section 4.3.5) contained 13904 and 14051 measurements for channels 1 and 2, respectively.

Second, timing corrections were made. HMJD(UTC) values were converted to the HJD(TDB) timing standard. In both channels the conversion from UTC to TDB was applied by adding 67.184s to the times. The light travel time of the system was also accounted for. This effect causes the secondary eclipse to appear later than the $\phi = 0.5$ prediction by $\frac{2a}{cP}$ (for a circular orbit). To account for this effect a value of 53.112s was subtracted from the times in both channels, with the value for $\frac{2a}{cP}$ being determined from an initial MCMC analysis, where the light travel correction had not been made to the IRAC data. The resulting datasets from this aperture photometry analysis were used as inputs for the MCMC parameter fitting (see Section 4.3.3).

In Figures 4.4 and 4.5, I have plotted the raw fluxes extracted from the aperture photometry analysis (for $r_{\text{source}} = 2.5$ pixels) against the x and y PSF

centroid and also against time. Figure 4.4 shows that the channel 1 data display the intra-pixel sensitivity variations that are common in observations with the IRAC instrument (see Section 2.1.4). The variations here are dominated by the correlation of the flux with the y position. Note that the time-scale of the variations in flux is around half of what it was for the WASP-3 channel 1 data (see Figure 3.9). This is due to an increase in the cycling frequency of the heater that causes *Spitzer's* pointing wobble⁴. In channel 2 (Figure 4.5), while there is no obvious sign of the intra-pixel sensitivity variations in the raw data, these effects in this channel are typically below 0.5%. The RMS on the raw fluxes here is 0.7%, so it may be that the trends are hidden in the point-to-point scatter of the raw fluxes. Indeed, in the testing of detrending functions to remove the intra-pixel sensitivity variations (see Section 4.3.4), I found that functions with terms in both x and y are favoured.

4.3.3 WASP-21 MCMC set up

With the aperture photometry analysis complete, secondary eclipse light curve solutions were explored using the Markov Chain Monte Carlo code MCMCTRANSIT, described in Section 2.4. Briefly, from a set of proposal parameters (T_0 , P , δ , T_{14} , b , K_1 , T_{eff} , $[\text{Fe}/\text{H}]$, $\sqrt{e} \cos \omega$, $\sqrt{e} \sin \omega$, $\Delta F_{3.6\mu\text{m}}$, $\Delta F_{4.5\mu\text{m}}$), secondary eclipse light curves in each channel were modelled by evaluating the visible fraction of the planetary disk scaled by the day-side planet-to-star flux ratio, for the given channel. The planetary flux contribution was then removed from the measured flux to give the stellar flux, using equation 3.8. The stellar flux, which contained the IRAC intra-pixel sensitivity variations, was then modelled using polynomial fits (i.e. subsets of equation 2.1), with the coefficients being obtained using singular value decomposition [Press et al., 1992].

The constraints on parameters in MCMCTRANSIT that are normally applied through model fits to transit data were instead applied here using priors (see Section 2.4.2.3). For WASP-21, the ephemeris and transit light curve priors were taken from Ciceri et al. [2013], in which an analysis of all the transit data for this object (along with two new light curves from medium-class telescopes) was performed. The ephemeris values from this study were $T_0 = \text{BJD}(\text{TDB}) 2454743.04052 \pm 0.00071$ and $P = 4.3225186 \pm 0.0000030$ days. Again, note that while the *Spitzer* IRAC data uses HJD times, the difference between BJD and HJD is at most around 1 s, so it is negligible in this work.

The transit light curve priors used were $\frac{R_*}{a} + \frac{R_p}{a}$, $\frac{R_p}{R_*}$ and i - the parameters used in the transit light curve fitting in Ciceri et al. [2013]. The values I used for

⁴<http://ssc.spitzer.caltech.edu/warmmission/news/21oct2010memo.pdf>

these priors were:

$$\frac{R_{\star}}{a} + \frac{R_p}{a} = 0.1169 \pm 0.0031$$

$$\frac{R_p}{R_{\star}} = 0.1055 \pm 0.0023$$

$$i = 86.97 \pm 0.33^{\circ}$$

The T_{eff} and $[\text{Fe}/\text{H}]$ prior values, used to determine the stellar mass and thus set the scale of the system (see Section 2.4.1), were taken from a recent study by Mortier et al. [2013] in which a uniform analysis of the properties of 90 exoplanet host stars was performed. The values used were $T_{\text{eff}} = 5924 \pm 55$ K and $[\text{Fe}/\text{H}] = -0.22 \pm 0.04$. Note that while the new metallicity estimate is higher than the old estimate $[[\text{Fe}/\text{H}] = -0.46 \pm 0.11$; Bouchy et al. 2010], WASP-21 is still very much a low metallicity system.

Radial velocity data were taken from Bouchy et al. [2010], which comprises measurements from four instruments (FIES, CORALIE, SOPHIE and HARPS). Timing corrections were made to these data, as for the IRAC data, to account for a change in the timing system used (from UTC to TDB) and the light travel time of the system. Since the data span a few years, the UTC to TDB conversion varied between 66.184–67.184 s. The light travel time correction for radial velocities was $\frac{a}{c}$ since the radial velocity signal originates from the star (as explained in Section 3.3.5). For WASP-21 this correction was applied by subtracting 26.559 s from the radial velocity times.

Using the set up described above, I carried out an initial set of MCMC runs to ensure that a robust parameter set solution could be obtained for WASP-21. This was done by running 6 independent MCMC runs and checking that the resulting posterior distributions had converged and were well mixed. This was done using the Gelman-Rubin statistic [Gelman & Rubin, 1992], which analyses the intra-chain parameter variances and compares them to the inter-chain variances. Gelman-Rubin statistic values close to unity imply good convergence and mixing, and for each of the proposal parameters used in MCMCTRANSIT this was indeed the case. For these runs, the detrending function used to correct the intra-pixel sensitivity variations was:

$$F_{\star, \text{model}} = \hat{F}_{\star} + a_0 + a_x dx + a_y dy + a_{xx} dx^2 + a_{yy} dy^2 + a_t dt, \quad (4.1)$$

and the production run comprised 10^4 (successful) jumps. In both channels the IRAC datasets used were those for which $r_{\text{source}} = 3$ pixels.

4.3.4 WASP-21 detrending functions tests

With the above MCMC set-up, I conducted a series of tests to determine which detrending functions (i.e. which subsets of equation 4.1) optimised the secondary eclipse fits in each channel. For these tests I used aperture photometry datasets for which $r_{\text{source}} = 3$ pixels had been used for the flux measurements, in both channels. The MCMC production runs for each test used 10^4 steps.

To assess the relative quality of the fits given by the various detrending functions I used the Bayesian Information Criterion [BIC; Schwarz, 1978], as described in Section 2.5. Since the BIC is based on χ^2 , the photometric errors on the secondary eclipse datasets were not scaled during these detrending tests.

Initially MCMC runs were carried out on the data from the two IRAC channels separately. For each channel a variety of subsets of equation 4.1 were tested and the best two detrending functions (according to their BIC values) for each channel were used for further testing. If the BIC value for the next best detrending function also provided a reasonable fit then it was also tested (i.e. if the difference in the BIC to the best fit model was less than around 6; see Table 2.3).

Combined tests were then carried out, where MCMC runs were made using data from both of the IRAC channels. Different combinations of the detrending functions chosen from the individual fits were tested in order to find the best overall fit. The results of these combined tests are given in Table 4.2.

The results show that for channel 1 the favoured detrending function was:

$$F_{\star, \text{model}} = \hat{F}_{\star} + a_0 + a_x dx + a_y dy + a_{yy} dy^2 + a_t dt \quad (4.2)$$

while in channel 2 the favoured detrending function was:

$$F_{\star, \text{model}} = \hat{F}_{\star} + a_0 + a_x dx + a_{xx} dx^2 + a_y dy + a_{yy} dy^2 \quad (4.3)$$

In channel 1, the two detrending functions assessed gave similar eclipse depths. In channel 2 the agreement was not as good, with the detrending function that was linear in x and y giving a slightly shallower eclipse than that given by the functions with quadratic terms. However, the eclipse depths were within 1σ agreement and since the latter function is favoured by the BIC ($\Delta\text{BIC} = 6.1$; see Table 2.3), I proceeded in using this for the channel 2 detrending function. The chosen detrending functions were then used in subsequent tests of the optimal source aperture radius.

Table 4.2: Eclipse depths and BIC values for fits to the WASP-21 channel 1 and 2 data using different combinations of detrending functions. Note that these are ordered in terms of their BIC values (lowest first). The detrending functions chosen for these combined tests were those that performed the best in MCMC fits to individual datasets. The eclipse depths for channel 1 are consistent, while in channel 2 the eclipse depths for the function using only the linear terms in x and y are lower by around 1σ . However, with the BIC value being higher by 6.1 than the best combination, there is no reason to favour the linear detrending function.

Channel 1		Channel 2	
Detrending function ($-\hat{F}_*$)	Eclipse depth (%)	Detrending function ($-\hat{F}_*$)	Eclipse depth (%)
$a_0 + a_x dx + a_y dy + a_{yy} dy^2 + a_t dt$	0.073 ± 0.011	$a_0 + a_x dx + a_x dx^2 + a_y dy + a_{yy} dy^2$	0.062 ± 0.013
$a_0 + a_x dx + a_y dy + a_{yy} dy^2 + a_t dt$	0.073 ± 0.011	$a_0 + a_x dx + a_y dy$	0.048 ± 0.013
$a_0 + a_x dx + a_x dx^2 + a_y dy + a_{yy} dy^2 + a_t dt$	0.074 ± 0.011	$a_0 + a_x dx + a_y dy$	0.048 ± 0.013
$a_0 + a_x dx + a_y dy + a_{yy} dy^2 + a_t dt$	0.073 ± 0.011	$a_0 + a_x dx + a_y dy + a_{yy} dy^2$	0.060 ± 0.013
$a_0 + a_x dx + a_x dx^2 + a_y dy + a_{yy} dy^2 + a_t dt$	0.073 ± 0.011	$a_0 + a_x dx + a_x dx^2 + a_y dy + a_{yy} dy^2$	0.061 ± 0.014
$a_0 + a_x dx + a_x dx^2 + a_y dy + a_{yy} dy^2 + a_t dt$	0.073 ± 0.010	$a_0 + a_x dx + a_y dy + a_{yy} dy^2$	0.060 ± 0.013
			BIC = $\chi^2 + k \ln N$
			31260.5
			31266.6
			31272.2
			31274.4
			31279.5
			31279.9




Figure 4.6: Plots of the normalised secondary eclipse residual RMS and eclipse depth values for WASP-21, in channels 1 and 2, as a function of the source aperture radius used in the aperture photometry analysis (Section 4.3.2). In both channels, the residual RMS value is minimised for $r_{\text{source}} = 2.5$ pixels, and worsens steadily as r_{source} increases. The channel 2 eclipse depths are stable, except for $r_{\text{source}} = 2.0$ pixels. In channel 1, there is a clear trend in the eclipse depth with source aperture radius, which is discussed in the text.

4.3.5 WASP-21 aperture radius tests

With the optimal detrending functions selected, the optimal r_{source} value used in the aperture photometry was then determined. This was done by running MCMC runs on the aperture photometry datasets extracted using $r_{\text{source}} = 2\text{--}6$ pixels in steps of 0.5 pixels, using the detrending functions in equations 4.2 and 4.3. In each channel the optimal r_{source} value was determined as that which minimised the RMS on the normalised residuals of the secondary eclipse fit for that channel. Since the photometric errors on the secondary eclipse data do not contribute to the RMS value, they were rescaled for these runs (to match Q to the number of degrees of freedom).

Figure 4.6 shows the results of these tests. The left column of this plot shows

the results for channel 1, while the right column shows the results for channel 2. The top row shows the RMS values against r_{source} , while the bottom row gives the eclipse depths versus r_{source} . In both channels the normalised residual RMS is minimised for $r_{\text{source}} = 2.5$ pixels. In channel 2 the eclipse depth is reasonably stable across the range of r_{source} tested, showing the eclipse depth measurement is robust. The eclipse depth for $r_{\text{source}} = 2.0$ pixels is somewhat discrepant in comparison to the other values in this channel. This may be a sign that this small aperture is cutting too far into the PSF and could be introducing additional systematic trends into the flux measurements.

In channel 1 there is clearly a problem. A systematic trend of the eclipse depths with r_{source} is seen, suggesting the eclipse depths measurements are not robust with respect to the source aperture radius. In the next section I carry out tests to try to characterise the causes of this systematic trend.

Figure 4.7 shows plots (for both channels) of the binned normalised residual RMS values as a function of bin width, for a range of source aperture radii. The expectation for these plots, for the case of white noise residuals, is that the binned RMS values will decrease as the square-root of the number of points contained in each bin (shown as the dotted lines). Significant deviations away from this as binning increases suggests red noise in the residuals, which could bias the secondary eclipse values. Perhaps unsurprisingly, there is a deviation away from the white noise expectation for channel 1 (red noise signals can also be seen in the light curve residual plot of Figure 4.19), further motivating additional tests into this dataset. Note, however, that the red noise component is not suppressed at larger radii, where the RMS is degraded. Therefore, there is no motivation from this analysis to move to larger source aperture radii in channel 1.

Channel 2 shows better behaved residuals, following the white noise expectation for all but the largest bin sizes (and even here the discrepancy is not significant). As for channel 1 there is no motivation here to change from the $r_{\text{source}} = 2.5$ pixel aperture selected by RMS minimisation.

4.3.6 WASP-21 additional tests

To explore the trend seen in the eclipse depth in channel 1 (Figure 4.6) I first tested if the same trend was found when different detrending functions were used. I tried a variety of alternative functions and found that the same trend appeared, suggesting that it is not a result of the detrending function I have used. The detrending functions tested were:

$$F_{\star, \text{model}} = \hat{F}_{\star} + a_0 + a_y dy \quad (4.4)$$




Figure 4.7: Binned, normalised residual RMS values as a function of bin width for WASP-21 secondary eclipse fits. The top panel is for channel 1 and the bottom panel is for channel 2. Black lines are for the optimal source aperture radius found in Section 4.3.5 (2.5 pixels), while the red, blue and green lines are for 3.0, 4.0 and 5.0 pixel radii, respectively. The dotted lines show the $1/\sqrt{n}$ expectation, where n is the number of data points per bin. This expectation is fixed to the unbinned RMS value for the optimal source aperture radius. A red noise signal appears for channel 1 for the larger bin widths, reflecting the systematic trends seen in the residuals for this channel in Figure 4.19. The residuals for channel 2 are much better behaved, staying close to the white noise expectation for all but the largest bin widths. Note that vertical offsets between the different source aperture radii reflect what is seen in the top panels of Figure 4.6.

$$F_{\star, \text{model}} = \hat{F}_{\star} + a_0 + a_y dy + a_t dt \quad (4.5)$$

$$F_{\star, \text{model}} = \hat{F}_{\star} + a_0 + a_x dx + a_{xx} dx^2 + a_y dy + a_{yy} dy^2 + a_t dt \quad (4.6)$$

These were chosen as I felt they bracketed the plausible detrending functions for the channel 1 data, given the correlations in Figure 4.4.

It is interesting to see what the eclipse fits look like for the different r_{source} values and how they differ. In Figure 4.8, I show the eclipses (i.e. the raw flux with the detrending model taken out) for $r_{\text{source}} = 5$ pixels and $r_{\text{source}} = 2.5$ pixels. The bottom two panels show the differences in the raw fluxes and the detrended fluxes for these two cases. A systematic slope in the flux difference spanning the eclipse ingress (highlighted in Figure 4.8) seems to be driving a suppression of the eclipse depth for the $r_{\text{source}} = 5$ pixel case.

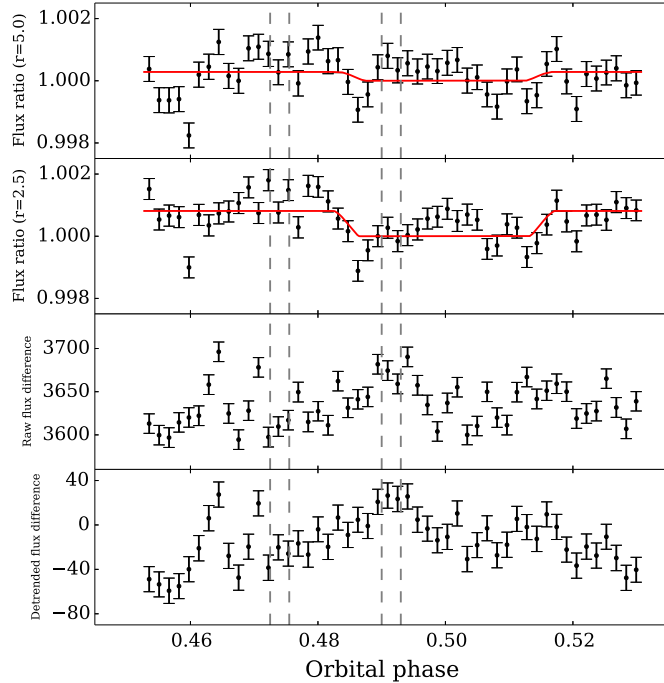


Figure 4.8: The top two panels show the eclipse light curves (i.e. the raw flux with the detrending model taken out) for the channel 1 data of WASP-21, for $r_{\text{source}} = 5$ pixels and $r_{\text{source}} = 2.5$ pixels. The third panel from the top shows the difference in the raw flux for these two apertures, while the bottom panel gives the difference in the detrended fluxes. A systematic slope in the flux difference spanning the eclipse ingress is highlighted by the dashed grey lines. This feature in the difference light curves is what appears to be driving the variation in eclipse depth with r_{source} seen in Figure 4.6.

One potential cause of this could be that there is a temporally changing bad pixel (or pixels) affecting the photometry at larger r_{source} values. To test this I created separated stacked images from the extremes of the systematic slope seen in Figure 4.8 (within the phases marked with the dashed lines) and differenced these. The resulting difference image is shown in Figure 4.9 along with 2.5 and 5.0 pixel apertures. In the annulus between these two apertures there is no sign of any areas of systematically low or high counts, suggesting the eclipse depth trend is not the result of bad pixels.

Another possible source for the systematic in the bottom panels of Figure 4.8 is that there are variations in the background flux in the 2.5 – 5.0 pixel annulus that are not accounted for by the background subtraction. To test this I performed




Figure 4.9: A difference image of WASP-21 in IRAC channel 1. Median stacked images were created at the extremes of the systematic slope highlighted in Figure 4.8 and then differenced. I only used images where the determined centroid values were similar (i.e. within a 0.01×0.01 pixel box) in an attempt to suppress the effects of the movement of the PSF. Also plotted are 2.5 and 5.0 pixel apertures (red lines), between which there are no signs of any bad pixels that could drive the eclipse depth trend seen in Figure 4.6.

aperture photometry on a blank part of the frame using a 2.5 and a 5.0 pixel aperture. The flux values were background subtracted in the same way as for WASP-21 and a difference light curve was made. This is shown in Figure 4.10 and should be compared to the third panel of Figure 4.8. Interestingly, the flux differences for the blank part of the frame show systematic features with amplitudes comparable to those seen for WASP-21, suggesting that spatial structure in the background variations could be the cause of the problem. If they are, correcting the eclipse depth trend would be very difficult, if not impossible. As noted in Section 4.3.2, I checked that the photometry was not significantly affected (i.e. to the levels of the systematics here) by different sigma-clipping levels for the background. However, this does not guard against the local variations in the background that may be a problem here.

Since a solution to the eclipse depth trend is not apparent, I will proceed with the WASP-21 channel 1 data by setting an upper limit on the eclipse depth.




Figure 4.10: Difference light curve for a blank part of the WASP-21 channel 1 frames. Aperture photometry was performed on this region using 2.5 and 5.0 pixel apertures and background values were subtracted in the same way as for WASP-21. The resulting difference light curve shows comparable systematic features to those in the third panel of Figure 4.8 - suggesting that local variations in the background could be driving the trend seen in Figure 4.6.

4.3.7 WASP-21 optimal MCMC analysis

From the analyses in Sections 4.3.5 and 4.3.6 I chose to set an upper limit on the eclipse depth for channel 1. I chose a value of 0.11%, which reflected both the upper limits of the systematic trend seen in Figure 4.6 and also a prayer bead analysis (see Section 3.3.10) I carried out for this channel.

For channel 2, using the detrending function and aperture radius selected in Sections 4.3.4 and 4.3.5, I ran a final, longer MCMC fit using 10^5 steps in the production run, to ensure the parameter space was being fully explored. I did not include the channel 1 data in this analysis. The final eclipse depth derived for this channel was $0.064 \pm 0.014\%$. The optimal system parameter values and uncertainties given in the results section (Section 4.6 and Table 4.5) were taken directly from the output of this MCMC run; no additional error assessments were made. The errors on the channel 2 eclipse data were scaled by 1.05 for this analysis, suggesting that the fit to the data was reasonable before the rescaling of the errors.

Figure 4.19 shows binned time series for WASP-21 in channels 1 and 2. The channel 1 plots are for illustrative purposes only. The data and model fits shown for this channel (using $r_{\text{source}} = 2.5$ pixels) were only used in deciding the upper limit on the eclipse depth. The top row of this plot shows the raw flux light curves along with the eclipse model (which includes the detrending function). The middle row shows the fluxes with the detrending functions removed and normalised to the flux of the star, along with the eclipse model. The model residuals are shown in the bottom row.

4.4 Analysis of WASP-28

4.4.1 WASP-28 pre-aperture photometry

The pre-aperture photometry analysis for WASP-28 followed a very similar form to that applied to WASP-21 (see section 4.3.1) and so is described only briefly here, except for significant differences in the analyses.

The main difference between the WASP-21 and WASP-28 data is that the WASP-28 data were taken in full array mode, so modifications associated with the sub-array data were not required e.g. the splitting of data cubes into individual images. In particular the background trends seen in Figure 4.2 and 4.3 were a sub-array feature and so no frames were required to be removed because of this.

Pixel values were converted from MJy/sr to electrons as for WASP-21. Appropriate masking frames were not supplied for each individual image for WASP-28, so a semi-static pixel mask (which flags permanent and semi-permanent bad pixels and regions) provided for each channel was applied.

Cosmic ray hits, noisy pixels and column pull-down/pull-up were masked using the same technique as for WASP-21. The only difference was that instead of all the pixels in the frame being analysed, only a 50×50 pixel box centred on WASP-28 was analysed. The windows used for the median smoothed light curves in the cosmic ray hit masking were 100 observations in both channels 1 and 2.

The only manual masking that was applied was for channel 1, where the positive residual bias pattern effect, seen in the WASP-21 channel 1 data, was present. As for WASP-21, none of the source apertures used in the aperture photometry sampled this masked region.

4.4.2 WASP-28 aperture photometry

As with the pre-aperture photometry, many aspects of the aperture photometry analysis for WASP-28 were in common with those described for WASP-21 (Section 4.3.2). The following values were selected for the inputs to the IDL photometry software, using the same reasoning as for WASP-21. Note that the values were the same for channels 1 and 2:

centroiding box size: 7×7 pixels

$r_{\text{in}} = 15$ pixels

$r_{\text{out}} = 25$ pixels

background sigma-clipping level: 3σ




Figure 4.11: A plot of the WASP-28 channel 1 raw flux against dx and dy and also against time. Clear correlations of the raw flux with x and y positions are present, while the $\sim 1\%$ variation of flux in time can also be seen.

Using these inputs, the IDL aperture photometry software was run on the data for each channel using source apertures from $r_{\text{source}} = 2\text{--}6$ pixels, in steps of 0.5. As for WASP-21, the values extracted from each frame were: the HMJD(UTC) time at mid-exposure; the background subtracted source flux; the photometric error and the x and y centroid position from the Gaussian fit. Sigma clipping of these flux and x and y values was again carried out at the 4σ level.

The times were converted from HMJD(UTC) to HJD(TDB), with the conversion from UTC to TDB applied by adding 67.184 s to the times for both channels. The light travel time effect was corrected for by subtracting 45.836 s from the times in both channels.

A ramp in the background values of both channels was found, like that found for WASP-21, and as a result 200 points were removed from the start of the channel 1 data, and 50 were removed from the start of the channel 2 data. For the final datasets (for which $r_{\text{source}} = 2.5$ pixels, see Section 4.4.5), channel 1 contained 3712 measurements and channel 2 contained 3867 measurements.

Figures 4.11 and 4.12 show the raw fluxes extracted from the aperture photometry (using $r_{\text{source}} = 2.5$ pixels) against the x and y PSF centroid and also against time. For channel 1, clear trends in the flux with x and y position can be seen that give rise to the $\sim 1\%$ systematic variations in the raw flux with time. In channel 2, like for WASP-21, no clear correlations between the raw flux and x and y posi-




Figure 4.12: An equivalent plot to Figure 4.11, except here for the channel 2 data for WASP-28. As for the channel 2 data for WASP-21, there are no obvious correlations between the x and y positions and the flux. However, simple detrending functions are favoured in the tests to remove the intra-pixel sensitivity variations, as described in Section 4.4.4.

tions are seen, but such trends are favoured in the tests to remove the intra-pixel sensitivity variations from the data (see Section 4.4.4).

4.4.3 WASP-28 MCMC set up

The MCMC analysis for WASP-28, using `MCMCTRANSIT`, was carried out as described for WASP-21 (Section 4.3.3), except for the differences noted here. Ephemeris and light curve priors were taken from Anderson et al. [2014]. The transit light curve priors were applied to the quantities δ , b and T_{14} , in accordance with the parameters used by Anderson et al. [2014]. As for WASP-21, the T_{eff} and $[\text{Fe}/\text{H}]$ prior values were taken from the study of exoplanet host stars by Mortier et al. [2013]. The values used for these priors were:

$$T_0 = \text{BJD}(\text{TDB}) 2455290.40596 \pm 0.00031$$

$$P = 3.4088300 \pm 0.000006 \text{ d}$$

$$\delta = 0.01300 \pm 0.00027$$

$$b = 0.21 \pm 0.1$$

$$T_{14} = 0.1349 \pm 0.0010 \text{ d}$$

$$T_{\text{eff}} = 6134 \pm 34 \text{ K}$$

$$[\text{Fe}/\text{H}] = -0.12 \pm 0.03$$

Note that Anderson et al. [2014] give T_0 in BJD(UTC), so the appropriate correction of 66.184 s was added to this quantity to convert to BJD(TDB). Again, the difference between HJD and BJD was considered negligible. Also note that, as for WASP-21, $[\text{Fe}/\text{H}]$ for WASP-28 has increased from the initial estimate used to motivate this study.

Radial velocity data were also taken from the study by Anderson et al. [2014], with measurements coming from the CORALIE and HARPS spectrographs. Timing corrections were applied as usual, with the UTC to TDB correction applied by adding 67.184 s to the times and the light travel time correction applied by subtracting 22.918 s from the times.

Tests of the robustness of the parameter solution for this object were carried out, using 6 independent MCMC runs and checking the Gelman-Rubin statistic [Gelman & Rubin, 1992] for the jumps parameters. As for WASP-21, the statistic for each parameter implied that the 6 MCMC chains had converged and were well mixed. For these runs, equation 4.1 was used for the systematic detrending in both channels and 10^4 jumps were again used in the production runs.

4.4.4 WASP-28 detrending functions tests

With the set-up described above, I ran tests to determine which detrending functions optimised the secondary eclipse fits in each channel, using the BIC to distinguish between different model fits.

Table 4.3 gives the results of MCMC runs using combinations of the best detrending functions (which were determined separately for each channel). The results show that for both channel 1 and channel 2 the favoured detrending function was:

$$F_{\star, \text{model}} = \hat{F}_{\star} + a_0 + a_x dx + a_y dy \quad (4.7)$$

In channel 1 the eclipse depths are in good agreement between the two best detrending functions chosen for that channel. The different detrending functions used for channel 2 show a poorer agreement between the eclipse depths, and the difference in the BIC for the best two functions (1.9) does not strongly support one function over the other (see Table 2.3). However, because these eclipse depths are consistent at 1σ and the linear model is the best fit, I chose this detrending function for the final analysis. Therefore, for the subsequent tests of the optimal source aperture radius

the detrending function given in equation 4.7 was used for both channels.

Table 4.3: Eclipse depths and BIC values for fits to the WASP-28 channel 1 and 2 data using different combinations of detrending functions. The detrending functions chosen for these combined tests were those that performed the best in MCMC fits to individual datasets. The eclipse depths for channel 1 are consistent between a model that is linear in x and y and an alternate model that has an additional quadratic term in x . For channel 2 there is a slight difference in the eclipse depths for the best two models: one that is linear in x and y and another which has an additional quadratic term in y . The difference in the BIC (1.9) does not strongly support one function over the other. However, because the eclipse depths are consistent at 1σ and the linear model is the best fit I chose this detrending function for the final analysis.

Channel 1		Channel 2	
Detrending function ($-\hat{F}_*$)	Eclipse depth (%)	Detrending function ($-\hat{F}_*$)	Eclipse depth (%)
$a_0 + a_x dx + a_y dy$	0.069 ± 0.017	$a_0 + a_x dx + a_y dy$	0.113 ± 0.021
$a_0 + a_x dx + a_y dy$	0.066 ± 0.018	$a_0 + a_y dy + a_{yy} dy^2$	0.133 ± 0.023
$a_0 + a_x dx + a_y dy$	0.067 ± 0.018	$a_0 + a_x dx + a_y dy + a_{yy} dy^2$	0.135 ± 0.024
$a_0 + a_x dx + a_x dx^2 + a_y dy$	0.069 ± 0.017	$a_0 + a_x dx + a_y dy$	0.115 ± 0.023
$a_0 + a_x dx + a_x dx^2 + a_y dy$	0.066 ± 0.018	$a_0 + a_y dy + a_{yy} dy^2$	0.133 ± 0.025
$a_0 + a_x dx + a_x dx^2 + a_y dy$	0.066 ± 0.018	$a_0 + a_x dx + a_y dy + a_{yy} dy^2$	0.133 ± 0.023
			BIC = $\chi^2 + k \ln N$
			7420.5
			7422.4
			7427.5
			7429.6
			7431.2
			7436.0




Figure 4.13: Plots of the normalised secondary eclipse residual RMS and eclipse depth values for WASP-28, in channels 1 and 2, as a function of the source aperture radius. In both channels, the residual RMS value is minimised for $r_{\text{source}} = 2.5$ pixels. In channel 1 there is a trend of eclipse depth with r_{source} , but there is agreement with the optimal eclipse depth at 1σ across all r_{source} values. A similar trend can be seen in channel 2, with very marginal disagreement with the optimal eclipse depth occurring at $r_{\text{source}} = 5$ and 6 pixels, where the RMS has degraded substantially.

4.4.5 WASP-28 aperture radius tests

With the optimal detrending functions selected, the optimal r_{source} values were then determined. MCMC runs were carried out for datasets extracted using $r_{\text{source}} = 2\text{--}6$ pixels and in each channel the optimal r_{source} value was found as that which minimised the RMS on the normalised residuals of the secondary eclipse fit for that channel. As for WASP-21, the errors on the secondary eclipse data were rescaled for these runs.

Figure 4.13 shows the results of these tests. For channel 1 the normalised residual RMS is minimised for $r_{\text{source}} = 2.5$ pixels. There is some structure in the eclipse depths as a function of r_{source} , but all values agree with the $r_{\text{source}} = 2.5$ eclipse depth to better than 1σ , so I adopt this r_{source} value. In channel 2 the




Figure 4.14: An equivalent plot to Figure 4.7, but here shown for the WASP-28 channel 1 and 2 data. In both channels, the binned RMS values follow the white noise expectation down to the largest bin sizes.

optimal r_{source} value is again 2.5 pixels. Structure in the eclipse depths is present at a similar level to channel 1, with the most different eclipse depths occurring at large r_{source} , where the RMS values have significantly degraded. As in channel 2 for WASP-21, the eclipse depth for $r_{\text{source}} = 2.0$ pixels is somewhat discrepant, perhaps due to systematic effects associated with using too small a source aperture.

Figure 4.14 shows the tests of red noise in the light curve residuals for a range of source aperture radii. In both channels, the binned RMS values follow the white noise expectation down to the largest bin sizes. Therefore, there was no motivation to increase the secondary eclipse uncertainties e.g. by using a prayer bead analysis. The uncertainty values outputted from the MCMC parameter estimation were adopted.

4.4.6 WASP-28 optimal MCMC analysis

Using the detrending functions and aperture radii selected in Sections 4.4.4 and 4.4.5 I carried out a longer MCMC run, using 10^5 steps in the production run, to ensure the parameter space was being fully explored. The optimal parameter values and uncertainties used in the results section (Section 4.6 and Table 4.5) were taken directly from the output of this MCMC run; no additional error assessments were made. The errors on the channel 1 eclipse data were scaled by 0.98, while for channel 2 no rescaling was necessary, implying the secondary eclipse models provide a good

fit to the data. Figure 4.20 shows channel 1 and 2 binned time series for the raw fluxes, the eclipse light curves and the residuals associated with this final MCMC analysis. The final eclipse depths derived for WASP-28 were $0.077 \pm 0.017\%$ (for channel 1) and $0.112 \pm 0.022\%$ (for channel 2).

4.5 Analysis of WASP-37

4.5.1 WASP-37 pre-aperture photometry

The pre-aperture photometry analysis for WASP-37 was almost identical to that described for WASP-28 (Section 4.4.1). The only significant difference in the analysis was the manually masked regions. In both channels a star roughly 25 pixels from WASP-37 was masked. Additionally, in channel 1, as for WASP-21 and WASP-28, the positive residual bias pattern effect caused a systematic brightening of the corner of the chip (near to WASP-37) and so this region was masked. In this channel there was also a streak of brightened pixels just below WASP-37's PSF which was also masked. This artifact is known as a *positive slew residual* and is caused by the presence of a bright star being on the detector while the telescope is slewing between target positions⁵. The presence of this masked region limited the usable source apertures in the aperture photometry to sizes of $r_{\text{source}} = 5$ or less. In channel 2, additional masking was applied to a permanently pulled up column in the background region.

4.5.2 WASP-37 aperture photometry

As with the pre-aperture photometry, the aperture photometry analysis for WASP-37 was almost identical to that for WASP-28 (Section 4.4.2). The inputs to the IDL aperture photometry software were as given for WASP-28, with 4σ clipping also being applied to the extracted flux and x and y values.

For the conversion of times from the UTC to TDB timing standard a value of 66.184s was added to the extract times, while for the light travel time correction a value of 45.764s was subtracted.

The background effects seen in WASP-28 were also seen for WASP-37. Ramps in the background values led to 200 points being removed from the start of the channel 1 data and 50 points being removed from the start of the channel 2 data. For the final datasets (with $r_{\text{source}} = 2.0$; see Section 4.5.5), channel 1 contained 3664 measurements and channel 2 contained 3813 measurements.

⁵<http://irsa.ipac.caltech.edu/data/SPITZER/docs/irac/warmfeatures/>




Figure 4.15: A plot of the WASP-37 channel 1 raw flux against dx and dy position and also against time. A correlation of the raw flux with the y centroid position can be seen. Unlike WASP-21 and WASP-28 in channel 1, the flux variations in time are not so clear here, due to an increased point-to-point scatter on the raw flux data.

Figures 4.15 and 4.16 show the raw fluxes extracted from the aperture photometry (using $r_{\text{source}} = 2.0$ pixels) against the x and y PSF centroid position and also against time. In channel 1, a clear trend of the raw flux with y position can be seen, while in channel 2 there are no clear correlations. As with WASP-21 and WASP-28, tests of the removal of intra-pixel sensitivity variations (see Section 4.5.4) suggest there are such correlations, but we are unable to see them here due to the large point-to-point scatter in this channel.

4.5.3 WASP-37 MCMC set up

The MCMC analysis for WASP-37 followed the same form as those for WASP-21 and WASP-28. Ephemeris and light curve priors were taken from the discovery paper for this object, by Simpson et al. [2011]. Transit light curve priors were applied to the quantities δ , b and T_{14} , in accordance with the parameters used by Simpson et al. [2011]. Prior values for T_{eff} and $[\text{Fe}/\text{H}]$ were also taken from the spectroscopic analyses of this study (no values were given for WASP-37 in the more recent study of host stars by Mortier et al. [2013]). The values used for these priors were:

$$T_0 = \text{HJD(TDB)} 2455338.6196 \pm 0.0006$$

$$P = 3.577469 \pm 0.000011 \text{ d}$$




Figure 4.16: An equivalent plot to Figure 4.15, but for channel 2 data. Similarly to WASP-21 and WASP-28 there are no obvious correlations between flux and the x and y positions, but detrending functions containing positional terms are favoured in the tests in Section 4.5.4.

$$\delta = 0.01427 \pm 0.00027$$

$$b = 0.198 \pm 0.130$$

$$T_{14} = 0.1304 \pm 0.0018 \text{ d}$$

$$T_{\text{eff}} = 5800 \pm 150 \text{ K}$$

$$[\text{Fe}/\text{H}] = -0.40 \pm 0.12$$

T_0 is given by Simpson et al. [2011] in HJD(UTC) and so a correction of 66.184s was added to convert to the TDB timing system.

Radial velocity data were also taken from Simpson et al. [2011], with the measurements coming from the SOPHIE and CORALIE instruments. Timing corrections were applied as usual, with the UTC to TDB correction applied by adding 66.184s to the times and the light travel time correction applied by subtracting 22.882s.

Tests of the robustness of the parameter solution for this object were then carried out as for WASP-21 and WASP-28, using 6 independent MCMC runs and checking the Gelman-Rubin statistic for the jump parameters. The statistics for each parameter confirmed that the 6 MCMC chains had converged and were well

mixed. Again, equation 4.1 was used for the detrending of the intra-pixel sensitivity variations in both channels and 10^4 jumps were used in the production runs.

4.5.4 WASP-37 detrending functions tests

With this set-up, I ran tests to determine which detrending function optimised the secondary eclipse fits in each channel, in the same way as for WASP-21 and WASP-28. Table 4.4 gives the results of MCMC runs using combinations of the best detrending functions. The results show that for channel 1 the favoured detrending function is:

$$F_{\star,\text{model}} = \hat{F}_{\star} + a_0 + a_x dx + a_y dy + a_{xx} dx^2 + a_{yy} dy^2, \quad (4.8)$$

while for channel 2 the favoured function is:

$$F_{\star,\text{model}} = \hat{F}_{\star} + a_0 + a_x dx, \quad (4.9)$$

Channel 1 shows some variation in eclipse depth between the detrending functions that provide reasonable fits. Similarly to WASP-21 in channel 2, the detrending function that is quadratic in x and y gives a slightly larger eclipse depth than that which is linear in x and y . The two do agree at 1σ and since the quadratic function is favoured by the BIC ($\Delta\text{BIC} = 5.4$; see Table 2.3) it was used in subsequent tests. In channel 2 there is good agreement, in terms of the eclipse depths, between the three detrending functions tested.

Table 4.4: Eclipse depths and BIC values for fits to the WASP-37 channel 1 and 2 data using different combinations of detrending functions. The channel 1 eclipse depths are consistent to within 1σ for the three functions assessed here. The largest difference to the best fit eclipse depth is for the function that is linear in x and y . This fit has a $\Delta\text{BIC} = 5.4$ compared to the optimal fit and since the eclipse depths are consistent, I select the quadratic function for further analysis. For channel 2 there is good agreement across the three detrending functions assessed.

Channel 1		Channel 2	
Detrending function ($-\hat{F}_s$)	Eclipse depth (%)	Detrending function ($-\hat{F}_s$)	Eclipse depth (%)
$a_0 + a_x dx + a_{xx} dx^2 + a_y dy + a_{yy} dy^2$	0.098 ± 0.023	$a_0 + a_x dx$	0.088 ± 0.028
$a_0 + a_x dx + a_{xx} dx^2 + a_y dy + a_{yy} dy^2$	0.102 ± 0.023	$a_0 + a_x dx + a_y dy$	0.083 ± 0.027
$a_0 + a_x dx + a_{yy} dy$	0.078 ± 0.023	$a_0 + a_x dx + a_y dy$	0.083 ± 0.030
$a_0 + a_x dx + a_y dy$	0.078 ± 0.022	$a_0 + a_x dx$	0.086 ± 0.028
$a_0 + a_x dx + a_{xx} dx^2 + a_y dy$	0.083 ± 0.023	$a_0 + a_x dx + a_y dy$	0.089 ± 0.028
$a_0 + a_x dx + a_{xx} dx^2 + a_y dy$	0.081 ± 0.023	$a_0 + a_x dx$	0.089 ± 0.027
$a_0 + a_x dx + a_{xx} dx^2 + a_y dy + a_{yy} dy^2$	0.099 ± 0.024	$a_0 + a_x dx + a_y dy + a_{yy} dy^2$	0.080 ± 0.027
$a_0 + a_x dx + a_y dy$	0.080 ± 0.023	$a_0 + a_x dx + a_y dy + a_{yy} dy^2$	0.086 ± 0.028
$a_0 + a_x dx + a_{xx} dx^2 + a_y dy$	0.083 ± 0.022	$a_0 + a_x dx + a_y dy + a_{yy} dy^2$	0.081 ± 0.028
			BIC = $\chi^2 + k \ln N$
			7419.8
			7420.7
			7425.2
			7426.5
			7427.2
			7429.3
			7435.1
			7438.3
			7440.3




Figure 4.17: Plots of the normalised secondary eclipse residual RMS and eclipse depth values for WASP-37, in channels 1 and 2, as a function of the source aperture radius. In both channels the optimal source aperture radius is $r_{\text{source}} = 2.0$ pixels and there is good agreement across the radii tested.

4.5.5 WASP-37 aperture radius tests

Source aperture radius tests were then carried out for WASP-37. Because of the positive slew residual systematic, noted in Section 4.5.1, only r_{source} values up to 5 pixels were tested. This was not a concern, since the RMS tends to be significantly worse for $r_{\text{source}} > 5$ pixels in comparison to smaller apertures.

Figure 4.17 shows the results of these tests. In both channels the eclipse depths are stable across the range of radii tested, suggesting robust secondary eclipse solutions. The normalised residual RMS is minimised for $r_{\text{source}} = 2.0$ pixels in both channels. In the previous objects analysed, the eclipse depths for $r_{\text{source}} = 2.0$ pixels have been marginally discrepant, so this result was treated with some caution. Visual inspection of the raw flux derived with different r_{source} values show that the systematic variability is not changed significantly at the smaller radii (for both channels). This, combined with the facts that light curve residuals are close to white at smaller




Figure 4.18: An equivalent plot to Figure 4.7, but here shown for the WASP-37 channel 1 and 2 data. The black line is for the optimal source aperture size $r_{\text{source}} = 2$, while the red, blue and green lines are for $r_{\text{source}} = 3, 4$ and 5 pixels, respectively. In channel 1, the binned RMS values follow the $1/\sqrt{n}$ white noise expectation for all but the largest bin sizes, except for $r_{\text{source}} = 5$ pixels. In channel 2, the binned RMS values follow the white noise expectation down to the largest bin sizes.

r_{source} values (see Figure 4.18) and that the eclipse depths are stable with r_{source} , suggests that the $r_{\text{source}} = 2.0$ eclipse depths are trustworthy. I therefore proceed using these results, with the eclipse depth uncertainties determined by the MCMC.

4.5.6 WASP-37 optimal MCMC analysis

The detrending functions and aperture radii selected in Sections 4.5.4 and 4.5.5 were used in an MCMC run using 10^5 steps in the production run. As for WASP-28, the optimal parameter values and uncertainties used in the results section were taken directly from the output of this MCMC run. The photometric errors on both the channel 1 and channel 2 eclipse data were scaled by a factor 0.99, suggesting the secondary eclipse model provides a good fit to both datasets. Figure 4.20 shows channel 1 and 2 binned time series for the raw fluxes, the eclipse light curves and the residuals associated with this final MCMC analysis. The final eclipse depths derived for WASP-37 were $0.097^{+0.023}_{-0.025}$ % (for channel 1) and 0.090 ± 0.030 % (for channel 2).

Table 4.5: Selected results from the MCMC analyses for WASP-21, WASP-28 and WASP-37, as presented in Sections 4.3.7, 4.4.6 and 4.5.6. For WASP-21 results are derived from a fit to the channel 2 data only, except for $\Delta F_{3.6\mu\text{m}}$ and $T_{3.6\mu\text{m}}$ which come from the upper limit place on the eclipse depth in channel 1. Values for WASP-28 and WASP-37 come from joint fits to the channel 1 and 2 data.

Parameter	Symbol	WASP-21	WASP-28	WASP-37	Unit
3.6 μm eclipse depth	$\Delta F_{3.6\mu\text{m}}$	< 0.11	0.077 ± 0.017	$0.097_{-0.025}^{+0.023}$	%
4.5 μm eclipse depth	$\Delta F_{4.5\mu\text{m}}$	0.064 ± 0.014	0.112 ± 0.022	0.090 ± 0.030	%
3.6 μm brightness temperature	$T_{3.6\mu\text{m}}$	< 1680	1420 ± 110	1460 ± 130	K
4.5 μm brightness temperature	$T_{4.5\mu\text{m}}$	1200 ± 100	1410 ± 120	1230 ± 160	K
3.6–4.5 μm slope	ζ	> -0.097	-0.002 ± 0.031	-0.051 ± 0.042	$\% \mu\text{m}^{-1}$
	$e \cos \omega$	$0.0002_{-0.0014}^{+0.0015}$	$-0.0020_{-0.0022}^{+0.0019}$	$0.0015_{-0.0020}^{+0.0043}$	
Eclipse phase	ϕ_{E}	0.5001 ± 0.0009	$0.4987_{-0.0014}^{+0.0012}$	$0.5009_{-0.0013}^{+0.0027}$	
Orbital separation	a	0.0531 ± 0.0002	0.0459 ± 0.0001	0.0458 ± 0.0005	au
Orbital period	P	4.322519 ± 0.000003	3.408830 ± 0.000006	3.577469 ± 0.000011	days
Planetary radius	R_{p}	1.22 ± 0.04	1.25 ± 0.04	1.19 ± 0.05	R_{J}
Planetary mass	M_{p}	0.31 ± 0.01	0.93 ± 0.05	1.89 ± 0.09	M_{J}
Stellar radius	R_{\star}	1.19 ± 0.03	1.12 ± 0.03	1.03 ± 0.04	R_{\star}
Stellar mass	M_{\star}	1.07 ± 0.01	1.11 ± 0.01	1.00 ± 0.04	M_{\star}
Stellar effective temperature	T_{eff}	5920 ± 60	6130 ± 40	5800 ± 150	K
Stellar metallicity	[Fe/H]	-0.22 ± 0.04	-0.12 ± 0.03	-0.40 ± 0.12	dex

4.6 Results

The final results from the MCMC analyses for WASP-21, WASP-28 and WASP-37 are presented in Table 4.5, while binned versions of the raw fluxes, secondary eclipse light curves and the residuals are shown in Figures 4.19–4.21.

Secondary eclipse depths and brightness temperatures

The main result from this chapter is that significant detections of planetary thermal emission have been made in five of the six secondary eclipse datasets analysed. The depths of these secondary eclipses, which represent day-side planet-to-star flux ratios, are: $\Delta F_{4.5\mu\text{m}} = 0.064 \pm 0.014\%$ (for WASP-21b); $\Delta F_{3.6\mu\text{m}} = 0.077 \pm 0.017\%$ and $\Delta F_{4.5\mu\text{m}} = 0.112 \pm 0.022\%$ (for WASP-28b), and $\Delta F_{3.6\mu\text{m}} = 0.097_{-0.025}^{+0.023}\%$ and $\Delta F_{4.5\mu\text{m}} = 0.090 \pm 0.030\%$ (for WASP-37b). The only dataset for which a significant secondary eclipse detection was not made was for WASP-21b in channel 1. As described in Section 4.3.7, I have placed an upper limit on the eclipse depth here of $\Delta F_{3.6\mu\text{m}} < 0.11\%$.




Figure 4.19: Binned time series for WASP-21 in channels 1 and 2 of IRAC, from the final analysis of these data as described in Section 4.3.7. The top row shows the raw flux light curves for the optimal source aperture radii ($r_{\text{source}} = 2.5$ pixels) along with the full secondary eclipse model, which includes the detrending functions found in Section 4.3.4. The middle row shows the fluxes with the detrending functions removed and normalised to the flux of the star. Residuals are shown in the bottom row. Note that only an upper limit has been placed on the eclipse depth in channel 1 - the light curves are shown here only for illustration. In channel 2 the eclipse depth was found to be $0.064 \pm 0.014\%$.




Figure 4.20: An equivalent plot to Figure 4.19, but for the analysis of WASP-28 presented in Section 4.4.6. The eclipse depths found from these fits were $0.077 \pm 0.017\%$ and $0.112 \pm 0.022\%$ for channel 1 and 2, respectively.




Figure 4.21: An equivalent plot to Figures 4.19 and 4.20 but for the analysis of WASP-37 presented in Section 4.5.6. The eclipse depths found from these fits were $0.097^{+0.023}_{-0.025}$ % and 0.090 ± 0.030 % for channel 1 and 2, respectively.

As described in Section 2.6.1, these eclipse depths can be used to derive planetary brightness temperature estimates. Following the method described there, I find brightness temperature estimates of: $T_{3.6\mu\text{m}} < 1680$ K and $T_{4.5\mu\text{m}} = 1200 \pm 100$ K (for WASP-21b); $T_{3.6\mu\text{m}} = 1420 \pm 110$ K and $T_{4.5\mu\text{m}} = 1410 \pm 120$ K (for WASP-28b) and $T_{3.6\mu\text{m}} = 1460 \pm 130$ K and $T_{4.5\mu\text{m}} = 1230 \pm 160$ K (for WASP-37b). For comparison, the equilibrium temperatures for these planets (with $A_B = 0$ and $\varepsilon = 1$) are: 1350 K for WASP-21b; 1460 K for WASP-28b and 1330 K for WASP-37b. The brightness temperatures I have derived are comparable to these values, suggesting plausible eclipse depths for the three objects have been found. I note here that the upper limit on the brightness temperature for WASP-21b in channel 1 ($T_{3.6\mu\text{m}} < 1680$ K) is close to the maximum effective temperature allowed for this planet, $T_d = 1730$ K (using equation 2.37 with $A_B = 0$ and $\varepsilon = 1$).

Eccentricity constraint

Previous analyses of the radial velocity data from the three systems studied in this chapter have shown no evidence for significantly eccentric orbits [Bouchy et al., 2010; Anderson et al., 2014; Simpson et al., 2011]. I can test this further, because the combination of radial velocity and secondary eclipse data I have analysed for each of the objects places constraints on the eccentricity of the system (as described in Section 2.4.2.3). Specifically, the central phase of the secondary eclipse constrains the quantity $e \cos \omega$ through [Charbonneau et al., 2005]:

$$\Delta\phi \equiv \phi_E - 0.5 \simeq \frac{2}{\pi} e \cos \omega, \quad (4.10)$$

Figure 4.22 shows the distribution of the argument of periastron (ω , which defines the orientation of the orbit) against the distribution of the eccentricity, resulting from the final MCMC analysis for each object. In each case, the $e \cos \omega$ value is constrained to be small (see Table 4.5) meaning that in general the eccentricity of these systems can only be large if orientated in a narrow range of ω .

For WASP-21, the $e \cos \omega$ value was determined using the fit to the channel 2 data only. Given that the channel 1 eclipse could not be found reliably and the residuals from the best fit displayed significant amounts of red noise (Figure 4.7), I chose not to use the channel 1 data to constrain $e \cos \omega$. With $e \cos \omega = 0.0002_{-0.0014}^{+0.0015}$, the secondary eclipse timing is not significantly different to the circular orbit expectation. I found that for e to be larger than 0.03, ω was restricted to the narrow range of $87.7 < |\omega| < 92.1$, at 1σ . I also performed a Lucy-Sweeney test [Lucy & Sweeney, 1971], which is an F-test where the eccentric fit is tested against a more simple cir-

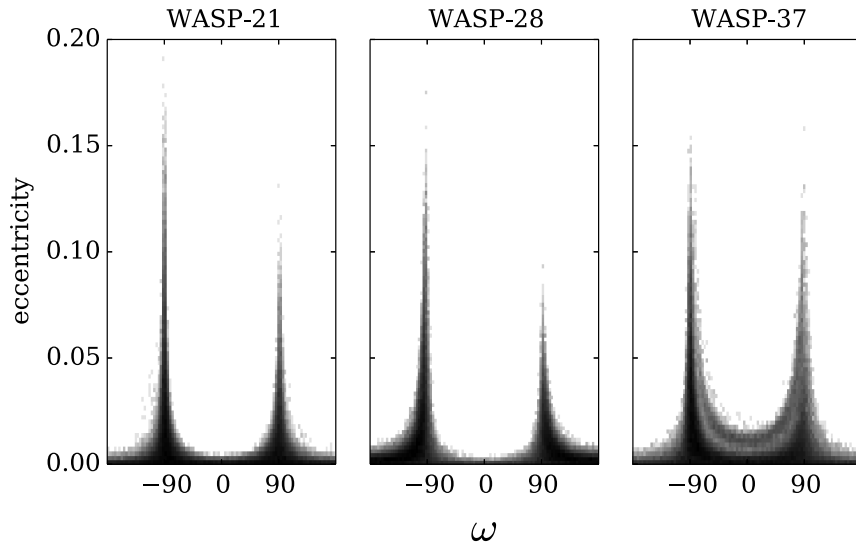


Figure 4.22: 2D histograms of the MCMC chain values for e and ω in the final analyses of WASP-21, WASP-28 and WASP-37 (from left to right, respectively). In each case the eccentricity is restricted to small values, except for when $|\omega| \sim 90^\circ$. This is a result of the small differences between the measured eclipse central phases (see Table 4.5) and the expectation from a circular orbit of $\phi_E = 0.5$.

cular fit. In the case of WASP-21, the circular orbit fit cannot be ruled out. Hence, the evidence for this system is that it is entirely consistent with a circularised orbit.

For WASP-28 and WASP-37, simultaneous fits to the secondary eclipse data in both IRAC channels were used to constrain $e \cos \omega$. For WASP-28, I obtain a value of $e \cos \omega = -0.0020^{+0.0019}_{-0.0022}$, suggesting a very marginal offset of the eclipse timing from the circular orbit expectation. The restriction on ω for $e > 0.03$ reflects this, with $91.1 < |\omega| < 97.9$ (at 1σ). The Lucy-Sweeney test for WASP-28 implies that, like WASP-21, a circular orbit cannot be ruled out. For WASP-37, the constraint on $e \cos \omega$ is not as strong, as can be seen in Figure 4.22 where a larger range of eccentricities are allowed. Here, $e \cos \omega = 0.0015^{+0.0043}_{-0.0020}$, with ω restricted to $80.7 < |\omega| < 91.1$, for $e > 0.03$ (at 1σ). Again, the Lucy-Sweeney test implies the results are consistent with that of a circular orbit.

All of these conclusions are consistent with the previous eccentricity analyses for these systems.

4.7 Discussion

Albedo and energy redistribution

Using the brightness temperatures derived in Section 4.6, I estimated the day-side effective temperature of each of the planets studied in this chapter following the method of Cowan & Agol [2011], as described in Section 2.6.2. The values found were: $T_d = 1200 \pm 190$ K for WASP-21b (using only the channel 2 brightness temperature); $T_d = 1410 \pm 170$ K for WASP-28b and $T_d = 1420 \pm 220$ K for WASP-37b.

In the context of the work by Cowan & Agol [2011] (described in Section 1.3.2.2), these planets fall into the group of less highly irradiated planets that show a range of T_d/T_0 values and as such do not provide much leverage. Nevertheless, with $T_{\varepsilon=0}$ in the range 1700 – 1900 K, the T_d/T_0 values for my objects of 0.63 ± 0.10 (for WASP-21b), 0.68 ± 0.08 (for WASP-28b) and 0.76 ± 0.12 (for WASP-37b), are consistent with what has been found for other planets with similar irradiation levels.

Strong constraints cannot be placed on A_B and ε for these planets, given their moderate T_d/T_0 values and sizeable uncertainties (see equation 2.37). As an example of this Figure 4.23, shows the 2D PDF in A_B and ε for WASP-28b (calculated using the method described in Section 2.6.2). It shows that only very high albedo values can be ruled out, while ε is unconstrained. It is interesting to note, however, that very low values for both A_B and ε are not favoured for this planet.

The 3.6–4.5 μm slope

As described in Sections 1.3.2.3 and 3.5.2, Knutson et al. [2010] introduced a model-independent empirical measure of the 3.6–4.5 μm slope in planetary spectra, with respect to a best fitting blackbody. This has been found to correlate with the presence or absence of temperature inversions, because planetary CO and H₂O features sampled in the 4.5 μm band switch between being in absorption (for the non-inverted atmospheres) and being in emission (for the inverted atmospheres). Following Anderson et al. [2011], I calculated ζ values (as defined in Section 3.5.2) for WASP-21b, WASP-28b and WASP-37b. Respectively, the measured values were $\zeta > -0.097\% \mu\text{m}^{-1}$ (as a lower limit due to the upper limit on the channel 1 eclipse depth), $\zeta = -0.002 \pm 0.031\% \mu\text{m}^{-1}$ and $\zeta = -0.051 \pm 0.042\% \mu\text{m}^{-1}$. Strongly negative (positive) ζ values imply that CO and H₂O features are being seen in absorption (emission). The values for the three planets studied here imply neither strong emission nor absorption.

In the introduction to this chapter I speculated that in low metallicity systems the unknown, inversion-driving absorber could be deficient in the planetary




Figure 4.23: The 2D PDF for WASP-28b found using its T_d/T_0 value with equation 2.39, assessed over a grid of allowable albedo and ε values. Darker regions are of higher probability than lighter regions. The distribution shows the full range of ε values are allowable by the data its heat redistribution properties are unconstrained. In terms of the Bond Albedo, only very high values ($A_B > 0.9$) are ruled out, while it is interesting to note that very low values for both A_B and ε are not favoured.

atmosphere, and that we would preferentially see strong absorption for these planets. A hint of a trend of ζ against $[\text{Fe}/\text{H}]$ for the existing data seems to qualitatively agree with this. The low metallicity systems sampled in this study clearly had the potential to consolidate such a trend, however my ζ values for WASP-21b, WASP-28b and WASP-37b do not provide support for it. This can be seen clearly in Figure 4.24, where the ζ values all lie above the best fit linear model to the previous data and are all reasonably consistent with a blackbody (which would have $\zeta = 0$). Additionally, the Pearson correlation coefficient for $[\text{Fe}/\text{H}]$ against ζ , with the addition of the new data (excluding WASP-21b), is $r = 0.26$. This is essentially unchanged from the value of $r = 0.27$ determined from the existing data in Figure 4.1.

It is interesting to see how the planets studied here also fit into the activity-inversion relation, as suggested by Knutson et al. [2010] and discussed in Section 3.5.3. In Figure 4.25 I show $\log R'_{\text{HK}}$ (an index that describes the chromospheric activity of the host star) against ζ , including the systems WASP-21 and WASP-28




Figure 4.24: Equivalent plot to Figure 4.1, but here with the measured ζ values for WASP-21, WASP-28 and WASP-37 found in this chapter. Note that the ζ value for WASP-21 is a lower limit, because only an upper limit was placed on the channel 1 eclipse depth for this object. The grey line is the best fit linear model found for the data, excluding my values. A hint of a correlation between ζ and $[\text{Fe}/\text{H}]$ from the existing data is not supported by the values for the three planets studied in this chapter. The Pearson correlation coefficient for these data (excluding WASP-21b) is $r = 0.26$.

(WASP-37 does not have a reliable $\log R'_{\text{HK}}$ measurement).

My measurement of ζ for WASP-28 is interesting in this context. There are two potential reasons why we might have expected WASP-28b to display strong $4.5\ \mu\text{m}$ absorption *a priori*: the low metallicity and relatively high activity of its host star. The fact that it does not show strong absorption, and is in fact consistent with a blackbody, may be because the inversion-inducing absorber survives in the atmosphere of WASP-28b (because the activity and metallicity of the host are not extreme enough to remove it). Alternatively, there could be a different dependency on metallicity to that speculated above, where the opacity sources that drive the absorption/emission features in the $4.5\ \mu\text{m}$ band are suppressed in low metallicity systems, giving rise to weaker features, rather than a stronger absorption. The consistency of the WASP-28b ζ value with a blackbody is consistent with this. The prediction for future observations would be that low metallicity systems never devi-




Figure 4.25: The correlation between $\log R'_{\text{HK}}$ and ζ as presented in Figure 3.21, but here with additional measurements for WASP-21 and WASP-28 (plotted as blue points). No reliable $\log R'_{\text{HK}}$ measurement exists for WASP-37. The errors on $\log R'_{\text{HK}}$ for WASP-21 and WASP-28 reflect the errors on a single measurement, rather than the range of values found over a few years, as is the case for WASP-3 (highlighted in red). Like WASP-3, WASP-28 orbits a moderately active star, but does not seem to have strong absorption in the $4.5 \mu\text{m}$ band (IRAC channel 2) - its ζ value is consistent with a blackbody. This object, along with WASP-48 (highlighted in green) is discussed in the text. With the addition of the data for WASP-28 and WASP-48, the Pearson correlation coefficient here is $r = -0.58$, suggesting little change from the value of $r = -0.59$ found in Section 3.5.3.

ate far from a blackbody slope, whether or not they exhibit temperature inversions.

The planet WASP-48b [O'Rourke et al., 2014] is also interesting in this context. Due to its high irradiation and quiet host star, this planet was expected to show strong emission at $4.5 \mu\text{m}$. However, as can be seen from Figure 3.5.3, it too is consistent with a blackbody. As for WASP-28, perhaps the low metallicity of this system prevents strong $4.5 \mu\text{m}$ emission being seen, due to a lack of opacity in this band.

4.8 Conclusions

In this chapter I have presented warm *Spitzer* IRAC secondary eclipse observations of three low metallicity systems: WASP-21, WASP-28 and WASP-37. These observations were made at $3.6\ \mu\text{m}$ and $4.5\ \mu\text{m}$ for each object. Five of the six secondary eclipses observed were detected, revealing day-side planet-to-star flux ratios of: $\Delta F_{4.5\mu\text{m}} = 0.064 \pm 0.014\%$ (for WASP-21b); $\Delta F_{3.6\mu\text{m}} = 0.077 \pm 0.017\%$ and $\Delta F_{4.5\mu\text{m}} = 0.112 \pm 0.022\%$ (for WASP-28b), and $\Delta F_{3.6\mu\text{m}} = 0.097^{+0.023}_{-0.025}\%$ and $\Delta F_{4.5\mu\text{m}} = 0.090 \pm 0.030\%$ (for WASP-37b). An upper limit of 0.11% was found at $3.6\ \mu\text{m}$ for WASP-21b. Temperature estimates derived for these planets allow for a range of planetary reflectance and heat redistribution properties, while the timings of the eclipses show the orbits of these systems are consistent with being circular.

Exoplanet systems with existing *Spitzer* IRAC measurements show a hint of a correlation between the host star's metallicity ($[\text{Fe}/\text{H}]$) and the empirical $3.6\text{--}4.5\ \mu\text{m}$ spectral slope of the planet (ζ ; see Figure 4.1). However, the current population of hot Jupiter exoplanets does not cover a wide range of metallicities. The WASP systems analysed in this chapter were chosen to extend the metallicity range of systems with IRAC secondary eclipse measurements, and to put stronger constraints on the weak correlation seen in Figure 4.1. From the flux measurements for the three planets studied here, I found no evidence to support the trend of ζ with host star metallicity.

Chapter 5

Ground-based secondary eclipse observations of WASP-33

5.1 Introduction

The planet WASP-33b was the first planet discovered around an A star and remains one of only two such transiting systems [the other being Kepler-13Ab; Shporer et al., 2014]. The planet was first identified in WASP photometry by Christian et al. [2006], but its host was found to be rapidly rotating, precluding it from the normal mass estimate from RV follow-up analyses. As a result, Collier Cameron et al. [2010] performed a more sophisticated analysis using line-profile tomography, in which the planetary signal was detected travelling through the host’s composite spectral line profile during transit. A joint analysis with follow-up transit photometry revealed a $1.5 R_J$ planet in a retrograde orbit of period 1.22 d ($a = 0.026$ au). Radial velocity data to measure the host’s reflex motion was also taken and, while a definite mass estimate could not be made, a 3σ upper limit of $4.1 M_J$ was found, confirming WASP-33b’s planetary nature. The line-profile tomography analysis of Collier Cameron et al. [2010] also showed evidence of non-radial pulsations in the host star, with the authors suggesting it is a γ Dor-type star. Herrero et al. [2011] subsequently discovered millimagnitude (mmag) photometric oscillations from the star, with a 69 min period. From this, and the stellar properties presented by Collier Cameron et al. [2010], Herrero et al. [2011] re-classified WASP-33 as a δ Scuti star. These stars are short-period (< 0.3 d) pulsators lying in the instability strip, with spectral types ranging from A2–F2. The pulsations can be non-radial and as such multiple resonant modes can be excited, leading to a rich frequency spectrum. Subsequent analyses of long baseline time-series photometry have indeed revealed

multiple pulsation frequencies for WASP-33 [Kovács et al., 2013; von Essen et al., 2014]. The most recent analysis from von Essen et al. [2014] detected 8 distinct frequencies from $8 - 35 \text{ d}^{-1}$ with amplitudes from $0.4 - 1.0 \text{ mmag}$.

In addition to being the only known planet orbiting a δ Scuti star, WASP-33b is also the most highly irradiated planet known - a result of its very tight orbit ($a = 0.026 \text{ au}$) around an A5 star ($R_{\star} = 1.44 R_{\odot}$, $T_{\text{eff}} = 7430 \text{ K}$). With an equilibrium temperature of $T_{\text{d}} = 2700 \text{ K}$ ($A_{\text{B}} = 0$, $\varepsilon = 1$), it is well suited to measurements of its thermal emission through secondary eclipses observations. Several of these have already been made: one at $0.91 \mu\text{m}$ [Smith et al., 2011]; two in the K_{s} band [Deming et al., 2012; de Mooij et al., 2013] and one in each of the $3.6 \mu\text{m}$ and $4.5 \mu\text{m}$ bands of *Spitzer's* IRAC instrument [Deming et al., 2012]. The results have indicated WASP-33b is very hot, with both simple blackbody fits and more detailed modelling finding the planet is highly absorbing and redistributes heat inefficiently [Deming et al., 2012; de Mooij et al., 2013]. Models by Madhusudhan [2012] fitted to the data of Deming et al. [2012] and Smith et al. [2011] do not distinguish clearly between two broad model types: one set have an inverted atmosphere with a solar composition while the other has a non-inverted vertical temperature structure, but requires a carbon-rich atmosphere. The latter is slightly favoured, particularly by the K_{s} band measurement of Deming et al. [2012], which was subsequently confirmed by de Mooij et al. [2013].

In this chapter I present another secondary eclipse measurement for WASP-33b using the ULTRACAM instrument [Dhillon et al., 2007] on the William Herschel Telescope (described in Section 2.2). The eclipse was measured in the three channels of the ULTRACAM instrument simultaneously, using the filters SDSS u', 'Blue continuum' ($\lambda_{\text{c}} = 5149 \text{ \AA}$) and SDSS z'. The use of the ULTRACAM instrument for these observations has many advantages. For example, it is a very stable photometer, with no moving parts, which aids in the accurate removal of systematic effects (including the pulsations of WASP-33) which is crucial for a detection of the secondary eclipse. Since WASP-33 is a bright star ($V = 8.3$), the high duty cycle of ULTRACAM is an advantage because it allows for the detection of many more photons than a standard photometer (suppressing the photon noise). Additionally, the simultaneous, multi-wavelength observations give extra leverage in the fitting of WASP-33's pulsations, as these can be modelled across the three ULTRACAM channels simultaneously.

The removal of these pulsation signals has formed an important part of the previous secondary eclipse detections for WASP-33 and it is typically done by assessing the periodogram of the dataset. In this chapter I have used a slightly different

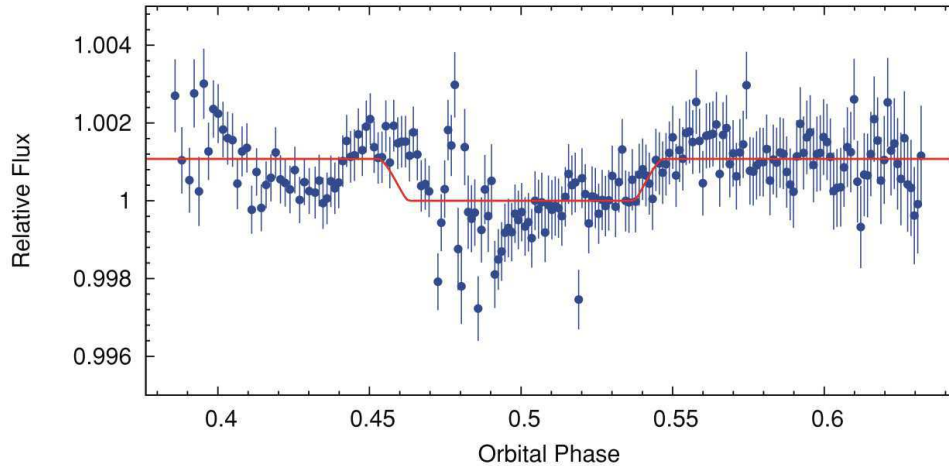


Figure 5.1: The secondary eclipse light curve of WASP-33 in the S[III] band, by Smith et al. [2011], after the removal of the host’s pulsation signals (modelled using a sum of four sinusoids). The resulting planet-to-star flux ratio estimate was $0.109 \pm 0.030\%$. Systematic noise is still clearly visible in the residuals of the eclipse model, perhaps due to residual pulsations or thin cirrus cloud that was present during the observation. The secondary eclipse observations of WASP-33b presented in this chapter will be useful for testing the robustness of this result.

method of fitting the pulsations in the three channels of ULTRACAM using the frequencies from the detailed study by von Essen et al. [2014].

A determination of the z' band eclipse depth is important for a number of reasons. First, with a measured effective temperature estimate of 3290 K [de Mooij et al., 2013], the spectral peak for WASP-33b will be $\lambda_{\max} \sim 880$ nm. The z' band will sample this and provide an important constraint on the planet’s bolometric luminosity. Second, while a similar measurement has already been made in the S[III] band ($\lambda_c = 0.91 \mu\text{m}$), red noise is still clearly present in the pulsation corrected eclipse light curve (see Figure 5.1). A repeat measurement will test the robustness of this eclipse depth.

5.2 Observations

WASP-33 was observed using the ULTRACAM instrument on the William Herschel Telescope (WHT) on 8 October 2012 over a period of 8 hours around the predicted time of secondary eclipse. The ULTRACAM instrument provides simultaneous photometry in three wavelength bands, referred to as the blue, green and red chan-




Figure 5.2: Response functions for the SDSS u' (blue curve), Blue continuum (green curve) and SDSS z' (red curve) filters used in the ULTRACAM secondary eclipse observations of WASP-33b presented in this chapter. These filters were used in the blue, green and red channels of ULTRACAM, respectively. Also shown are the quantum efficiencies of the CCDs used in the blue and green channels (dashed grey line) and the red channel (dotted grey line).

nels. For our observations, these channels were fitted with SDSS u' ($\lambda_c = 3557 \text{ \AA}$, FWHM= 599 \AA), 'Blue continuum' ($\lambda_c = 5149 \text{ \AA}$, FWHM= 158 \AA) and SDSS z' ($\lambda_c = 9097 \text{ \AA}$, FWHM= 1370 \AA) filters, respectively (see Figure 5.2 for the response curves of these filters). The CCDs were set up in a windowed mode, with two pairs of windows of size 230×230 pixels and 300×370 pixels (with a pixel scale of $0.3''/\text{pixel}$). The window sizes, which dictate the exposure time of ULTRACAM (see Section 2.2.2), were set to avoid saturation of WASP-33 in the green channel, whilst still allowing the detection of two comparison stars within $4'$ of the target. The resulting exposure time was 0.977 s .

The telescope was also defocussed to give images with FWHM of around $4''$, which allowed for a longer exposure time (and therefore larger window sizes) while avoiding saturation. Defocussing has the added advantage of reducing the noise introduced by the flat field. For an in-focus stellar image, a lot of weight is put onto a small number of pixels. In this case, errors in the flat field correction combined with movement of the stellar image (due to imperfect telescope tracking) will introduce spurious signals into the light curves. Defocussing decreases this noise source by reducing the weighting on individual pixels.

The relatively narrow 'Blue continuum' filter (see Figure 5.2) was chosen so

that count levels in the red and green channels were comparable. Though the SDSS z' filter is very wide, the quantum efficiency of the red channel CCD is poorer than the green CCD, and in addition WASP-33 is a blue star ($B-V=0.27$). If a wider filter had been used in the green channel (e.g. SDSS g') then the counts would have been very much higher than in the red channel. In order to avoid saturation in the green channel, the exposure time would have been very short and therefore the window sizes would have been very small (hindering the possibility of imaging a suitable comparison star). Since one of the primary focuses of this project was to detect an eclipse in the z' band data, we were happy to use the narrow ‘Blue continuum’ filter to give slightly increased photon noise in the green channel in return for more flexibility in terms of comparison stars.

Although WASP-33 is a blue star, it is still significantly fainter in the u' band of ULTRACAM. As a result we made use of the option to co-add frames in the blue, where the frame transfer and read out process of ULTRACAM is only carried out once for every N frames taken in the red and green channels. This is done to reduce the read out noise in the blue channel data. For our observations of WASP-33 we used $N=4$. Over the 8 hours of observations, 29082 exposures were taken in the green and red channels, while 7270 exposures were taken in the blue channel.

The observations were affected throughout the night by small amounts of cloud, as can be seen in the raw light curves for WASP-33 in Figure 5.4. Also, for the last hour of the observations significant amounts of cloud were present and have clearly affected the data.

Calibration frames were taken before and after the science observations of WASP-33. Flat field frames were taken at both dusk and dawn, however the clouds present at the end of the night also affected the dawn flats, so they were not used in this analysis. Bias frames for the flat fields were taken in the afternoon before the observations, while bias frames for the science data were taken in the afternoon of the following day (since the set-up of the instrument for the science data was not known beforehand).

5.3 Analysis

5.3.1 ULTRACAM Pipeline processing

Calibration frames

The first stage in the analysis of the WASP-33 data was to calibrate the science data and then perform aperture photometry on the resulting frames to obtain the

required flux measurements from the target and the comparison stars.

The first step in the calibration involved creating master bias frames. Bias frames are 0 s exposures taken to measure the offset signal applied during the readout of the CCDs. Due to the fact that the science data was taken in a windowed mode, whilst the flat field frames were not, separate bias frame sets were taken for these two cases (since the bias signal is expected to be different for these different modes). In both cases, separate bias frame sets were taken for each of the ULTRACAM channels, resulting in 6 separate bias frame sets, each containing 101 individual frames. Master bias frames were created for each set by averaging the individual frames using a 4σ clipped mean. This process reduces the effects of read out noise on the master bias. The master biases were then subtracted from individual flat field and science data frames to correct for the bias offset.

Master flat field frames were then created, again one for each channel of the instrument. Flat fields are frames that have been exposed to an illumination that is constant across the field-of-view. Variations in the pixel signals in these frame are then interpreted as variations in the pixel response (e.g. due to varying quantum efficiencies between pixels). Thus, by dividing the science frames by a flat field one can remove the pixel response variations from the science data. The flat fields in this study were measured by taking multiple images of the sky near the zenith around the time of sunset, which provides a suitably evenly illuminated field. Flats were also taken at sunrise, but they were affected by clouds and were not used here.

Master flats were created for each channel in the following way: first, the relevant bias frame was removed from each of the images. Then frames with a mean count level above or below specified limits were removed. The upper limits were set so that non-linear detector effects were avoided, while the lower limits were set to avoid needlessly low counts and significant stellar contamination. These flats were then split into groups of 11 frames with similar mean levels. Within each group the individual frames were normalised to their mean levels and a ‘group master’ was created using median pixel values. Since the field-of-view was continually being rotated with respect to the sky during the flat exposures, this process should have removed any stars present in the flats as they would illuminate different pixels in each frame. The various ‘group master’ frames were then scaled to the mean level of their input groups and added together. The resulting frame was then normalised to its mean level to create the master flat. The final steps here assign a higher weighting to the brighter flats (which have smaller fractional noise) than the fainter flats, giving a better quality master flat.

Setting up apertures

With the calibration frames made, the apertures used for aperture photometry on the science frames were then set up. Stacked images, using 100 individual frames, were created for each channel and apertures were specified on these images for the three stars in the field of view: WASP-33 ($V = 8.1$), a bright comparison star ($V = 9.4$, referred to here as ‘comparison 1’) and a fainter comparison ($V = 11.3$, ‘comparison 2’). Moffat profile fitting was carried out on each of these stars (see Section 2.3.1), in each channel, to provide an initial centroid position in each case.

At this stage I decided on the sizes of the sky annulus radii, which were the same for each object and each channel. The inner sky radius was chosen to be at 35 pixels, such that the sky estimate was not affected by the wings of the defocussed stellar PSFs. This was assessed using visual inspection of the frames and the PSFs from Moffat fitting. The outer sky radius was set such that the uncertainty in the sky estimate did not give a significant contribution to the overall photometric error. However, the outer radius could not be too large otherwise large scale variations in the sky across the CCD could have biased the sky estimate. A value of 50 pixels was chosen to provide a good balance between these two considerations.

Aperture photometry

With the initial aperture positions specified, I ran the REDUCE program (see Section 2.3.1) to extract time-series information from the science data frames, most notably flux measurements of the three stars in each of the three ULTRACAM channels. Moffat fitting was carried out here to determine the centroid positions of each star, as described in Section 2.3.1 (and equation 2.6). A number of user specified parameters could be chosen to customise the running of this program. Some of the key parameter values used here were:

$r_{\text{inner}} = 35 \text{ pix}$	Inner sky radius
$r_{\text{outer}} = 50 \text{ pix}$	Outer sky radius
$\text{profile_fit_fwhm} = 14 \text{ pix}$	Initial FWHM used in Moffat profile fitting
$\text{profile_fit_beta} = 18 \text{ pix}$	Initial beta parameter used in Moffat profile fitting
$\text{profile_fit_hwidth} = 25 \text{ pix}$	Half-width of box used for Moffat profile fitting

Customisation of the last 3 parameters here was important as their values were chosen to reflect the highly defocussed nature of the PSFs (which is not typical for ULTRACAM observations).




Figure 5.3: Signal-to-noise ratios for a typical image from the red (top panel), green (middle) and blue (bottom) ULTRACAM channels. The signal-to-noise is maximised in these channels at $r_{\text{source}} = 20$, 18 and 18 pixels, respectively. The source aperture size I used for the aperture photometry was $r_{\text{source}} = 25$ pixels (marked by the vertical dotted lines), which was chosen to avoid potential systematic effects associated with using a smaller source aperture.

A key choice in the aperture photometry reduction was the size of the source aperture radius, r_{source} . I used fixed apertures throughout, with a size chosen such that it contained most of the stellar flux, without the background significantly contributing to the photometric error budget. Figure 5.3 shows the signal-to-noise ratios for a typical image in each of the three ULTRACAM channels, as a function of source aperture radius. Maxima in the signal-to-noise ratios occur at $r_{\text{source}} = 18$ – 20 pixels. However, after visual inspection of the images and WASP-33’s radial flux profile, I decided to use a slightly larger source aperture, with $r_{\text{source}} = 25$ pixels (marked in Figure 5.3 with the vertical dotted lines). This was done to ensure that systematic effects (e.g. flux spilling out of the aperture due to variable seeing) were minimised, whilst still maintaining a decent signal-to-noise level. This source aperture size was used for each of the three stars in each of the three ULTRACAM channels.

Using this set up, the times, centroid positions, fluxes, flux errors and sky background estimates were extracted for each star in each channel, as described in Section 2.3.1. Before the science frames were evaluated, the calibration frames, described in earlier in this section, were applied to the science data by subtracting the master bias and then dividing by the master flat.

5.3.2 Post-aperture photometry processing

From the output of the REDUCE program, separate files were created for the time-series data in the three channels. These contained the MJD(UTC) times of the observations and the raw fluxes of the three stars along with their errors. In addition the files contained differential fluxes calculated as the WASP-33 flux divided by the flux from comparison 2, along with the propagated errors. These differential fluxes were the key data set used in this study. The reason why comparison 2 was used rather than comparison 1 (or some weighted combination of these stars) was because of a systematic feature that was present in the differential fluxes when using comparison 1 (see Section 5.3.4).

Sigma clipping was then applied to the time-series data in each channel in the following way. Smoothed versions of the flux time-series were found using median filtering, with a window width of 15 min in the red and green channels and 20 min in the blue channel. The smoothed time-series were then subtracted from the original flux. Individual points in this smoothed-subtracted time-series that were more than 3σ from 0 were removed (where σ was found using $1.48 \times \text{MAD}$, Ruppert 2011). This process was carried out iteratively until no more points were clipped. The clipping was applied to the differential flux as well as to the raw fluxes of WASP-33 and comparison 2. If a point was clipped in any of these time-series, it was clipped from them all.

Additionally, it was apparent from the raw light curves (Figure 5.4) that towards the end of the night significant amounts of cloud was present. As a result the final hour of data was removed. Also, some frames were not evaluated by the REDUCE program due to failures in the Moffat profile fitting. 565, 50 and 37 frames were affected in this way in the red, green and blue channels, respectively. Visual inspection of these frames did not show any obvious features that may have caused these failures. However it is likely that they arose from the discrepancy in form between the defocussed PSFs of the stars and the Moffat profile.




Figure 5.4: Raw light curves for WASP-33 in the red, green and blue channels. The grey points are the unbinned raw fluxes, while the coloured points are the binned versions of the sigma-clipped fluxes (with the final hour of data also removed). The effect of clouds in the final hour of the observations is clear. Shorter time-scale variations in the raw flux are suggestive of passing cirrus cloud. The long time-scale curvature in these light curves is a result of atmospheric extinction, which correlates with the airmass (shown in the bottom panel).

5.3.3 Reduced light curves

Figure 5.4 shows the raw fluxes from WASP-33 along with binned versions of the sigma-clipped fluxes, measured in the way described in Sections 5.3.1 and 5.3.2. Also shown is the airmass value for WASP-33 throughout the night. The most obvious feature in the raw fluxes is the poor data at the end of the night, which had clearly been affected by clouds and which were removed. Systematic trends of the flux with airmass are also apparent, especially in the blue channel data, while short time-scale reductions in flux throughout the night are suggestive of passing cirrus clouds. There also appeared to be additional systematic features, e.g. around time 56209.0 (MJD, UTC) in the red channel. These may have resulted from clouds, variations in H_2O absorption (especially in the red channel) or even the δ Scuti pulsations




Figure 5.5: The sky background in the red, green and blue channels of ULTRACAM, measured in an annulus from 35–50 pixels centred on WASP-33. The grey points shows the raw background fluxes, while the coloured points are binned with a width of 8 minutes. The red channel data shows a large amount of variation, on both long and short time-scales. The green and blue channel background light curves are quite different, showing no signs of the short time-scale variability and a different long time-scale trend.

of WASP-33. The clipped and binned fluxes are, unsurprisingly, better behaved with the large deviations due to cloud being suppressed. However, the airmass and absorption/pulsation systematics are still present.

Figure 5.5 shows the sky background estimate in each channel, measured in the annulus around WASP-33. The red channel data here shows a large amount of variation, with short time-scale variability combined with a longer time-scale trend that does not seem to correlate with airmass. In the green and blue channels there appears to be no short time-scale variations, while the long time-scale trends are much weaker and do not share the same shape as in the red channel. Visual inspection of both the raw fluxes and the sky background light curves led me to test whether the raw flux residuals from a simple airmass fit (i.e. $a_0 + a_1X$, where X represents the airmass) correlated with the sky background. I found that in the red channel data this was indeed the case (but see the end of Section 5.3.4), while for




Figure 5.6: Normalised differential light curves for WASP-33, using comparison 2 as a calibrator, after sigma-clipping and removal of the final hour of data. The grey points are binned to around 30 s (for clarity), while the coloured points are binned with a width of 4 minutes. The light curves show that many of the systematic effects seen in the raw light curves (Figure 5.4) have been removed, leaving behind $\sim 0.2\%$ δ Scuti pulsation signals from WASP-33. Note that the mean values for the un-normalised differential light curves were 12, 31 and 63 in the red, green and blue channels, respectively. This highlights the colour difference between WASP-33 ($B - V = 0.27$) and comparison 2 ($B - V = 0.97$).

the green and blue channels no such correlations were found.

A first-order removal of the systematic effects seen in the raw fluxes (except for WASP-33’s pulsations) was carried out by dividing the (un-binned) clipped fluxes for WASP-33 by the corresponding fluxes from comparison 2 (see Section 5.3.4 for why this comparison star was used). Figure 5.6 shows these differential fluxes in the three channels. Clearly much of the noise seen in the raw light curves have been removed in this process. As a result the pulsations of WASP-33 are easily seen, with the variations clearly being correlated across the three bands. The pulsations have amplitudes of $\sim 0.2\%$, similar to those found in previous work on this system [e.g. Herrero et al., 2011; Smith et al., 2011; Deming et al., 2012; de Mooij et al., 2013].

Although Figure 5.6 shows that most of the systematic effects in the raw light curves have been removed, there may well still be smaller amplitude, second order

effects present. For example, there is a linear slope in the blue channel, possibly due to a second-order trend in airmass (caused by the different effective wavelength of the blue channel observations for WASP-33 and comparison 2 due to their colour differences). These potential second order effects are explored in the following Section.

5.3.4 Testing systematic effects

In the raw light curves for WASP-33 (Figure 5.4) it is clear there is a trend in the flux that correlates with airmass. Although much of this effect is removed in the differential light curves (Figure 5.6), second-order trends with airmass are still expected. To characterise this dependence, I used the fluxes from the two comparisons stars, which are assumed not to be variable, in order to avoid any confusing effects from the pulsations of WASP-33.

Figure 5.7 shows the raw flux for both comparison 1 and comparison 2, along with the differential flux, as functions of both airmass and time (this figure is for the red channel data; equivalent plots for the green and blue channels can be found in the Appendix, Figures A.1 and A.2). The raw fluxes in this figure show a clear trend with airmass, as expected. A simple model for this trend would be

$$\log_{10}(F) = a + bX, \quad (5.1)$$

where X is the airmass. Such a model is physically motivated. Under the assumption that the atmosphere is a grey absorber, the flux received at the Earth's surface will vary exponentially with the path length through the atmosphere, which is proportional the airmass. The quantity b is referred to as the *extinction coefficient*.¹

The plots in the left column of Figure 5.7 (and also Figures A.1 and A.2) give logarithmic fluxes against airmass, so on these plots the model in equation 5.1 would be a straight line. Clearly this would not be a good representation of the raw data; a more complex model is needed. The models plotted in these figures are for an extinction coefficient that varies linearly (red) and quadratically (blue) in time i.e.

$$\log_{10}(F) = a + (b + c\Delta t)X \quad (5.2)$$

$$\log_{10}(F) = a + (b + c\Delta t + d\Delta t^2)X, \quad (5.3)$$

where $\Delta t = \text{Time(MJD)} - 56209.0$. For the raw fluxes, the quadratic model provides a better fit than the linear model. This is particularly apparent in the plots against

¹<http://www.starlink.rl.ac.uk/docs/sc6.htx/node15.html>




Figure 5.7: Red channel raw fluxes for comparison 1 (top row) and comparison 2 (middle row) along with the differential flux for these two stars (bottom row). Each of these is plotted against airmass (left column) and time (right column). The raw fluxes show similar trends in time to the red channel raw flux light curve of WASP-33. The plots of the raw fluxes against airmass show that a linear trend with airmass (equation 5.1) would not be appropriate. Models containing extinction coefficients that vary linearly (red line) and quadratically (blue line) in time do better, with the quadratic model clearly providing a better fit. However, even the quadratic model does not reproduce the variations in the differential flux light curve (bottom right plot). This poor fit is discussed in the text.

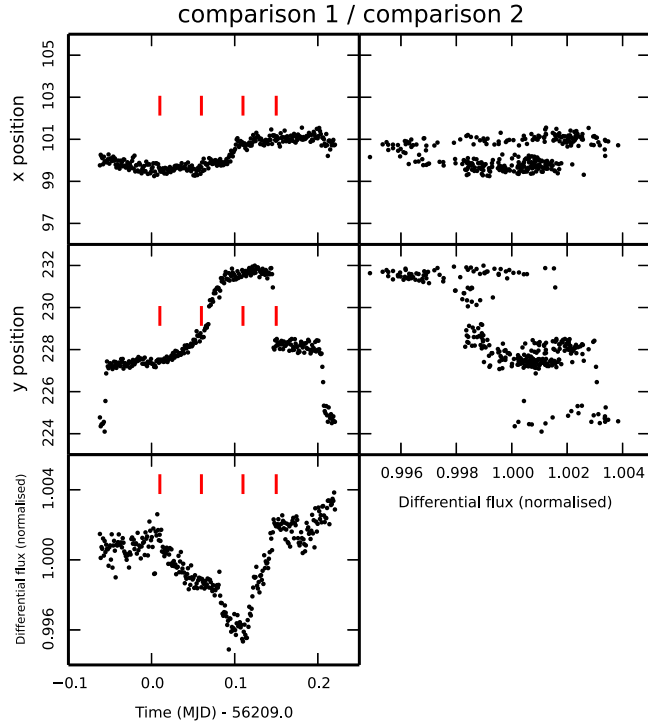


Figure 5.8: Red channel x and y positions for comparison 1 (top and middle rows, respectively) along with the differential flux for comparison 1 and comparison 2. These are plotted against time (left column) and differential flux (right column). A large drift can be seen in the y position, with a smaller drift in the x position. These drifts do not clearly correlate with the variation in the differential flux seen in the bottom panel. For example, between the third and fourth red markers the x and y positions remain stable, while the differential flux increases rapidly.

airmass, especially for the red channel data. This also appears to be true in the differential fluxes (the bottom row plots), however it is clear from the differential flux light curves that even the quadratic model does not provide a good fit. In fact there appears to be a systematic dip in the differential flux of around 0.5% centred around 56209.1 (MJD, UTC) that is common to all three channels (note that the differential flux for comparison 1 / comparison 2 is not expected to have variations due to astrophysical effects, e.g. stellar variability). Given that this feature is larger than the expected eclipse depth ($\sim 0.5\%$ compared to $\sim 0.1\%$), I investigated it further to assess if either of the comparison stars were unsuitable as calibrators for WASP-33.

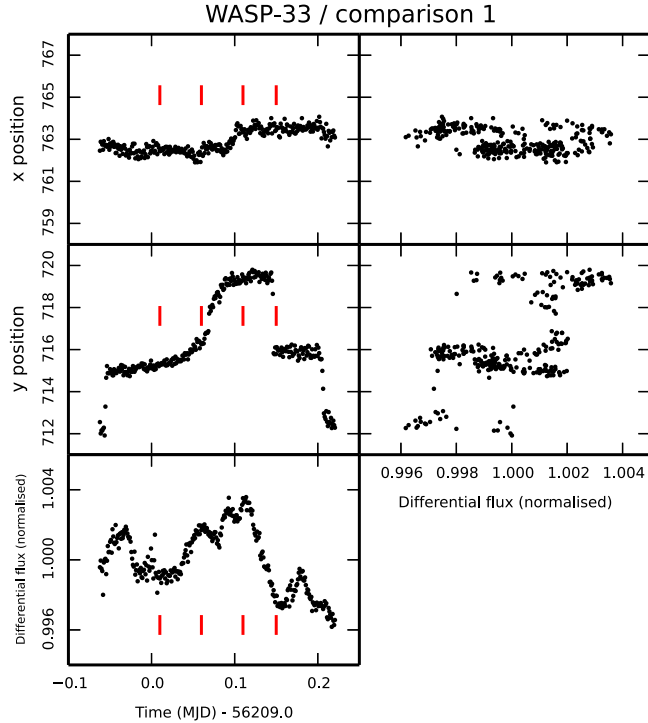


Figure 5.9: An equivalent plot to Figure 5.8, but here plotted for WASP-33 and comparison 1. The differential flux light curve (bottom left plot) shows a similar qualitative behaviour to that in Figure 5.8, but in the opposite sense (and complicated somewhat by the pulsations of WASP-33). Given that comparison 1 acts as the denominator in the differential flux here, while in Figure 5.8 it was the numerator, it seems that comparison 1 is driving the systematic bump seen in these light curves.

One possible source of this systematic is drifting of the stellar PSFs across the CCDs. A drift can clearly be seen in the upper left panels of Figure 5.8. This Figure shows the x and y centroid positions for comparison 1 in time along with the differential flux for comparison 1 / comparison 2. Also shown are the x and y positions against differential flux. (Using the positions for comparison 2 gives a very similar plot, since the positions of the three stars on the CCDs are highly correlated). These quantities are plotted for the red channel data only, where the systematic in the differential flux is most prominent. Figure 5.8 shows that while the flux systematic occurs at roughly the same time as the positional drifts, a clear correlation between these quantities is not present. Using the red markers as a guide,

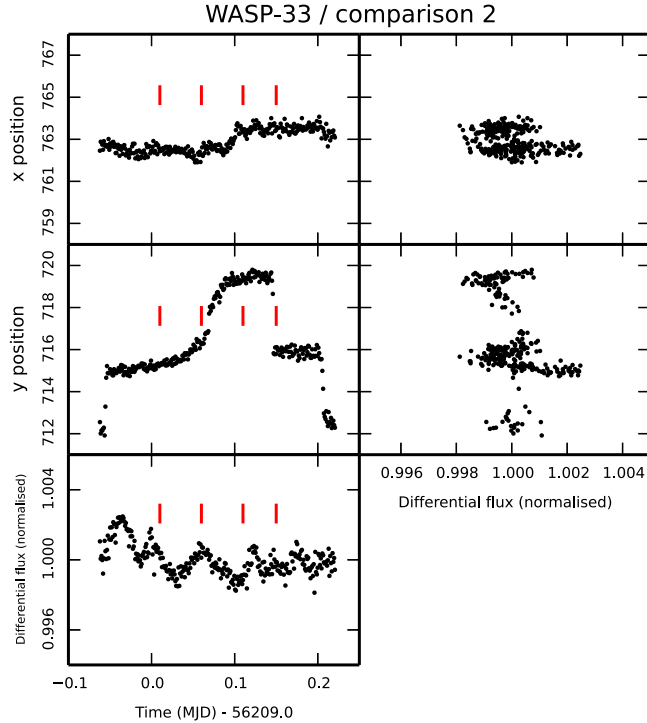


Figure 5.10: An equivalent plot to Figures 5.8 and 5.9, but here plotted for WASP-33 and comparison 2. The systematic bump seen in the differential flux for Figures 5.8 and 5.9 is no longer present, and the 0.2% pulsations of WASP-33 are now clear.

one can see that sometimes a correlation is present e.g. between the first and second red markers the initial drop in flux occurs as the y position begins to drift. However at other times no correlation is seen e.g. between the third and fourth markers there is very little change to the positions but a rapid change in the flux.

I also considered the possibility that the variations in the differential flux are associated with focussing issues around the meridian passage, which occurred at around MJD(UTC) 56209.1. A systematic could be caused by variable amounts of light spilling out of the source apertures as the shape of the PSF changes with the focus. To test this I reduced the data using a very large source aperture ($r_{\text{source}} = 85$ pix) that extended well beyond the wings of PSF, ensuring that any flux variations resulting from a variable PSF would be negligible. The resulting differential flux plot for comparison 1 / comparison 2 showed exactly the same qualitative behaviour as seen in Figure 5.8 (but with increased noise), suggesting the systematic in the differential flux is not the result of light spilling out of the source apertures.

Figures 5.9 and 5.10 show equivalent plots as Figure 5.8, but for the differential fluxes of WASP-33 / comparison 1 and WASP-33 / comparison 2 (also, the positions are given for WASP-33, but these are again tightly correlated with the comparison star positions). Focussing on the bottom-left plot of Figure 5.9, the differential flux appears to show the expected δ Scuti variations. However, while I believe the variations at the beginning and end of the time-series are δ Scuti variations, I believe the central ‘bump’ to be the same systematic seen in Figures 5.7 and 5.8. The bump has a similar behaviour to that seen in Figure 5.8, but in the opposite sense. Again, using the red markers as a guide, there are rises in the flux between the first and third markers, before a quick drop in flux between the third and fourth markers (though these features are complicated somewhat by the pulsation signals from WASP-33). Since comparison 1 acted as the numerator in Figure 5.8, while in Figure 5.9 it acts as the denominator, it seems clear that comparison 1 is driving the systematic. Indeed, the differential flux for WASP-33 / comparison 2 shown in Figure 5.10 shows no signs of the $\sim 0.5\%$ systematic, leaving clear pulsation signals at the 0.1–0.2% level.

Although dropping the use of comparison 1 increased the white noise component in the differential fluxes, it would not have been sensible to use it given the size of the systematic. As a result I used only comparison 2 for the subsequent analysis of the differential flux.

Turning back to the airmass correction, Figure 5.11 is equivalent to Figure 5.7, except this time using WASP-33 and comparison 2 (equivalent plots in the green and blue channels can be found in the Appendix, Figures A.3 and A.4). Though the trends with airmass in these plots are complicated somewhat by the δ Scuti pulsations of WASP-33, these features are at a lower level than the systematic associated with comparison 1. As in Figure 5.7, the model with the extinction coefficient that varies quadratically in time (equation 5.3) fits the raw fluxes better than the model where the extinction coefficient varies linearly in time (equation 5.2). However, the differential light curves are well fitted by the linearly time varying extinction and the quadratic model does not offer much of an improvement. The linear time term is clearly still required however, particularly in the blue channel (see Appendix, Figure A.4), since a straight line fit to the differential flux against airmass would still not be appropriate. Therefore the airmass model with an extinction coefficient that varies linearly in time is used in the subsequent analysis of the data.

Given the correlation between the red channel raw flux and the sky background seen in Section 5.3.3, I tested if there were any correlations between the sky background and the differential flux (after removal of the selected airmass model).




Figure 5.11: Equivalent plot to Figure 5.7, but here shown for WASP-33 and comparison 2. The model with the extinction coefficient that varies quadratically in time (equation 5.3) still fits the raw fluxes better than the model where the extinction coefficient varies linearly in time (equation 5.2). However, there is not a clear improvement in the quadratic fit over the linear fit to the differential flux data. As a result the latter model is used in the subsequent analysis of this data.

I found there to be no such trend.

5.3.5 The light curve model

In order to detect the flux signal from WASP-33b, modelling of the differential flux in each ULTRACAM channel was carried out. For each channel, the model can be thought of in three parts: the airmass trend; the visible flux from the planet (the secondary eclipse) and the pulsations of WASP-33. The airmass correction component is given in equation 5.2, although the model was applied to the flux in linear space, so the actual form of the airmass term was:

$$F_{\text{airmass}} = 10^{(b_1+b_2\Delta t)X} \quad (5.4)$$

where $\Delta t = \text{Time(MJD)} - 56209.0$.

The model for WASP-33b's planetary flux (i.e. the secondary eclipse light curve) was calculated by first modelling the visible fraction of the planet in time, $\eta(t)$. This was done following a similar method to that described for the MCMCTRANSIT program, as shown in Section 2.4 (i.e. using equation 2.21, along with equation 2.17 to convert between time and phase). In contrast to the model in MCMCTRANSIT, I used fixed values for the parameters entering equation 2.21. Specifically these parameters were: $\frac{R_p}{a}$, $\frac{R_p}{R_*}$, i , T_0 and P , with values taken from Collier Cameron et al. [2010]. The model also assumed a circular orbit, since all previous work on this system has shown no signs of a significant eccentricity. As a result, the timing of the secondary eclipse was not free to move away from the $\phi = 0.5$ expectation. I did, however, alter the model to account for the light travel time effect, which makes the eclipse appear later than the $\phi = 0.5$ expectation. I did this by adding 25 s to the times in the model, with the correction being found using the semi-major axis value given by Collier Cameron et al. [2010]. To give the planetary light contribution to the differential flux model, $\eta(t)$ was scaled by the planetary flux, F_p , which was the only free parameter in the secondary eclipse component of the light curve model.

The differential flux variations due to the pulsations of WASP-33 were modelled in general using a sum of sinusoids:

$$F_{\text{pulsations}} = \sum_{i=1}^{N_f} a_i \sin(2\pi\Delta t f_i + \Phi_i), \quad (5.5)$$

where N_f is the number of frequencies (i.e. the number of pulsation modes) used in the model. The details of the pulsation model (e.g. how many different frequencies were used and the restrictions on frequency and phase values) are given in Section

5.3.6.

The complete model for the differential flux light curve was found by adding the planetary light to the pulsation model, along with an offset and then scaling this by the model for the airmass, to give:

$$F_{\text{model}} = \left(a_0 + \sum_{i=1}^{N_f} a_i \sin(2\pi\Delta t f_i + \Phi_i) + \eta(t) F_p \right) \times 10^{(b_1 + b_2 \Delta t)X} \quad (5.6)$$

Separate instances of this model were used for each of the three data sets (from the three channels of ULTRACAM), however as discussed in Section 5.3.6 different forms of the model were tested, with frequencies and phase values fixed and/or linked across the fits to the different data sets.

5.3.6 Light curve fitting

Previous eclipse studies for WASP-33 have used periodogram analyses to pick out the strongest pulsation frequencies in their datasets [Smith et al., 2011; Deming et al., 2012; de Mooij et al., 2013]. However, it has been noted that with only a fraction of a days worth of data, the pulsations are not well defined and can lead to biases in the eclipse depth determinations [de Mooij et al., 2013]. For this study, I have used an alternative method that makes use of the frequencies determined in a dedicated and detailed study of WASP-33's pulsations by von Essen et al. [2014]. In that study, WASP-33's pulsation spectrum was evaluated at optical wavelengths, over a period of 2 years that included the date of the ULTRACAM observations presented in this chapter. As such, the frequencies determined by von Essen et al. [2014] should be appropriate for modelling the pulsations modes seen in the ULTRACAM data.

My method involved fitting sinusoids to the data in each channel (through equation 5.6), using fixed values for the frequencies determined by von Essen et al. [2014] (using their best fit values). The amplitudes (a_i) for each frequency mode were free to vary and to be different in each channel, because the pulsation amplitudes of δ Scuti stars can have significant colour dependencies. The phase value (Φ_i) for each frequency mode was also free, however it was restricted to be the same value across the three channels (tests for relaxing this restriction are discussed later in this section). The other free parameters (a_0 , F_p , b_1 , b_2) were allowed to be different for the three channels. The number of free parameters, for a given number of modelled frequency modes (N_f), was therefore $12 + 4N_f$.

Fits were optimised through χ^2 minimisation using a Levenberg-Marquardt algorithm [Press et al., 1992]. The starting parameter values for this optimisation

were all taken to be 0, except for a_0 and b_1 . The starting values for these parameters were determined with a fit to the differential flux in each channel using equation 5.6, with a_i , F_p and b_2 all fixed to 0. Note that the χ^2 value that was minimised in the Levenberg-Marquardt algorithm was from the joint fit to all the channels - fitting was not done in each channel separately (due to Φ_i being linked across the channels). At this stage errors on the parameters were not assessed - this was done using an MCMC algorithm along with a prayer bead analysis (see Section 5.3.7).

In order to select an appropriate model for the pulsations of WASP-33, I fitted a set of increasingly complex models to the data. Starting with 1 mode, the frequency (from von Essen et al. 2014) that provided the best fit to the data (according to χ^2) was selected. Then another mode was added by testing the remaining von Essen et al. [2014] frequencies, with the frequency that provided the biggest χ^2 improvement again being chosen. This was done until the full set of 8 frequencies from von Essen et al. [2014] were used. Table 5.1 summarises the best set of frequencies for each stage of this process. The aim here was to track how the fit improved and how the eclipse depths varied in each channel. Figure 5.12 shows the results from this analysis, with plots of: the normalised residual RMS (binned to the ingress/egress timescale of the eclipse); the reduced χ^2 and the eclipse depths, as a function of the number of frequency modes used in the model (N_f). Each of these is given for the three ULTRACAM channels.

The eclipse depth (i.e. planet-to-star flux ratio, ΔF) in each channel was found as:

$$\Delta F = \frac{F_p}{a_0}, \quad (5.7)$$

where a_0 was used as the estimate of the average stellar flux in the channel. This was a reasonable assumption because the a_i values tended to be around 1000 times smaller than a_0 (since the δ Scuti pulsations are at the mmag level) and also the average of the sinusoid terms will tend to 0 over the time-series.

Figure 5.12 shows that there are improvements to the RMS in all channels up to $N_f = 7$, but that adding the eighth frequency makes little difference to the quality of the fit. The reduced χ^2 shows a similar trend, though in the green channel χ_r^2 does not show much improvement beyond $N_f = 5$. The fits are formally good ($\chi_r^2 \simeq 1$) in the green and blue channels, but this is not the case in the red channel. The red channel data has a higher signal-to-noise and so could be affected by things such as: higher order variability in the Earth's atmosphere (i.e. the airmass model is not quite appropriate); instrumental effects or higher frequency pulsations for WASP-33.

It can be seen from the bottom-left panel of Figure 5.12 that the red channel




Figure 5.12: Plots of the normalised residual RMS (with the residuals binned to the ingress/egress timescale of the eclipse; top row), the reduced χ^2 for the light curve model fit (middle row) and the eclipse depths (bottom row), as a function of the number of frequency modes used in the model (N_f). These plots are given for the red (left column), green (middle column) and blue (right column) ULTRACAM channels. The errors on the eclipse depths for the models using $N_f = 5-7$ come from the error analyses described in Section 5.3.7. Errors were not determined for other values of N_f . The most significant trend in these plots is the consistency of the eclipse detection in the red channel (bottom left plot). The eclipse depths vary between 0.039% and 0.063% across the entire range of N_f values, and all agree within the errors given for the $N_f = 5-7$ model fits.

Table 5.1: The frequencies used in the modelling of WASP-33’s pulsations, as described in Sections 5.3.5 and 5.3.6. The frequency values here were taken from an analysis of WASP-33’s pulsations by von Essen et al. [2014]. N_f represents the number of pulsation modes modelled (see equation 5.6). For each N_f , the frequency mode that was added to the previous ($N_f - 1$) model was that which gave the largest improvement to χ^2 .

N_f	Frequencies used (d^{-1})
1	9.8436
2	9.8436, 24.8835
3	9.8436, 24.8835, 20.1621
4	9.8436, 24.8835, 20.1621, 21.0606
5	9.8436, 24.8835, 20.1621, 21.0606, 34.1252
6	9.8436, 24.8835, 20.1621, 21.0606, 34.1252, 10.8249
7	9.8436, 24.8835, 20.1621, 21.0606, 34.1252, 10.8249, 8.3084
8	9.8436, 24.8835, 20.1621, 21.0606, 34.1252, 10.8249, 8.3084, 20.5353

eclipse depths have a consistent value. This is remarkable given the changes in the pulsations models used. The other channels show larger variations. In the green channel, the eclipse depths are around 0 up to $N_f = 5$ and then jump to being positive for $N_f = 7$ and 8, while for the blue channel the depths are consistently negative, before jumping to being around 0. In each channel the $N_f = 6$ fit is consistently low.

Since the $N_f = 8$ fit does not offer an improvement over the $N_f = 7$ fit, and since there is no significant change to the eclipse depths in any channel, the $N_f = 7$ fit is chosen as the best fit. However, there is a discontinuity in the eclipse depths in each channel between $N_f = 5 - 7$ and this is explored further in the error analysis in Section 5.3.7.

I also tested some variations on the modelling technique used. For example, rather than fixing the phases to be the same across the three channels, I allowed them to be free in the fit. This was motivated by the fact that, at different wavelengths, pulsation modes of δ Scuti stars can be offset in phase [e.g. Garrido & Rodriguez, 1990]. These fits produced results very similar results to those described above, for fits up to $N_f = 6$. The results agreed to a level well below the uncertainties for the $N_f = 5$ fit, as determined in Section 5.3.7. For more complex models, differences in phase (for a given frequency) across the three channels were becoming as large as $90-180^\circ$, which is not physical [typical phase offset between photometric band are $< 10^\circ$ Garrido et al., 1990]. Given the consistency of the eclipse depths up to $N_f = 6$, I favoured the more simple model with fixed phases across the three channels.

I also tried modelling using the pulsation frequencies found in the data itself,

as has been done in similar previous studies [Smith et al., 2011; de Mooij et al., 2013]. I followed the same pattern as before of increasing the complexity of the model and tracking changes to the fit quality and eclipse depths. The first frequency was assessed by determining the Lomb-Scargle periodogram of the residuals to a fit that used no sinusoidal terms. The dominant frequency in the periodogram was then used as a starting point for a Levenberg-Marquardt minimisation for $N_f = 1$, where the frequency (f_i) now entered the model as a free parameter (though, like Φ , the value was restricted to be the same across the three channels). The residuals from this fit were again assessed using a Lomb-Scargle periodogram, with the next dominant frequency being selected and used in the $N_f = 2$ fit. This process was repeated up to $N_f = 5$. The resulting eclipse depths showed good agreement with those derived from the original technique in the red and blue channels. The only disagreement for the green channel was for the model using $N_f = 2$, where there was a difference in the eclipse depths of $< 2\sigma$. Therefore, I saw no motivation to alter the technique I used to fit WASP-33’s pulsations.

5.3.7 Eclipse depth uncertainty estimates

Uncertainties for the eclipse depths in each channel, for the models using $N_f = 5 - 7$, were assessed in two ways: using a Markov Chain Monte Carlo (MCMC) routine and using a ‘prayer-bead’ method.

The MCMC routine I used here was different to that described and used in Chapters 2–4. I used the Python program EMCEE [Foreman-Mackey et al., 2013], which utilises a slightly different algorithm to the Metropolis-Hastings algorithm described in Section 2.4.2. Briefly, the algorithm simultaneously evolves an ensemble of points in the model parameter space (known as ‘walkers’). At each step in the chain the proposal step for a given walker is found by drawing randomly from the current state of all the other walkers in the ensemble and then perturbing from this, based on the parameter differences of the two walkers. Comparison of the posterior probabilities of the current and proposed model fits (given by equation 2.26) are then used to determine if the proposal is accepted, in a similar way to the Metropolis-Hastings algorithm. If the proposal is rejected, the state of the walker remains the same for the next step in the chain (note the slight difference here compared to the MCMCTRANSIT algorithm). All of the walkers in the ensemble are updated in this way before moving onto the next step.

The light curve model used to fit the WASP-33 data in the MCMC routine was as described in Sections 5.3.5 and 5.3.6, with separate instances of equation 5.6 being fitted to each channel, phase values being fixed across the three channels




Figure 5.13: Binned normalised residual RMS values as a function of bin width (dt). From top to bottom, plots are for the red, green and blue channels of ULTRACAM. The different line colours represent fits using different numbers of frequencies, N_f . Black is for $N_f = 5$, red is for $N_f = 6$ and blue is for $N_f = 7$. The dotted lines show the $1/\sqrt{n}$ expectation for white noise, where n is the number of data points per bin. This expectation is fixed to the unbinned RMS value for the fit using $N_f = 5$. In each channel the binned RMS values diverge from the white noise expectation, motivating the use of the prayer bead error analysis in Section 5.3.7.

and frequencies fixed to the values given in Table 5.1. The MCMC jump parameters were therefore the same as the free parameters used in the Levenberg-Marquardt fit in Section 5.3.6. In order to obtain appropriate eclipse depth uncertainties, the differential flux errors were rescaled such that the best fit (found in Section 5.3.6) gave a reduced χ^2 of 1 in each channel. I also note here that uniform priors were placed on all of the jump parameters.

I used an ensemble of 500 walkers for which the parameter values were initialised using the Levenberg-Marquardt minimisation solution (from Section 5.3.6). Each walker was perturbed from this, in each dimension of the model parameter space, by a Gaussian random number with standard deviation of 5σ (where σ values were found for each jump parameter from an initial MCMC run). A burn-in

Table 5.2: Eclipse depth and uncertainty estimates derived from light curve models using $N_f = 5-7$ (see equation 5.6 and Table 5.1). The eclipse depths are taken from the Levenberg-Marquardt fit (Section 5.3.6), while uncertainty estimates derived from the MCMC and prayer bead analyses (Section 5.3.7) are given in the fourth and fifth columns, respectively. The adopted uncertainty value was taken as the largest value from these two analyses.

Channel	N_f	Eclipse depth (%)	MCMC (%)	Prayer (%)	Adopted values (%)
Red	5	0.052	± 0.006	$+0.016$ -0.017	$0.052^{+0.016}$ -0.017
Red	6	0.039	± 0.007	$+0.020$ -0.018	$0.039^{+0.020}$ -0.018
Red	7	0.057	± 0.007	$+0.008$ -0.017	$0.057^{+0.008}$ -0.017
Green	5	-0.012	± 0.008	$+0.017$ -0.015	$-0.012^{+0.017}$ -0.015
Green	6	-0.033	± 0.010	$+0.016$ -0.039	$-0.033^{+0.016}$ -0.039
Green	7	0.037	± 0.017	$+0.066$ -0.003	$0.037^{+0.066}$ -0.017
Blue	5	-0.065	± 0.021	$+0.071$ -0.058	$-0.065^{+0.071}$ -0.058
Blue	6	-0.107	± 0.023	$+0.072$ -0.060	$-0.107^{+0.072}$ -0.060
Blue	7	-0.001	± 0.028	$+0.024$ -0.020	-0.001 ± 0.028

phase of 2×10^3 steps was carried out before a production run of 2×10^4 steps. From the resulting posterior distribution, the uncertainties on the eclipse depths in each channel were found from the 68% confidence intervals of the marginalised distributions for these parameters (i.e. their histograms). Separate MCMC runs were carried out for models using $N_f = 5, 6$ and 7 . The resulting uncertainty values can be found in Table 5.2.

Along with the assessment of uncertainties from the MCMC routine, I was motivated to perform a ‘prayer bead’ analysis [Gillon et al., 2007] after assessing the trend in the binned residual RMS values in each channel, as shown in Figure 5.13. The RMS values tend to diverge from the white noise expectation in each of the channels - a sign that the data is affected by red noise. Note that the red noise signal is suppressed for fits using more frequency modes, suggesting that residual pulsations are driving the red noise signal.

The prayer bead analysis I performed here was similar to that described for WASP-3 in Section 3.3.10, where the residuals to the light curve fits were subtracted from the raw differential flux data and then added back in with a cyclic offset in phase. I used 5000 residual shifts for this analysis and for each shift the model described in Section 5.3.6 (and equation 5.6) was fitted to the data using a Levenberg-Marquardt χ^2 minimisation algorithm [Press et al., 1992]. From the

Table 5.3: Final eclipse depth results for the secondary eclipses in each of the ULTRACAM channels.

Channel	Filter	Eclipse depth
Red	z'	$0.057^{+0.008}_{-0.017} \%$
Green	Bcont	$< 0.155 \%$ (3σ)
Blue	u'	$< 0.069 \%$ (3σ)

resulting distribution of eclipse depths, the 1σ eclipse depth uncertainty was found as the 68% confidence interval centred on the value from the original fit (that had no residual shift applied). This analysis was applied to each of the ULTRACAM channels separately, e.g. to assess the prayer bead errors in the red channel eclipse depth, the blue and green channel data were not modified. As with the MCMC analysis, prayer bead errors were derived for models using $N_f = 5, 6$ and 7 . The resulting eclipse depth errors from this prayer bead analysis can be found in Table 5.2.

The final column of Table 5.2 gives the adopted error values in each channel, for each N_f . The adopted error was chosen as the largest error resulting from the MCMC and prayer bead analyses. The chosen errors are plotted on the eclipse depth values for the $N_f = 5 - 7$ fits in Figure 5.12.

The main conclusion from these error analyses is that the eclipse in the red channel (the z' band) is detected significantly and consistently across the model fits using $N_f = 5 - 7$. The z' band eclipse depths for other N_f values are also consistent within the errors derived here. Because the $N_f = 7$ solution gives the best fit to the data in the red channel, I use the eclipse depth and uncertainty values from this fit as the final value.

For the green and blue channels, the error analysis shows that eclipse depths are not reliably found across the $N_f = 5 - 7$ fits, and are not clearly distinguished from 0. As a result, I only place upper limits on the eclipse depths for the green and blue channels.

5.4 Results

The final results from the analysis described above are given in Table 5.3. The main result I report in this chapter is the detection of a secondary eclipse signal from WASP-33b in the z' band. The eclipse depth is $\Delta F_z = 0.057^{+0.008}_{-0.017} \%$ and this corresponds to a brightness temperature of $T_z = 3170^{+90}_{-190}$ K (derived using the

method described in Section 2.6.1).

The results for the other two ULTRACAM bands do not show clear detections of an eclipse signal and so I have placed 3σ upper limits on the eclipse depth values here. The upper limits were taken as the 3σ limit of the distribution resulting from either the MCMC or the prayer bead analysis, whichever was largest. The model fits using $N_f = 7$ were used here. For the Blue continuum band (green channel) the limit is $\Delta F_{\text{Bcont}} < 0.155\%$ (resulting from the prayer bead analysis), while for the u' band (blue channel) it is $\Delta F_u < 0.069\%$ (resulting from the MCMC analysis).

The top row of Figure 5.14 shows the normalised differential fluxes in each channel along with the $N_f = 7$ model, which was used to derive the eclipse depth in the z' band. Note that the models shown for the Blue continuum and u' bands do not reflect the final results in these bands - upper limits were given on the eclipse depths here instead. The middle row of Figure 5.14 shows the normalised differential fluxes that have had the pulsation and airmass model components removed, leaving the planetary flux signal. In the z' band the eclipse is clearly seen. A non-zero eclipse depth is also seen in the Blue continuum band, but this is not adopted due to the variations seen across the fits using $N_f = 5 - 7$ (but see Section 5.5.3). Finally, in the u' band no eclipse is seen.

5.5 Discussion

5.5.1 Comparison to previous measurements

Detection of exoplanetary eclipses is difficult and for WASP-33b this is especially the case due to its pulsations. Previous eclipse detections for WASP-33b have typically assessed the dominant pulsations frequencies using the eclipse datasets i.e. the data is detrended against itself. Here I have used a different technique, basing the pulsation model for WASP-33 on *a priori* determined frequencies found in a detailed study by von Essen et al. [2014]. These were applied to the simultaneous multi-band data provided by the ULTRACAM instrument. In the Blue continuum and u' bands I found formally good fits to the data, with the fit using 7 of the frequencies found in the von Essen et al. [2014] study being the best. This suggests the pulsations are well modelled using this method. The reduced χ^2 was higher for the z' band, a result of the higher signal-to-noise in this band. However, the eclipse depth solutions in this band were remarkably stable, with extreme values of 0.039% and 0.063% and agreement within errors across the fits using different numbers of pulsations modes (N_f). Therefore I believe that the method used here and the z' band eclipse depth derived are robust.




Figure 5.14: Plots of: the (normalised) differential flux, along with the light curve model using $N_f = 7$ (top row); the differential flux that has had the pulsation and airmass model components removed and has been normalised to the stellar flux, leaving the secondary eclipse signal (middle row); the residuals from the light curve model shown in the top row (bottom row). These are given for the red (z'), green (Blue continuum) and blue (u') channels, from left to right. The eclipse signal for the red channel is clearly seen and has a depth of $0.057^{+0.008}_{-0.017}$ %. For the final results in the green and blue channels I set upper limits on the eclipse depths, however the best fitting model for the green channel, shown here, does hint at a detection of an eclipse, with a depth of $0.037^{+0.066}_{-0.017}$ %. This is discussed in Section 5.5.3.

A similar secondary eclipse measurement has been derived previously by Smith et al. [2011]. This measurement was taken in the narrow S[III] filter, using the ACAM instrument on the WHT. This filter has a central wavelength of $\lambda_c = 0.9077 \mu\text{m}$, which is very similar to the central wavelength of the z' band observation reported here, $\lambda_c = 0.9058 \mu\text{m}$. Smith et al. [2011] found the S[III] eclipse depth to be $0.109 \pm 0.030 \%$, in marginal agreement with my value of $0.057_{-0.017}^{+0.008} \%$. However, the eclipse measured by Smith et al. [2011] is heavily affected by red noise (see Figure 5.1), possibly a result of the thin cirrus cloud that was present during the observation. The pulsations were modelled by testing combinations of frequencies derived from their own data and the dominant frequency found by Herrero et al. [2011] (21 d^{-1}). For their final result, Smith et al. [2011] use three frequencies from their own data, along with the 21 d^{-1} frequency, but they do not report eclipse depth values for other models they assessed (using different numbers of frequency modes). I feel, therefore, that my eclipse detection is more robust than that of Smith et al. [2011], so I will not use the latter in the subsequent discussion.

One point to note here is that, although the central wavelength of the z' and S[III] filters are similar, their widths are very different (FWHM = 1370 \AA and FWHM = 54 \AA , respectively). As a result one would not expect the eclipse depths to be exactly the same. A better comparison is provided by the brightness temperatures because these account for the filter responses (see Section 2.6.1). The brightness temperature estimate for my measurement is $T_z = 3170_{-190}^{+90} \text{ K}$, while for the Smith et al. [2011] measurement it is $3625_{-240}^{+210} \text{ K}$. These values are in marginal agreement, as were the eclipse depths, suggesting the differences between the filter responses do not have a significant effect (at least for a blackbody model for the planet).

5.5.2 Temperature estimates

Figure 5.15 shows the secondary eclipse depths (planet-to-star flux ratios) for WASP-33 as a function of wavelength. These data are taken from this chapter, Smith et al. [2011], Deming et al. [2012] (K_s band and *Spitzer* 3.6 and $4.5 \mu\text{m}$) and de Mooij et al. [2013] (K_s band). The K_s band depth used here is a weighted average of the measurements by Deming et al. [2012] and de Mooij et al. [2013], which were consistent. Also note that while the Smith et al. [2011] eclipse depth is plotted it is not used in the following analyses.

Also plotted are flux ratio models using a blackbody spectrum for the planet and a stellar flux model from Kurucz [1993]. The middle model is for the best fitting planetary blackbody, found by minimising χ^2 . The model comparisons to the eclipse depths were found as the passband integrated blackbody model divided by

the passband integrated stellar flux (shown as grey squares for the best fit). The best fitting blackbody temperature is 3200 K, confirming WASP-33b as the hottest known exoplanet. I have also plotted models using equilibrium temperatures (as described in Section 2.6.2) assuming $A_B = 0$, with $\varepsilon = 0$ (upper curve in Figure 5.15) and $\varepsilon = 1$ (lower curve). These correspond to temperatures of 3440 K and 2693 K, respectively. Since the best fitting blackbody is at the upper end of this temperature range it is clear that the infrared eclipse depths are suggestive of a planet that is absorbing a large fraction of the incident flux and is not redistributing this heat efficiently.

To test this further I carried out an analysis into the Bond albedo (A_B) and the global heat redistribution properties (ε) of the planet, following the methodology of Cowan & Agol [2011], as described in Section 2.6.2. To estimate the day-side effective temperature of the planet I first determined brightness temperature estimates for the eclipses observed by Deming et al. [2012] and de Mooij et al. [2013] (see the lower panel of Figure 5.15). The resulting values, along with my estimate for the z' band brightness temperature, were used in equation 2.36 to give a day-side effective temperature estimate of $T_d = 3170^{+80}_{-130}$ K. This value matches the z' band brightness temperature I determined in Section 5.4, in part because the z' band samples the expected spectral peak of WASP-33b ($\lambda_{\max} \sim 880$ nm), which strongly affects the planet's bolometric flux. This highlights the importance of this detection in setting the energy output for the planet. I note that a previous analysis of the effective temperature, including the Smith et al. [2011] measurement, gave a value of $T_d = 3300 \pm 70$ K [de Mooij et al., 2013]. My z' band detection has revised this down marginally.

Using the system parameters from Collier Cameron et al. [2010] the substellar equilibrium temperature for the planet is $T_0 = 3810 \pm 80$ K and so $\frac{T_d}{T_0} = 0.83^{+0.05}_{-0.06}$. With $T_{\varepsilon=0} = 3440$ K, the high $\frac{T_d}{T_0}$ value for WASP-33 supports the emerging trend pointed out by Cowan & Agol [2011] that highly irradiated planets have systematically high $\frac{T_d}{T_0}$ values, implying low values for A_B and ε .

The latter point can be seen more clearly in the 2D probability distribution for WASP-33, found using the method described in Section 2.6.2. While the degeneracy between the albedo and heat redistribution is apparent here, each parameter is restricted to low to moderate values, even when the other parameter is minimised. High values for both A_B and ε are strongly ruled out here. The marginal 1σ upper limits reflect these conclusions, with $\varepsilon < 0.408$ and $A_B < 0.255$.

I also note here that the previous eclipse depths for WASP-33b have been compared to the atmospheric models of Madhusudhan [2012]. In that paper it was found that the data could be explained by both an inverted atmosphere with a




Figure 5.15: The spectrum of the planet WASP-33b. The top panel shows planet-to-star flux ratio values for WASP-33b as a function of wavelength. Data is taken from this chapter (z' band measurement, highlighted in red), Smith et al. [2011], Deming et al. [2012] and de Mooij et al. [2013]. The grey curves are flux ratio models assuming a blackbody spectrum for the planet and a stellar flux model from Kurucz [1993]. The middle curve is the best fit to the data, excluding the Smith et al. [2011] S[III] band value (white point). The filter response curves are shown for each of the measurements. The bottom panel is equivalent to the top panel, but plotted in terms of the planetary brightness temperature.




Figure 5.16: The 2D PDF for WASP-33b found using its T_d/T_0 value with equation 2.39, assessed over a grid of allowable albedo and ε values. Darker regions are of higher probability than lighter regions. The high value for T_d/T_0 restricts the allowable values of A_B and ε such that both cannot be large. The 1σ upper limits on the marginal distributions are 0.408 for ε and 0.255 for A_B , reflecting the restriction to low values seen in the 2D PDF.

solar composition and a non-inverted, carbon-rich atmosphere (with the latter being slightly favoured by the Deming et al. [2012] K_s band eclipse depth). Unfortunately the distinction between these models in the z' band is not great and my value is consistent with both model predictions. As such my z' band eclipse depth does not support either model over the other.

5.5.3 Thermal emission or reflected light?

Up to this point I have been assuming that the thermal emission of WASP-33b dominates the reflected light in the z' band. However, this may not be the case. The reflected light contribution to the planet-to-star flux ratio is given by $A_{g,\lambda} \left(\frac{R_p}{a}\right)^2$. Here, $A_{g,\lambda}$ is the wavelength-dependent geometric albedo - the ratio of the planet's reflected day-side flux to that of an ideal, spherical reflector of the same size and orbital separation as the planet (see Section 1.3.1.1). As an indicative example of

the possible effect of a significant reflected light signal I adopt a value of $A_g = 0.33$. While this is a high albedo in comparison to hot Jupiters in general, it was the value measured for the planet Kepler-13Ab [Shporer et al., 2014], which is of similar radius to WASP-33b and also orbits closely to an A-type star. Assuming no wavelength dependence for A_g , the expected reflected light eclipse depth is 0.026 % (using the system parameter values from Collier Cameron et al. 2010). This forms a significant fraction of the 0.057 % eclipse depth derived for the z' band. Assuming this contribution, the brightness temperature for WASP-33b in the z' band would be 2830^{+120}_{-360} K.

I re-performed the analysis of WASP-33b's day-side effective temperature, given in Section 5.5.2. This time I used the brightness temperatures derived from each of the eclipse depths with a 0.026 % reflected light component removed. The resulting value for $\frac{T_d}{T_0}$ was $0.76^{+0.05}_{-0.06}$. In terms of the trend for highly irradiated planets to have high $\frac{T_d}{T_0}$, this value is still reasonably consistent, but it is lower than has been typically found for the highly irradiated planets (see Figure 1.14).

Given there is no direct albedo measurement for WASP-33b, the true contributions of reflected and thermally emitted planetary light to the z' band cannot yet be determined. While the analysis above was made for illustrative purposes, it does highlight the need to consider reflected light at this wavelength, and as such provides an important caveat to the results presented in Section 5.5.2.

Finally, it is interesting to consider the results from the other two bands I have analysed. In the u' band, the upper limit of 0.069 % on the eclipse depth only provides a weak constraint on WASP-33b's geometric albedo, at $A_{g,u} < 0.88$. This is much larger than has been found for any hot Jupiter planet so far. The upper limit for the Blue continuum band, at 0.155 %, does not allow for any constraint to be placed on $A_{g,Bcont}$. However, the best fit in this band (using the $N_f = 7$ model) gives an eclipse depth of $0.037^{+0.066}_{-0.017}$ %. The value of A_g required to give such an eclipse depth is 0.47. Although at the top end of the albedo values being found for hot Jupiters, this is still plausible [e.g. Evans et al., 2013]. Additionally, the thermal contribution of WASP-33b to the Blue continuum band is 0.007 % (found using the best fitting blackbody temperature of 3200 K). Accounting for this, the geometric albedo required to match the Blue continuum eclipse depth drops to 0.38 - comparable to that found for Kepler-13Ab. Of course, the variations in eclipse depth with N_f for the Blue continuum channel means that the eclipse depth of $0.037^{+0.066}_{-0.017}$ % is not robust and I do not claim a detection here. However, a repeat observation of the secondary eclipse of WASP-33 in the Blue continuum (or similar) band would be able to shed more light on this tentative result.

5.6 Conclusions

In this chapter I have analysed secondary eclipse observations of the star-planet system WASP-33 using the simultaneous three-channel photometer, ULTRACAM, on the WHT. The δ Scuti pulsations of the host star were detected in each of the ULTRACAM channels. These pulsations were modelled using a sum of sinusoids, with frequencies based on those found by von Essen et al. [2014], in a joint fit that also included the planetary flux (secondary eclipse) signal. The pulsations were removed effectively using this method and a robust eclipse detection in the z' band was made. This signal was found consistently for models using different numbers of pulsation modes, and the best fit eclipse depth was found to be $\Delta F_z = 0.057^{+0.008}_{-0.017}$ %.

Under the assumption that this flux is a result of the thermal emission of the planet, the z' band brightness temperature was found to be $T_z = 3170^{+90}_{-190}$ K. The high temperatures found for this planet suggest it absorbs much of the incident stellar flux and does not redistribute this energy efficiently to its night-side. However, it is plausible that a significant fraction of ΔF_z results from reflected light, rather than thermal emission, which would lower the T_z estimate.

Although only upper limits could be placed on the eclipse depths for the Blue continuum and u' band, the best fit eclipse in Blue continuum band did give a plausible value of $\Delta F_{\text{Bcont}} = 0.037^{+0.066}_{-0.017}$ %. Further observations will be needed to confirm or reject this tentative result.

Chapter 6

Conclusions and future work

In this thesis I have presented secondary eclipse detections of five WASP planets, at optical and infra-red wavelengths using ground- and space-based observations. Summaries of these analyses are given below, before a discussion of potential follow-up studies resulting from my work and a more general outlook for the future of the atmospheric characterisation of exoplanets.

6.1 Chapter summaries

6.1.1 The thermal emission of the exoplanet WASP-3b

In Chapter 3, I analysed the infra-red secondary eclipses of WASP-3b, using *Spitzer's* IRAC instrument, at 3.6 μm , 4.5 μm and 8.0 μm . The data in each of the IRAC channels were affected by well known detector related systematics, which were corrected for using simple functions of the centroid position of the target's PSF, and time. The eclipse light curves were modelled simultaneously using the Markov Chain Monte Carlo code of Collier Cameron et al. [2007a] and the resulting eclipse depths for WASP-3 were: $\Delta F_{3.6\mu\text{m}} = 0.209^{+0.040}_{-0.028} \%$, $\Delta F_{4.5\mu\text{m}} = 0.282 \pm 0.012 \%$ and $\Delta F_{8.0\mu\text{m}} = 0.328^{+0.086}_{-0.055} \%$. The corresponding brightness temperatures for the planet were: $T_{3.6\mu\text{m}} = 2280^{+210}_{-150} \text{ K}$, $T_{4.5\mu\text{m}} = 2400 \pm 80 \text{ K}$ and $T_{8.0\mu\text{m}} = 2210^{+390}_{-250} \text{ K}$.

The high temperatures derived for the planet suggest that it absorbs much of the incident radiation from its host, whilst not efficiently redistributing this energy to its night-side. Weak energy redistribution was also implied in the comparisons of the eclipse depths to the models of Fortney et al. [2008]. In addition, these models favoured the the presence of a temperature inversion in the atmosphere of WASP-3b

WASP-3 is a moderately active star ($\log R'_{\text{HK}} = -4.87$) that probes the cut-off (at $\log R'_{\text{HK}} \sim -4.9$) in the proposed activity-inversion relation suggested

by Knutson et al. [2010]. My detection of an inverted atmosphere for WASP-3b is in contrast to the weakly/non-inverted atmospheres of two other planets orbiting similarly active stars (WASP-4b and HAT-P-6b), suggesting the cut-off with activity may not be as well defined as initial data suggested.

6.1.2 Testing the effects of metallicity on the temperature structure of exoplanet atmospheres

In Chapter 4, I presented *Spitzer* IRAC secondary eclipse observations of three low metallicity systems at 3.6 μm and 4.5 μm , with the aim of testing the effects of metallicity on the spectral properties of these planets. The systems observed were: WASP-21, WASP-28 and WASP-37.

As for WASP-3, the flux measurements from these observations showed signs of IRAC’s intra-pixel sensitivity systematic, which were modelled by testing a variety of polynomial detrending functions based on the target’s centroid position. Similarly to the analysis for WASP-3, the MCMC code of Collier Cameron et al. [2007a] was used to extract planet-to-star flux ratios for the three planets at 3.6 μm and 4.5 μm . These values were: $\Delta F_{4.5\mu\text{m}} = 0.064 \pm 0.014 \%$ (for WASP-21); $\Delta F_{3.6\mu\text{m}} = 0.077 \pm 0.017 \%$ and $\Delta F_{4.5\mu\text{m}} = 0.112 \pm 0.022 \%$ (for WASP-28), and $\Delta F_{3.6\mu\text{m}} = 0.097^{+0.023}_{-0.025} \%$ and $\Delta F_{4.5\mu\text{m}} = 0.090 \pm 0.030 \%$ (for WASP-37). The 3.6 μm eclipse depth for WASP-21 could not be determined reliably and so an upper limit of 0.11 % was given as the result here. The temperature estimates derived from these flux ratios allow for a range of planetary reflectance and heat redistribution properties, in keeping with other planets that experience a similar amount of irradiation from their hosts.

Exoplanet systems with existing *Spitzer* IRAC measurements showed a weak correlation between the host star metallicity and the empirical 3.6 – 4.5 μm spectral slope of the planet (ζ). The low metallicity systems studied in this chapter had the potential to consolidate this weak trend, however the measured ζ values for WASP-21, WASP-28 and WASP-37 did not provide any support for it.

6.1.3 Ground-based secondary eclipse observations of WASP-33

In Chapter 5, I analysed ground-based secondary eclipse observations of WASP-33, taken using the ULTRACAM instrument on the William Herschel Telescope (WHT). This instrument provided simultaneous observations in three wavelength bands from the optical to the near infra-red (centred on 3557 \AA , 5149 \AA and 9097 \AA). The $\sim 0.2 \%$ δ Scuti pulsations of the host star were detected and modelled using a sum of sinusoids, with frequencies that had been measured *a priori* in a dedicated

study of WASP-33 by von Essen et al. [2014].

The removal of these pulsation signals revealed a robust z' band eclipse detection. Consistent z' band eclipse depths were found across a range of models, using different amounts of sinusoids for the pulsation model component. The best fit z' band eclipse depth was found to be $\Delta F_z = 0.057^{+0.008}_{-0.017}$ %, corresponding to a brightness temperature of $T_z = 3170^{+90}_{-190}$ K. This detection strongly constrains the bolometric flux of the planet and the resulting day-side effective temperature for WASP-33b ($T_d = 3170^{+80}_{-130}$ K) suggests the planet absorbs much of the incident stellar flux and does not redistribute this energy efficiently to its night-side. However, it was found that there could be a significant reflected light component to the detected eclipse depth, which would result in lower temperature estimates for the planet.

Upper limits were placed on the optical secondary eclipses, since consistent depths could not be found reliably for the different pulsation models tested. However the best fit eclipse in the Blue continuum band ($\lambda_c = 5149 \text{ \AA}$) did reveal a tentative detection of $\Delta F_{\text{Bcont}} = 0.037^{+0.066}_{-0.017}$ %.

6.2 Future work

There are some interesting follow-up observations that could be made to characterise some of the tentative results presented in this thesis. In Chapter 3, I highlighted that simultaneous monitoring of the IRAC eclipse depths and $\log R'_{\text{HK}}$ values for three systems (WASP-3, WASP-4 and HAT-P-6) could help to characterise the activity cut-off in the proposed activity–inversion correlation of Knutson et al. [2010]. From Chapter 5, the tentative detection of an eclipse in the Blue continuum filter will need a repeat observation in order to be confirmed or rejected. These observations do, however, highlight the exciting possibilities of ground-based reflected light studies, which could place important constraints on the global energy properties of large numbers of hot Jupiter exoplanets.

Although my work into the effects of metallicity on the spectral properties of exoplanets returned a null result, it did highlight the potential of searching for trends in the current population of exoplanets with atmospheric detections. With the continuing success of ground-based transit surveys, such as WASP and HATNet, more planets will be discovered that push to the extremes in quantities such as the host star metallicity and activity, and the planet's irradiation. These systems will help to characterise and clarify current trends being seen, not just in the secondary eclipse observations that I have focused on, but also in transmission spectroscopy,

phase curves and direct imaging studies. Continuing lines of enquiry will include: the nature of the activity–inversion correlation of Knutson et al. [2010]; the prevalence of TiO and VO in hot Jupiter atmospheres; reflecting and heat redistribution properties of giant planets and the deficiency of CH₄ in directly imaged planets.

Current ground- and space-based facilities (e.g. WHT, VLT, Hubble and Spitzer) have given us our first glimpses into the properties of exoplanet atmospheres, but with the next generation of telescopes and instruments there should be a great leap forward in the field. A new class of 30–40 metre telescopes (e.g. the E-ELT) will allow for high resolution transmission spectroscopy and phase curve observations at near infra-red and optical wavelengths [Snellen, 2013; Madhusudhan et al., 2014], while for directly imaged planets there will be a push to lower mass planets (i.e. sub-Jupiter and terrestrial planets) at moderate separations (< 5 au) from their hosts.

The use of extremely large aperture telescopes is not critical for secondary eclipse and transmission studies of giant planets, since the limitations here tend to be detector systematics rather than photon noise. For these objects, the extremely stable environment and moderate aperture (6.5 m) that will be offered by the James Webb Space Telescope (JWST) will allow atmospheric characterisation to an unprecedented level of detail [Clampin, 2010]. Unambiguous detections of several atmospheric species (e.g. H₂O, CO, CH₄, Na, K), Rayleigh scattering signatures and temperature structures will be possible for many planets in the current population. With JWST, atmospheric characterisation studies will also begin to include sub-Jupiter and terrestrial mass planets. Space-based transit surveys such as TESS and PLATO will provide ideal targets here, as they will be sensitive to bright systems hosting planets down to super-Earth sizes. However, detailed studies of such systems in their habitable zones, even for the most promising targets, will be difficult with JWST [Seager & Deming, 2010].

Even further into the future, the ultimate goal of exoplanet atmospheric work is to characterise an Earth-analog i.e. an Earth-like planet, orbiting at 1 au from a Sun-like star. This is not within the reaches of current and near-future instrumentation; atmospheric signals will be at the 10⁻⁶ level. However, techniques such as transmission spectroscopy, secondary eclipses and direct imaging, may well be the methods we use to detect the first evidence of habitability and life on exoplanets e.g. through vegetation signatures such as the red edge or biosignature gases such as O₂ and O₃.

Appendix A

Appendix

A.1 WASP-33 airmass trend plots




Figure A.1: Equivalent plot to Fig. 5.7, but here given for the green channel of ULTRACAM. The variations in the differential flux light curve seen in the red channel are also seen here.




Figure A.2: Equivalent plot to Figs. 5.7 and A.1, but here given for the blue channel of ULTRACAM. The differential flux variation seen in the red and green channels is present, with an additional trend due to the colour difference of comparison 1 and comparison 2.




Figure A.3: Equivalent plot to Fig. 5.11, but here given for the green channel of ULTRACAM. The $\sim 0.2\%$ pulsations of WASP-33 are clearly seen in the differential light curve (bottom left panel).




Figure A.4: Equivalent plot to Fig. 5.11 and A.3 but here given for the blue channel of ULTRACAM. The δ Scuti pulsations of WASP-33 can be seen in the differential light curve, with an additional slope due to the colour difference of WASP-33 and comparison 2. Note that in the plot of differential flux against airmass, a straight line fit (given by equation 5.1) would not be appropriate. This motivates the use of the airmass model with an extinction coefficient that varies linearly in time (equation 5.2).

Bibliography

- Agol, E., Cowan, N. B., Knutson, H. A., Deming, D., Steffen, J. H., Henry, G. W., Charbonneau, D., 2010, *ApJ*, 721, 1861
- Alonso, R., et al., 2004, *ApJ*, 613, L153
- Anderson, D. R., et al., 2010, *ApJ*, 709, 159
- Anderson, D. R., et al., 2011, *MNRAS*, 416, 2108
- Anderson, D. R., et al., 2013, *MNRAS*, 430, 3422
- Anderson, D. R., et al., 2014, *ArXiv e-prints*
- Auvergne, M., et al., 2009, *A&A*, 506, 411
- Bakos, G., Noyes, R. W., Kovács, G., Stanek, K. Z., Sasselov, D. D., Domsa, I., 2004, *PASP*, 116, 266
- Ballard, S., et al., 2010, *PASP*, 122, 1341
- Baraffe, I., Chabrier, G., Barman, T., 2010, *Reports on Progress in Physics*, 73, 016901
- Barman, T. S., Macintosh, B., Konopacky, Q. M., Marois, C., 2011, *ApJ*, 733, 65
- Barros, S. C. C., Pollacco, D. L., Gibson, N. P., Howarth, I. D., Keenan, F. P., Simpson, E. K., Skillen, I., Steele, I. A., 2011, *MNRAS*, 416, 2593
- Baskin, N. J., et al., 2013, *ApJ*, 773, 124
- Batalha, N. M., et al., 2011, *ApJ*, 729, 27
- Beerer, I. M., et al., 2011, *ApJ*, 727, 23
- Bennett, D. P., et al., 2007, *ArXiv e-prints*

Blecic, J., et al., 2013, *ApJ*, 779, 5

Borucki, W. J., et al., 2010, *Science*, 327, 977

Borucki, W. J., et al., 2012, *ApJ*, 745, 120

Bouchy, F., et al., 2010, *A&A*, 519, A98

Brown, D. J. A., Collier Cameron, A., Hall, C., Hebb, L., Smalley, B., 2011, *MNRAS*, 415, 605

Brown, T. M., 2001, *ApJ*, 553, 1006

Burrows, A., Hubeny, I., Budaj, J., Knutson, H. A., Charbonneau, D., 2007, *ApJ*, 668, L171

Burrows, A., Budaj, J., Hubeny, I., 2008, *ApJ*, 678, 1436

Butler, R. P., Marcy, G. W., 1996, *ApJ*, 464, L153

Butler, R. P., Marcy, G. W., Williams, E., Hauser, H., Shirts, P., 1997, *ApJ*, 474, L115

Butler, R. P., Marcy, G. W., Vogt, S. S., Apps, K., 1998, *PASP*, 110, 1389

Butler, R. P., Marcy, G. W., Fischer, D. A., Brown, T. M., Contos, A. R., Korzennik, S. G., Nisenson, P., Noyes, R. W., 1999, *ApJ*, 526, 916

Campbell, B., Walker, G. A. H., Yang, S., 1988, *ApJ*, 331, 902

Casertano, S., et al., 2008, *A&A*, 482, 699

Charbonneau, D., Brown, T. M., Latham, D. W., Mayor, M., 2000, *ApJ*, 529, L45

Charbonneau, D., Brown, T. M., Noyes, R. W., Gilliland, R. L., 2002, *ApJ*, 568, 377

Charbonneau, D., Knutson, H. A., Barman, T., Allen, L. E., Mayor, M., Megeath, S. T., Queloz, D., Udry, S., 2008, *ApJ*, 686, 1341

Charbonneau, D., et al., 2005, *ApJ*, 626, 523

Chilcote, J., et al., 2014, *ArXiv e-prints*

Christian, D. J., et al., 2006, *MNRAS*, 372, 1117

Christiansen, J. L., et al., 2011, *ApJ*, 726, 94

- Ciceri, S., et al., 2013, *A&A*, 557, A30
- Clampin, M., 2010, in Coudé du Foresto, V., Gelino, D. M., Ribas, I., eds., *Pathways Towards Habitable Planets*, vol. 430 of *Astronomical Society of the Pacific Conference Series*, p. 167
- Claret, A., 2000, *A&A*, 363, 1081
- Cochran, W. D., Hatzes, A. P., Butler, R. P., Marcy, G. W., 1997, *ApJ*, 483, 457
- Collier Cameron, A., Pollacco, D., Hellier, C., West, R., WASP Consortium, SOPHIE and CORALIE Planet-Search Teams, 2009, in Pont, F., Sasselov, D., Holman, M. J., eds., *IAU Symposium*, vol. 253 of *IAU Symposium*, p. 29
- 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
- Cosentino, R., et al., 2012, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, vol. 8446 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*
- Cowan, N. B., Agol, E., 2011, *ApJ*, 729, 54
- Cowan, N. B., Machalek, P., Croll, B., Shekhtman, L. M., Burrows, A., Deming, D., Greene, T., Hora, J. L., 2012, *ApJ*, 747, 82
- de Mooij, E. J. W., de Kok, R. J., Nefs, S. V., Snellen, I. A. G., 2011, *A&A*, 528, A49
- de Mooij, E. J. W., Brogi, M., de Kok, R. J., Snellen, I. A. G., Kenworthy, M. A., Karjalainen, R., 2013, *A&A*, 550, A54
- Deming, D., Seager, S., Richardson, L. J., Harrington, J., 2005, *Nature*, 434, 740
- Deming, D., Harrington, J., Seager, S., Richardson, L. J., 2006, *ApJ*, 644, 560
- Deming, D., et al., 2011, *ApJ*, 726, 95
- Deming, D., et al., 2012, *ApJ*, 754, 106
- Deming, D., et al., 2013, *ApJ*, 774, 95
- Demory, B.-O., et al., 2011, *ApJ*, 735, L12

Désert, J.-M., et al., 2011a, *ApJS*, 197, 11

Désert, J.-M., et al., 2011b, *ApJS*, 197, 14

Dhillon, V. S., et al., 2007, *MNRAS*, 378, 825

Doyle, L. R., et al., 2011, *Science*, 333, 1602

Dressel, L., 2012, *Wide Field Camera 3 Instrument Handbook for Cycle 21 v. 5.0*

Dumusque, X., et al., 2012, *Nature*, 491, 207

Dumusque, X., et al., 2014, *ApJ*, 789, 154

Enoch, B., Collier Cameron, A., Parley, N. R., Hebb, L., 2010, *A&A*, 516, A33

Enoch, B., Collier Cameron, A., Horne, K., 2012, *A&A*, 540, A99

Evans, T. M., et al., 2013, *ApJ*, 772, L16

Fazio, G. G., et al., 2004, *ApJS*, 154, 10

Fischer, D. A., Marcy, G. W., Butler, R. P., Vogt, S. S., Apps, K., 1999, *PASP*, 111, 50

Ford, E. B., 2005, *AJ*, 129, 1706

Ford, E. B., 2006, *ApJ*, 642, 505

Foreman-Mackey, D., Hogg, D. W., Lang, D., Goodman, J., 2013, *PASP*, 125, 306

Fortney, J. J., Lodders, K., Marley, M. S., Freedman, R. S., 2008, *ApJ*, 678, 1419

Fressin, F., Knutson, H. A., Charbonneau, D., O'Donovan, F. T., Burrows, A., Deming, D., Mandushev, G., Spiegel, D., 2010, *ApJ*, 711, 374

Garrido, R., Rodriguez, E., 1990, *Ap&SS*, 169, 205

Garrido, R., Garcia-Lobo, E., Rodriguez, E., 1990, *A&A*, 234, 262

Gelman, A., Rubin, D. B., 1992, *Statistical Science*, 7, 457

Gibson, N. P., Pont, F., Aigrain, S., 2011, *MNRAS*, 411, 2199

Gibson, N. P., et al., 2012, *MNRAS*, 422, 753

Gillon, M., et al., 2007, *A&A*, 471, L51

Grillmair, C. J., et al., 2008, *Nature*, 456, 767

Hansen, B. M. S., 2008, *ApJS*, 179, 484

Harrington, J., Luszcz, S., Seager, S., Deming, D., Richardson, L. J., 2007, *Nature*, 447, 691

Haswell, C. A., 2010, *Transiting Exoplanets*

Hatzes, A. P., et al., 2011, *ApJ*, 743, 75

Haywood, R. D., et al., 2014, *MNRAS*, 443, 2517

Hebb, L., et al., 2010, *ApJ*, 708, 224

Henry, G. W., Marcy, G. W., Butler, R. P., Vogt, S. S., 2000, *ApJ*, 529, L41

Herrero, E., Morales, J. C., Ribas, I., Naves, R., 2011, *A&A*, 526, L10

Holman, M. J., et al., 2010, *Science*, 330, 51

Huitson, C. M., Sing, D. K., Vidal-Madjar, A., Ballester, G. E., Lecavelier des Etangs, A., Désert, J.-M., Pont, F., 2012, *MNRAS*, 422, 2477

Huitson, C. M., et al., 2013, *MNRAS*, 434, 3252

Kass, R. E., Raftery, A. E., 1995, *Journal of the american statistical association*, 90, 773

Knutson, H. A., Charbonneau, D., Allen, L. E., Burrows, A., Megeath, S. T., 2008, *ApJ*, 673, 526

Knutson, H. A., Charbonneau, D., Burrows, A., O'Donovan, F. T., Mandushev, G., 2009, *ApJ*, 691, 866

Knutson, H. A., Howard, A. W., Isaacson, H., 2010, *ApJ*, 720, 1569

Knutson, H. A., et al., 2007, *Nature*, 447, 183

Knutson, H. A., et al., 2012, *ApJ*, 754, 22

Konopacky, Q. M., Barman, T. S., Macintosh, B. A., Marois, C., 2013, *Science*, 339, 1398

Kovács, G., et al., 2013, *A&A*, 553, A44

- Kurucz, R., 1993, ATLAS9 Stellar Atmosphere Programs and 2 km/s grid. Kurucz CD-ROM No. 13. Cambridge, Mass.: Smithsonian Astrophysical Observatory, 1993., 13
- Laughlin, G., Crismani, M., Adams, F. C., 2011, ApJ, 729, L7
- Léger, A., et al., 2009, A&A, 506, 287
- Lodders, K., 2003, ApJ, 591, 1220
- Lucy, L. B., Sweeney, M. A., 1971, AJ, 76, 544
- Machalek, P., McCullough, P. R., Burke, C. J., Valenti, J. A., Burrows, A., Hora, J. L., 2008, ApJ, 684, 1427
- Maciejewski, G., et al., 2010, MNRAS, 407, 2625
- Maciejewski, G., et al., 2013, AJ, 146, 147
- Madhusudhan, N., 2012, ApJ, 758, 36
- Madhusudhan, N., Seager, S., 2010, ApJ, 725, 261
- Madhusudhan, N., Amin, M. A., Kennedy, G. M., 2014, ArXiv e-prints
- Madhusudhan, N., et al., 2011, Nature, 469, 64
- Mahtani, D. P., et al., 2013, MNRAS, 432, 693
- Mandel, K., Agol, E., 2002, ApJ, 580, L171
- Marcy, G. W., Butler, R. P., 1996, ApJ, 464, L147
- Marois, C., Macintosh, B., Barman, T., Zuckerman, B., Song, I., Patience, J., Lafrenière, D., Doyon, R., 2008, Science, 322, 1348
- Mathur, S., et al., 2013, ArXiv e-prints
- Maxted, P. F. L., et al., 2013, MNRAS, 428, 2645
- Mayor, M., Queloz, D., 1995, Nature, 378, 355
- Mayor, M., et al., 2003, The Messenger, 114, 20
- Mayor, M., et al., 2011, ArXiv e-prints
- McCullough, P. R., Stys, J. E., Valenti, J. A., Fleming, S. W., Janes, K. A., Heasley, J. N., 2005, PASP, 117, 783

- Miller, G. R. M., et al., 2010, *A&A*, 523, A52
- Moffat, A. F. J., 1969, *A&A*, 3, 455
- Montalto, M., et al., 2012, *MNRAS*, 427, 2757
- Mortier, A., Santos, N. C., Sousa, S. G., Fernandes, J. M., Adibekyan, V. Z., Delgado Mena, E., Montalto, M., Israelian, G., 2013, *A&A*, 558, A106
- Nascimbeni, V., et al., 2013, *A&A*, 549, A30
- Nikolov, N., et al., 2014, *MNRAS*, 437, 46
- Noyes, R. W., Jha, S., Korzennik, S. G., Krockenberger, M., Nisenson, P., Brown, T. M., Kennelly, E. J., Horner, S. D., 1997, *ApJ*, 483, L111
- O'Rourke, J. G., et al., 2014, *ApJ*, 781, 109
- Perryman, M., 2011, *The Exoplanet Handbook*
- Pollacco, D., et al., 2008, *MNRAS*, 385, 1576
- Pollacco, D. L., et al., 2006, *PASP*, 118, 1407
- Pont, F., Aigrain, S., Zucker, S., 2011, *MNRAS*, 411, 1953
- Pont, F., Sing, D. K., Gibson, N. P., Aigrain, S., Henry, G., Husnoo, N., 2013, *MNRAS*, 432, 2917
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., Flannery, B. P., 1992, *Numerical recipes in FORTRAN. The art of scientific computing*
- Quintana, E. V., et al., 2014, *Science*, 344, 277
- Ranjan, S., Charbonneau, D., Désert, J.-M., Madhusudhan, N., Deming, D., Wilkins, A., Mandell, A. M., 2014, *ApJ*, 785, 148
- Rostron, J. W., Wheatley, P. J., Anderson, D. R., Collier Cameron, A., Fortney, J. J., Harrington, J., Knutson, H. A., Pollacco, D. L., 2014, *MNRAS*, 441, 3666
- Rowe, J. F., et al., 2006, *ApJ*, 646, 1241
- Rowe, J. F., et al., 2008, *ApJ*, 689, 1345
- Ruppert, D., 2011, *Statistics and data analysis for financial engineering*, Springer
- Sada, P. V., et al., 2012, *PASP*, 124, 212

- Schwarz, G., 1978, *Ann. Statist.*, 6, 461
- Seager, S., 2011, *Exoplanets*
- Seager, S., Deming, D., 2010, *ARA&A*, 48, 631
- Seager, S., Sasselov, D. D., 1998, *ApJ*, 502, L157
- Seager, S., Sasselov, D. D., 2000, *ApJ*, 537, 916
- Shporer, A., et al., 2014, *ApJ*, 788, 92
- Simpson, E. K., et al., 2010, *MNRAS*, 405, 1867
- Simpson, E. K., et al., 2011, *AJ*, 141, 8
- Sing, D. K., Vidal-Madjar, A., Désert, J.-M., Lecavelier des Etangs, A., Ballester, G., 2008a, *ApJ*, 686, 658
- Sing, D. K., Vidal-Madjar, A., Lecavelier des Etangs, A., Désert, J.-M., Ballester, G., Ehrenreich, D., 2008b, *ApJ*, 686, 667
- Sing, D. K., et al., 2011, *A&A*, 527, A73
- Sing, D. K., et al., 2013, *MNRAS*, 436, 2956
- Smith, A. M. S., Anderson, D. R., Skillen, I., Collier Cameron, A., Smalley, B., 2011, *MNRAS*, 416, 2096
- Smith, A. M. S., et al., 2012, *A&A*, 545, A93
- Snellen, I., 2013, in *European Physical Journal Web of Conferences*, vol. 47 of *European Physical Journal Web of Conferences*, p. 11001
- Southworth, J., 2011, *MNRAS*, 417, 2166
- Spiegel, D. S., Silverio, K., Burrows, A., 2009, *ApJ*, 699, 1487
- Stevenson, K. B., et al., 2012, *ApJ*, 754, 136
- Sudarsky, D., Burrows, A., Pinto, P., 2000, *ApJ*, 538, 885
- Swain, M. R., Vasisht, G., Tinetti, G., 2008, *Nature*, 452, 329
- Swain, M. R., Vasisht, G., Tinetti, G., Bouwman, J., Chen, P., Yung, Y., Deming, D., Deroo, P., 2009, *ApJ*, 690, L114

- Todorov, K., Deming, D., Harrington, J., Stevenson, K. B., Bowman, W. C., Nymeyer, S., Fortney, J. J., Bakos, G. A., 2010, *ApJ*, 708, 498
- Todorov, K. O., et al., 2012, *ApJ*, 746, 111
- Todorov, K. O., et al., 2013, *ApJ*, 770, 102
- Tripathi, A., et al., 2010, *ApJ*, 715, 421
- Udry, S., Santos, N. C., 2007, *ARA&A*, 45, 397
- von Essen, C., et al., 2014, *A&A*, 561, A48
- Wakeford, H. R., et al., 2013, *MNRAS*, 435, 3481
- Wheatley, P. J., et al., 2010, *ArXiv e-prints*
- Wolszczan, A., Frail, D. A., 1992, *Nature*, 355, 145
- Zahnle, K., Marley, M. S., Freedman, R. S., Lodders, K., Fortney, J. J., 2009, *ApJ*, 701, L20
- Zhao, M., Milburn, J., Barman, T., Hinkley, S., Swain, M. R., Wright, J., Monnier, J. D., 2012, *ApJ*, 748, L8