PIRATE: A Remotely Operable Telescope for Research and Education

Thesis

PIRATE: A Remotely Operable Telescope for Research and Education

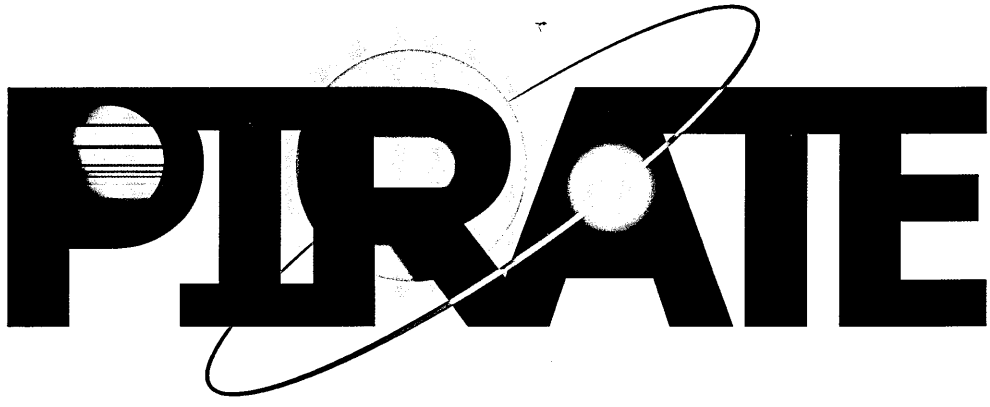

Stefan Holmes MPhys

Supervisors: Dr. Ulrich Kolb, Dr. Carole Haswell

PhD thesis submitted for the degree of Doctor of Philosophy

Department of Physical Sciences, The Open University

Submitted September 2012

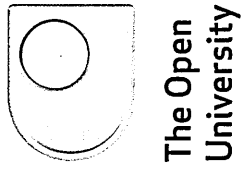

DATE OF SUBMISSION: 30 SEPTEMBER 2012

DATE OF AWARD: 4 JULY 2013

Abstract

This thesis introduces PIRATE, a remotely operable telescope facility for use in research and education, constructed from off-the-shelf hardware, operated by the Open University. It focuses on the PIRATE Mark 1 and PIRATE Mark 1.5 phases of operation; the telescope facility has been in the Mark 2 phase since September 2011. Situated at the Observatori Astronòmic de Mallorca, PIRATE is currently used to follow up potential transiting extrasolar planet candidates produced by the SuperWASP North experiment, as well as to hunt for novae in M31 and other nearby galaxies. It is operated by a mixture of commercially available software and proprietary software. In this thesis I discuss the hardware and its characterisation, problems with performing precision time-series photometry when using a German Equatorial Mount, and the PIRATE pipeline developed to assist with the fast and effective production of results from the follow-up observations of potential transiting extrasolar planets. I then go on to discuss the logistics of the follow-up program, and present the results from the PIRATE Mark 1 and 1.5 follow-up data, followed by the results from observational work on known transiting systems (such as WASP-12b). A discussion of areas for improvement and possible avenues for future work concludes the thesis.

Dedicated to my beautiful fiancée Kit for her boundless love and support.

Acknowledgements

I would firstly like to thank my supervisors Ulrich Kolb and Carole Haswell for all their efforts and assistance over the last few years. I couldn't have wished for better supervision; they did a sterling job of tempering my more tangential tendencies and always focused my attentions in the most profitable directions. My sincere thanks to you both.

I would also like to thank the PIRATE team and all of the contributors to the project: Ulrich for running the show, Vadim Burwitz, Rob Lucas, Juan Rodriguez, Jaume Andreu and all the staff at the OAM, other colleagues who have worked on the project at one time or another: Samantha Rolfe, John Barker, John Rostron; and my current colleagues now working on the project: Richard Busuttil and Jakub Bochinski - I leave it in safe hands.

Contents

List of Figures

Chapter 1

Introduction

1.1 The PIRATE project

PIRATE (Physics Innovations Robotic Astronomical Telescope Explorer)(Holmes et al., 2011) is a remote telescope facility situated on the Balearic island of Mallorca, near the village of Costitx, towards the middle of the island (2°57′03.34″E, 39°38′34.31″ N), approximately 162m above sea level. It is housed at the Observatori Astronòmic de Mallorca (OAM), an observatory with which The Open University (OU) has a long-held relationship, as it also hosts the OU's undergraduate astronomy residential schools. Funding for the PIRATE project was made available by a teaching innovation initiative (piCETL[1]) in order to create an affordable (~$150k) yet cutting-edge remote telescope facility built primarily from commercially available hardware and software, that would allow for real-time astronomy to be incorporated into current and future OU undergraduate courses. By building the facility from commercially available equipment and software, a blueprint was generated that can be copied across to similar facilities, with the intention of leading the expansion of real-time astronomy teaching at undergraduate level. A new neighbouring facility at the OAM, PTST, owned and operated by Hamburg Observatory, The University of Hamburg[2], was installed in June 2011 and uses similar hardware and the same software as PIRATE. The success of PIRATE as a tool for education helped secure funding for this project, and we hope that it will continue to lead the way in astronomy education at the undergraduate level for years to come.

[1]http://www8.open.ac.uk/opencetl/physics-innovations-centre-excellence-teaching-and-learning)
[2]http://www.hs.uni-hamburg.de

Alongside its use in education, the instrument has also proved to be a successful tool for astronomy research in the fields of transiting extrasolar planets, eclipsing binary stars, novae, and supernovae. This thesis focuses primarily on its use in research, though its use in education is detailed in the following subsection.

1.1.1 Use in education

PIRATE must cater for a relatively large student body with little or no prior experience in observational astronomy, in need of limited remote supervision by a tutor. Granting a sufficiently high level of access to the facility, such as the ability to initiate the opening and closing of the dome, is considered essential for the student learning experience. The adopted software solution, a combination of the commercial product ACP[3] (Denny, 2011) and the proprietary software described in Lucas and Kolb (2010), achieves the desired student access privileges without compromising the safe operation of PIRATE. Students familarise themselves with the telescope software by using the PIRATE simulator, which mimics the operation of the real facility, but uses only virtual hardware. After a tutor briefing session, conducted via audio communication and web interfaces, the student groups operate PIRATE successfully on their own. A (remote) night-duty astronomer is on call and contactable throughout the night.

PIRATE was deployed for OU undergraduate students for the first time in spring 2010 in a 10 week project which is part of the third level (third year) OU module S382 "Astrophysics". A total of about 30 students formed three groups with alternating access to PIRATE, for a total of 40 observing nights. Each group selected a suitable target source from the catalogue by Norton et al. (2007) of periodic variables found by SuperWASP and coincident with a ROSAT source. The groups then built up a long-term light curve of their target, such as the example shown in Fig. 1.1. Individual observing sessions were staffed by small observer teams of 2-4 students who kept in audio and text contact throughout the night. The collaboration of different observer teams in the larger groups on the same target source ensured the emergence of a usable data base to which all group members could develop a sense of ownership, even when occasional nights were clouded out and the observer teams on duty may not have succeeded in obtaining data themselves. A fuller account of the challenges and solutions for the PIRATE teaching project can be found in Kolb et al. (2010).

[3] www.dc3.com

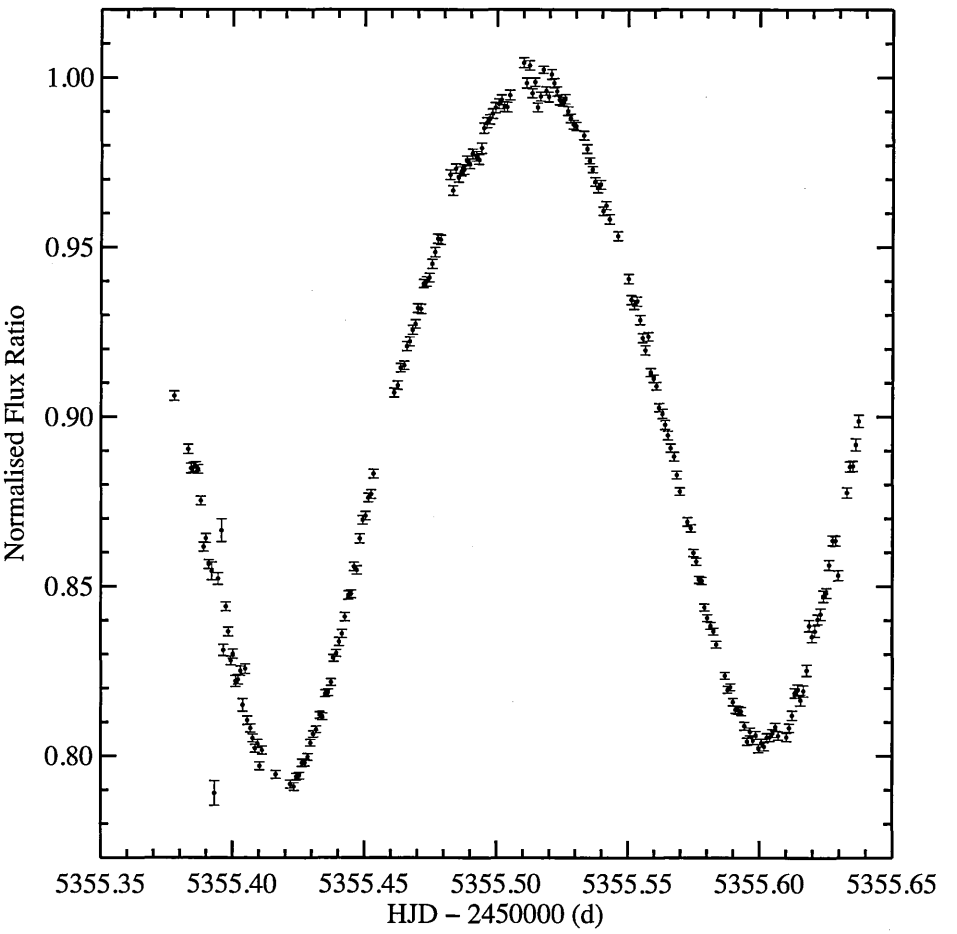

Figure 1.1: A light curve created from 206 I band images of 1SWASPJ164320.90+424433.4. The data were taken to fulfill project objectives for the first presentation of the OU's undergraduate module, S382, by OU undergraduate students, in April 2010.

In early 2012, PIRATE was also deployed for a second year module, SXP288, the PIRATE section of which forms part of a larger practical science program. This was very successful, and PIRATE's use in Open University teaching is being expanded further in the future.

1.1.2 Use in research

The nature of undergraduate module scheduling dictates that undergraduate students do not have year-round access to the telescope (they are currently limited to a 10 week run per annum for the 3rd level module, and a 2 week block for the 2nd level module - though this may be increased in the future), freeing up the rest of the year for

research applications. In this capacity PIRATE is currently used to perform follow-up photometry of SuperWASP Transiting Extrasolar Planet (TEP) candidates from the SuperWASP (Wide Angle Search for Planets) (Pollacco et al., 2006) survey, and participates in a nightly survey of M31, searching for extragalactic novae.

Wide-field ground-based transit surveys have proved a very effective tool for searching for TEPs. Their wide-field (many degrees) images of the night sky are used to produce light curves of millions of bright objects ($V \lesssim 15$), which are picked-through by transit-search algorithms (e.g. Collier Cameron et al. (2006)) that hunt for dips in brightness with depths and durations indicative of planetary transits. Such surveys have three main limitations in terms of hardware performance: photometric precision; temporal resolution, and angular resolution. These limitations arise from the ability to cover so much of the night sky regularly in the hunt for the aforementioned dips; the ability to do so inevitably leads to trade-offs. I will discuss the causes of such limitations in detail in Chapter 4. Unfortunately, they place restrictions on the likelihood that any individual planetary candidate discovered with the survey instrument is in fact a TEP. In fact, the vast majority of candidates uncovered in the survey data will prove to be false-positives. Small telescope facilities such as PIRATE, which have significantly better photometric performance, temporal resolution, and angular resolution (at the expense of the angular extent of the field-of-view) can be used to perform follow-up observations of TEP candidates, and either confirm their status as a false-positive, or show them to truly have a light-curve indicative of a possible TEP. Such follow-up work is always of great benefit then; one can either reject a candidate as a false-positive, and claw-back valuable large telescope spectroscopy observation time that would otherwise have been wasted attempting to confirm or deny the candidate, or one can possibly provide an enhancement to the initial parameter estimates of what may prove to be a new TEP.

1.1.3 The OAM

The OAM (Observatori Astronòmic de Mallorca) is an observatory on the island of Mallorca (2°57′03.34″E, 39°38′34.31″ N, 162m above sea level). On site is a planetarium which runs an educational service for schools on the island, and also hosts a weekly planetarium show for interested adults. Also at the observatory are a set of teaching domes that the Open University uses to host its residential schools, where students of the OU's undergraduate modules attend for week-long visits and use the on-site Meade telescopes to engage in practical astronomy. PIRATE's location at the

Month	Median Relative Humidity (RH)	σ_{RH}
January	72.4%	12.3%
February	69.6%	13.7%
March	70.8%	15.0%
April	66.5%	17.1%
May	60.8%	18.1%
June	58.6%	18.0%
July	60.0%	16.9%
August	58.8%	19.1%
September	68.2%	15.5%
October	69.9%	14.8%
November	76.4%	11.8%
December	71.4%	11.4%

Table 1.1: Monthly median relative humidity values and their corresponding standard deviation at the OAM

OAM was chosen because of this pre-existing relationship with the observatory. It was not chosen because the observatory is an excellent location for scientific observation. Fig. 1.2 is a reproduction of Fig. 3 from Bochinski et al. (2012). It details the breakdown of clear, cloudy, very cloudy, and rainy nights at the OAM in a calendar year, where the categories correspond to those defined in Table 2.3. As can be seen, the OAM gets approximately half a year's worth of clear nights per year. Unfortunately, this figure is typically cut down to around 100 observable nights per year due to excessive humidity (relative humidity $> 90\%$) at the site that forces the dome to remain shut with dehumidifers on. Monthly median humidities and their standard deviations can be found in Table 1.1. The site also suffers with light pollution to the South West from the island's capital, Palma, and a flightpath into Palma that passes directly overhead. On an excellent night, seeing of 1.5-2″ is achievable, though one is usually happy with 2-3″.

1.1.4 Other similar facilities

Remotely operated and robotic telescope facilities are becoming prominent in the world of professional astronomical research, due to the ever decreasing costs of small research-grade hardware. PIRATE has a a new neighbour in the form of PTST, which sits in a neighbouring dome and utilises a 24" PlaneWave CDK24 in a larger 4.5m diameter dome. It uses the same software as PIRATE, and the success of PIRATE helped secure funding for its owner and operator, The University of Hamburg. It was

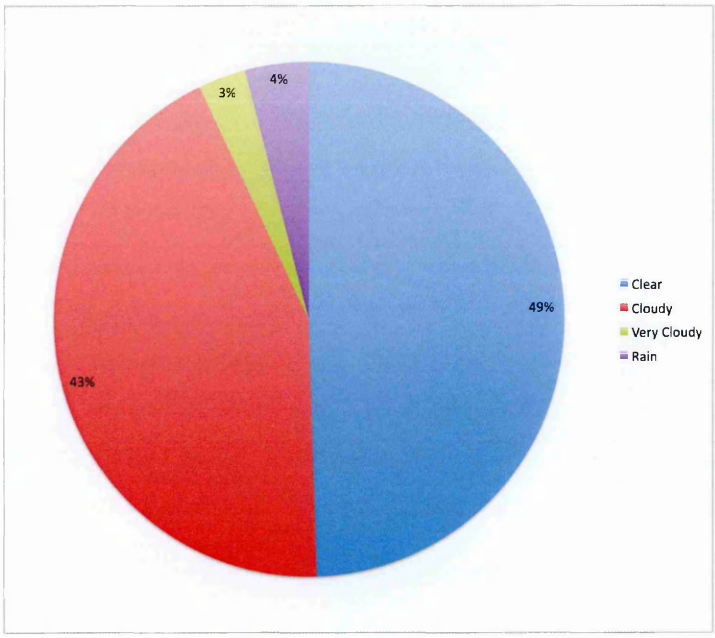

Figure 1.2: Reproduction of Fig. 3 from Bochinski et al. (2012) detailing the breakdown of nights by sky clarity according to the Boltwood classification scheme (see Table 2.3) in a calendar year.

installed in the Summer of 2011, and its primary use is intended to be schools and undergraduate teaching.

Similar follow-up photometry work is performed in the southern sky (using the data from the southern SuperWASP instrument, SuperWASP South) by TRAPPIST (Jehin et al., 2011). TRAPPIST was installed in April 2010 at the ESO La Silla Observatory. It has a 24" f/8 Ritchey-Chrétien optical tube assembly with a 2k × 2k Finger Lakes imager (resulting in a field-of-view of $22' \times 22'$) and employs the use of a special "$I + z$" filter that is well suited to transiting exoplanets.

1.1.5 Thesis summary

In the rest of this chapter I will briefly outline the history of exoplanet detection and introduce various methods employed to detect them.

In Chapter 2 I introduce the various iterations of PIRATE hardware that have been in place over the last 4 years, and describe the software required to operate PIRATE. I then detail the commissioning work performed to characterise the hardware.

Chapter 3 explains the process of generating the light curves produced from PIRATE data that are used to investigate potential SuperWASP planetary candidates.

In Chapter 4 I introduce the SuperWASP hardware and the experiment as a whole, and discuss the follow-up program used to determine astrophysical false-positive from genuine planetary candidates.

In Chapter 5 I detail the typical types of astrophysical false-positive seen in the SuperWASP candidate lists, and I reproduce the results from the observational follow-up program up to September 2011.

Chapter 6 describes the observations made of known planetary systems, including case studies of ephemeris refinement for WASP-12b, and modelling of the transit of WASP-10b.

In Chapter 7 I discuss project as a whole, including areas for future improvement to both the hardware and the data reduction routines.

1.2 Exoplanets and their discovery

1.2.1 Brief history of exoplanet science

The existence of planets around other stars in our Galaxy had long been suspected (Struve, 1952), if only from the natural desire to extend the properties of our own Solar System to that of other stars in our Galaxy. Their detection, however, had proved enormously elusive, and so their suspected ubiquity throughout the Galaxy had never been taken for granted. Before the first confirmed detection of an exoplanet in the 1990s, many unsuccessful claims were made to the first exoplanet detection. The most famous of these is the case of Peter van de Kamp's spurious claim of a Jovian-mass body orbiting Barnard's Star (van de Kamp, 1969). Van de Kamp laid claim to the detection of a planetary candidate orbiting Barnard's star through astrometric measurements: the measurement of a star's apparent position on the celestial sphere. He proposed that the proper motion of Barnard's star displayed a harmonic wobble - indicative of a reflex motion of the star due to the gravitational influence of a planetary body. This was later debunked when Hershey (Hershey, 1973) showed all of the stars studied in the van de Kamp photographic plates displayed the same systematic wobble - the reflex motion of Barnard's Star was in fact a manifestation of an instrumental defect. The first real detection of a planetary mass object outside of our Solar System came in the form of the planet-like objects around pulsar PSR B1257+12 (Wolszczan and Frail, 1992). These objects were discovered with radio-domain pulse time delay measurements, and belong to a system that remains unusual and unique

even today. Whilst these objects are clearly interesting, the discovery that truly gave birth to the field of exoplanetary astrophysics was of the Jupiter-mass object around 51 Pegasi (Mayor and Queloz, 1995). Mayor and Queloz conducted a radial velocity survey of 142 G and K dwarf stars with the ELODIE spectrograph in 1994, and found one of the stars (51 Peg) to exhibit a sinusoidal radial velocity variation suggestive of a companion planet. This would prove to be the first detection of an extrasolar planet orbiting a solar-type star. The technique, now most commonly referred to simply as 'The Radial Velocity Method' has proved to be especially fruitful, with the vast majority of the total of 777[4] known planets (as of 12th August 2012) planets detected by this method . The next significant milestone in the field (and the most pertinent to this discourse) came with the detection of the occultation of HD 209458 by its companion planet, HD 209458b (Charbonneau et al., 2000). The planet is said to 'transit' its host star, and thus this method of detection is known as 'The Transit Method'. This technique is typically performed with photometric measurements (although transmission spectroscopy has now been used to detect extrasolar planetary atmospheres, see e.g. (Charbonneau et al., 2002)). As the planet passes in front of its host star, it blocks some of its light, producing a periodic dimming synchronised with the orbital period of the planet - star system. The real challenge with such a detection is the scale of the dimming; typically a 2% dimming or less for a Jupiter-sized planet. The observed light curve must therefore demonstrate a point-to-point precision that is better than or at least comparable to the transit depth, a precision not usually needed for other types of variable star photometry. The obvious limitation of the transit method is that an orbital inclination edge-on to the observer's line of sight is a prerequisite for a transit to occur. The implication, therefore, is that TEPs are only a subset of the Galactic extrasolar planet population as a whole. This obviously makes them harder to find in a survey, and the detection rates of each method thus far uphold this statistical feature. The obvious way to counteract this statistical impairment is to observe more stars; wide-angle surveys are required. This limitation (the prerequisite for particular geometric alignment) is also the transit method's greatest strength, as it unlocks extra parameters for determination. The radial velocity technique, as its name implies, can only determine the reflex motion of a star along the observer's line of site. It therefore contains no information about the orbital inclination of the system, providing only an upper limit to the planetary mass. The transit method tightly constrains the orbital inclination, greatly refining the mass estimate. Secondly, the transit method unlocks the radius of the planet,

[4]http://exoplanet.eu

as the area of the stellar disc obscured by the planet determines the extent of flux loss throughout the transit. Thus, the two methods are complementary, together providing a planetary mass and radius; and consequently yielding an estimate for the density of the planet.

1.2.2 Methods of detecting exoplanets

Planetary bodies in our Solar System are visible to us primarily due to the light they reflect back at us from the Sun. They are thermally emissive bodies too, but generally we are able to pick them out in the night sky due to their reflected sunlight. As a proportion of the total emitted flux of the Sun, however, the amount of light reflected is miniscule. Let us compare the luminosity of Jupiter due to reflected sunlight with the total luminosity of the Sun. The incident flux (energy per unit area per unit time) at the distance of Jupiter from the Sun is given by:

$$F_{\odot} = \frac{L_{\odot}}{4\pi d_{\mathrm{J}}^2} \tag{1.1}$$

where $F_{\odot}$ and $L_{\odot}$ are the flux at distance d_{J} and luminosity of the Sun, respectively, and d_{J} is the distance of Jupiter from the Sun. Jupiter presents to the Sun a cross-sectional area, not its total surface area, so the peak luminosity of Jupiter due to reflected sunlight alone is:

$$L_J = A\pi R_J^2 F_{\odot} \tag{1.2}$$

where R_J is the radius of Jupiter, and A is the albedo of the planet, a quantity that represents the fraction of light reflected and not absorbed by the body. The ratio of the reflective luminosity of Jupiter to the luminosity of the Sun therefore is:

$$\frac{L_J}{L_{\odot}} = \frac{AR_J^2}{4d_J^2} \tag{1.3}$$

At a distance of 7.8×10^{11}m (its orbital semi-major axis) from the Sun, a radius of $\sim 7.1 \times 10^{7}$m, and opting for an albedo of 0.5, the luminosity of Jupiter is about a billionth of that of the luminosity of our Sun. We are still able to see Jupiter with the naked eye however; but only once the Sun is set, i.e. there is no direct sunlight incident on the Earth's night-side surface. We do not have such a luxury

when looking at extrasolar systems; we must always contend with direct starlight from the system's host star, and isolating a point source a billion times fainter than its host star is exceptionally troublesome. Consequently, most methods of detecting exoplanets are *indirect*. Such methods involve measuring the effect of the planetary body on the host star (and hence the host star's emitted light), rather than imaging the light from the planet itself. Direct imaging of exoplanets is possible and has been achieved in recent years (see e.g. 1RXS J160929.1–210524, Lafrenière et al. (2008)), but it is technologically limited, and remains at the frontier of the subject; it is not discussed in this introduction.

1.2.2.1 Astrometry

Astrometry is the science of measuring stellar positions on the celestial sphere. In order to infer the presence of a planet via astrometric measurements, one must look for the reflex motion of the star due to the planet, as the two objects orbit their common barycentre. Again, we don't expect to be able to isolate the light from the planet, but we can use its gravitational influence on its host star to infer its presence. Assuming a circular orbit, it is apparent that the orbital semi-major axis of the star about the barycentre is calculable from the maximum astrometric displacement of the star, so long as we know the distance to the star. Given that, for a simple two-body star-planet system

$$m_p a_p = m_* a_* \tag{1.4}$$

where m_p,a_p,m_*,a_* are the masses and semi-major axes of the planet and star respectively, and we are looking to detect the reflex motion of the star, it is evident that this method is most sensitive to high mass planets in large orbits around low mass stars. We can use the small angle approximation to arrive at:

$$\mu \simeq \frac{a_*}{d} \tag{1.5}$$

where μ is the astrometric displacement, and d is the distance to the star. We can, once again, insert some sample numbers into this expression to determine the level of astrometric precision required to detect the presence of Jupiter around the Sun, were we to view it from a distance of 10pc. Combining the above two equations we arrive at:

$$\mu \simeq \frac{m_p a_p}{m_* d} \tag{1.6}$$

Plugging in approximate values for the mass of Jupiter ($m_p \simeq 1.9 \times 10^{27}$kg), the Sun ($m_* \simeq 2 \times 10^{30}$kg), the semi-major axis of the orbit of Jupiter ($a_p = 7.8 \times 10^{11}$m), and $d = 10$pc ($d \simeq 3.09 \times 10^{17}$m), we find that $\mu \simeq 2.4 \times 10^{-9}$radians, or $\sim 4.9 \times 10^{-4}$ arcsec. This is clearly a challenging measurement. As will be discussed in more detail later, PIRATE's plate scale (or angular resolution) has typically been of order 1 arcsec pixel^{-1}(there have been various values throughout its lifespan thus far), so this astrometric displacement would be all but immeasurable by an instrument like PIRATE.

1.2.2.2 Gravitational lensing

In the event that a star passes behind a foreground star, the foreground star may act as a gravitational lens, bending the light from the background star in its gravitational field, resulting in the amplification of the background star. This effect is known as gravitational microlensing and was first predicted by Einstein (1936). In the event that the foreground star hosts an extrasolar planet, a further, secondary amplification of the light from the background star may occur, presenting an opportunity to detect an exoplanet around the foreground star. The light curve of such an event thus appears primarily as a peaked distribution (due to the primary amplification of the light from the background star) with a secondary, smaller peak during this event due to the lensing from the planet.

The first planet to be discovered through this method was during the microlensing event OGLE 2003-BLG-235/MOA 2003-BLG-53 (Bond et al., 2004), where a 7 day low amplitude deviation in the light curve due to a single-body lens system was detected. When modelled, this produced a mass ratio of 0.0039 for the lensing star-planet system, indicating a planetary mass of $1.5 M_J$ and an orbital separation of ~ 3 AU.

Gravitational microlensing is sensitive to low mass planets in large orbits, thus probing a unique part of exoplanet parameter space. However, each microlensing event is unique and never to be repeated, so follow up observations using alternative methods are required; though systems discovered through microlensing are usually distant and faint, and thus currently out of the reach of current photometric and spectroscopic follow-up instruments.

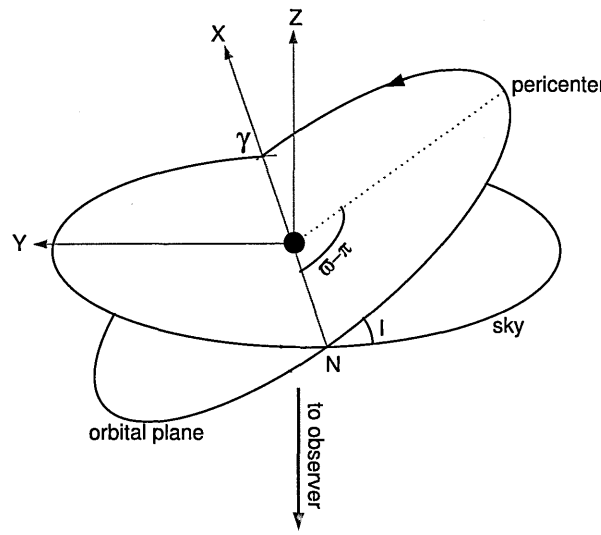

Figure 1.3: Diagram taken from Beaugé et al. (2007) representing the Cartesian coordinate system used in determining the radial velocity of the planetary body as a function of its orbital parameters.

1.2.2.3 Radial velocity measurements

The Radial Velocity method (hereafter 'RV') is analogous to the astrometric method in that we look solely at the reflex motion of the host star to infer the gravitational influence of a planet. The RV method differs from the astrometric method however in that we do not measure *the position* of the star around its barycentric orbit, but the component of its velocity both away from and towards us, along the line-of-sight. This is achieved by measuring the Doppler shift of known spectral lines, from which one can determine the radial velocity component along the line-of-sight. I briefly reproduce here Beaugé et al.'s (2007) description of the extraction of orbital parameters from RV measurements. Fig. 1.3 shows the Cartesian co-ordinate system adopted to describe an orbit of non-zero eccentricity, for a simple 2-body star-planet system. It describes an 'astrocentric' reference frame, where the centre of mass of the star is stationary, and thus we are currently seeking to describe the velocity component of the planet.

The z-axis is orientated along the line of sight between the central object and the observer. The pericenter is the point at which the planet is closest to the star. The inclination here is defined as the angle between the plane of the orbit and the plane of the sky. Consequently, the x-axis is defined by the two points at which the plane of the orbit and the plane of the sky intersect; the point where the planet crosses the sky plane and is moving towards the observer being γ. This is the 'ascending node';

the opposite point, labeled 'N' is the 'descending node'. The angle ϖ is the angle from the point γ to the pericenter, in the plane of the orbit. $(\varpi - \pi)$ is labeled on the diagram, and is known as the 'argument of periapsis'. Not shown in the diagram is the 'true anomaly', which we denote '$\theta(t)$'. This is the angle that measures how far around the orbit the orbiting body is. This is measured from the pericenter in the direction of travel of the body, and is the only quantity that is a function of time. The next step involves transforming the Cartesian unit vectors in the plane of the orbit to the above Cartesian coordinate system in terms of $\theta(t)$, ϖ and i, and subsequently determining the astrocentric velocity vector in terms of these quantities, which I will skip and simply reproduce the results of Beaugé et al.. The velocity component of the planet along the line of sight, in astrocentric coordinates is:

$$v_z = -\frac{2\pi a \sin i}{P\sqrt{1-e^2}}(\cos(\theta(t) + \varpi) + e\cos\varpi) \tag{1.7}$$

where a is the orbital semi-major axis and P is the orbital period. However, as previously mentioned, we do not wish to know the velocity of the planet, as we cannot detect its light next to the overbearing glare of its host star. We need to know the reflex velocity of the star due to the planet, and so we must make two further transformations. The first of these involves transforming from an astrocentric to a barycentric reference frame, so that we may determine the star's motion with respect to the centre of mass of the system. The second involves transforming to the frame of the observer, who observes not only the reflex motion of the star due to the planet, but also the radial velocity component of the star due to its motion around the Galaxy. We can take care of the first of these transformations with a few simple vector translations. The velocity vector of the planet in the astrocentric frame (of which the z-component was presented in Eq. 1.7) is easily broken down into barycentric planetary and stellar velocities as thus:

$$v = v_{p,bary} - v_{*,bary} \tag{1.8}$$

where v is the astrocentric planetary velocity, and $v_{p,bary}$ and $v_{*,bary}$ are the barycentric planetary and stellar velocity vectors, respectively. Differentiating the general form of Eq. 1.4 with respect to time yields:

$$m_* v_{*,bary} = m_p v_{p,bary} \tag{1.9}$$

We can then combine the above to relate the barycentric stellar velocity to the astrocentric planetary velocity, via the reduced mass of the planet:

$$v_{*,bary} = \left[\frac{m_p}{m_p + m_*}\right] v \tag{1.10}$$

The second translation to the frame of observer is very simple. Given that the orbital period of the planet is likely very small compared to the Galactic period of the star (the Sun's Galactic period - the time taken to do one circuit around the centre of the Galaxy - is around 250 million years), we can treat this is a constant offset throughout the duration of one planetary orbit. To investigate the reflex motion of the star due to the planet, we must therefore subtract the constant offset introduced by the star's orbit around the centre of the Galaxy. Bringing all of these elements together, we can now produce a final expression for the RV oscillations of a host star due to a single planet companion:

$$v_z(t) = v_{0,z} + \frac{2\pi a m_p \sin i}{(m_p + m_*)P\sqrt{1-e^2}}(\cos(\theta(t) + \varpi) + e\cos\varpi) \tag{1.11}$$

where $v_{0,z}$ is the constant offset due to the Galactic orbit of the host star, and the other symbols retain their previous meaning. On quick inspection, we can say a few things about Eq. 1.11. The first is that, all other parameters being fixed, eccentric orbits are easier to detect, as they have a greater RV amplitude. The second point is that we have one unknown in the equation, that we cannot determine from RV measurements alone, namely the inclination, i. Rather than being able to determine m_p from RV measurements, we may only find the quantity $m_p \sin i$, and thus a *lower limit* to the planetary mass. We are able to measure the inclination, but only if the planet transits its host star; more on this in subsection 1.2.2.4. Being able to determine an upper limit to the orbiting object's mass is still extremely useful, as it allows us to place the mass of the object within a planetary regime, or, alternatively, dismiss the object as a low mass star or brown dwarf.

As previously mentioned, Mayor and Queloz (1995) put this technique to work to discover the very first exoplanet orbiting a Solar-type main sequence star. They surveyed 142 G and K dwarf stars in 1994, using the fibre-fed Echelle spectrograph, ELODIE, at the Observatory d'Haute Provence (OHP). This gave them a typical precision of 13 $ms^{-1}$for each RV measurement. In Fig. 1.4 I repeat their phase-folded data, with the fit to $v_z(t)$ of the 4.23 day orbit of 51 Peg b, after the systemic velocity,

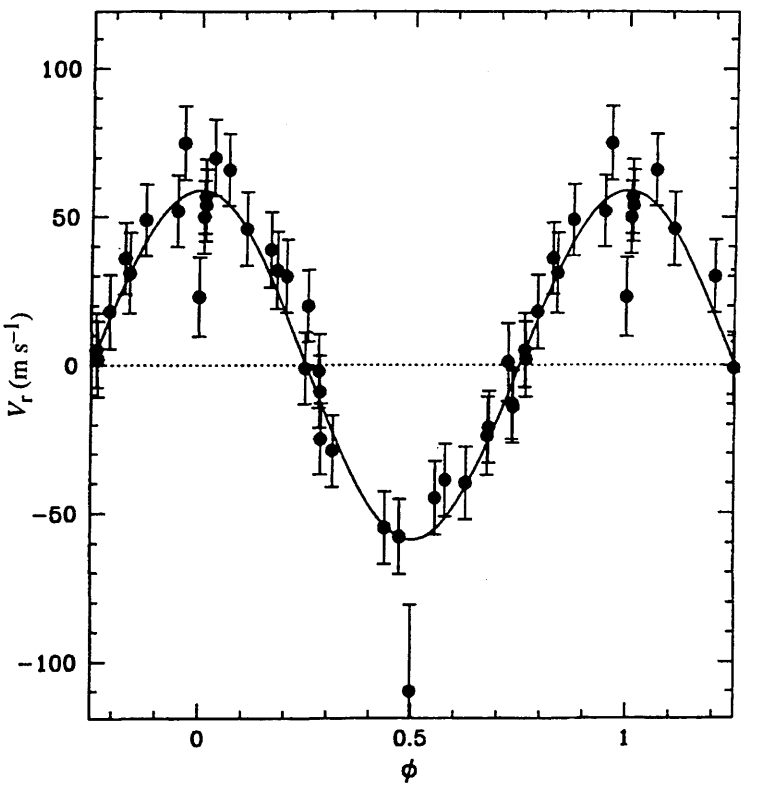

Figure 1.4: Orbital fit to $v_z(\phi)$, after removing the systemic velocity $v_{0,z}$ and phase-folding about the detected orbital period, for 51 Peg b, from Mayor and Queloz (1995).

$v_{0,z}$ was fitted and removed. Their fit yielded $m_p \sin i = 0.47 \pm 0.02$ M_J, which is remarkable for two reasons. Firstly, it represents a discovery that was the first of its kind, and gave birth to the field of exoplanetary science. Secondly, it represented a discovery that challenged all that was previously supposed about planetary formation theory. With such a short period, the Jovian mass object orbits around its host star at a distance of ~ 0.05 AU, approximately 6 times closer to its host star than Mercury is to the Sun. Many more planets of Jovian mass orbiting very close to their host stars were subsequently discovered in the wake of the initial discovery, which gave rise to a category of planets we now refer to as 'Hot Jupiters'; so called because of the high level of irradiation they experience from being so close to their host stars.

1.2.2.4 Transits

If the orbital plane of a distant exoplanet is aligned such that it passes in between the host star and the observer, the planet is said to 'transit' its host star. This offers

up another chance to indirectly detect the presence of a planet, as the planet blocks some of the light from the star during the transit, resulting in a periodic dimming of the star's brightness. This method is very advantageous, as it allows us to determine the radius of the exoplanet. If we assume the planet and star are at the same distance from the observer (a sound assumption as a typical Galactic interstellar distance is so much longer than the average orbital distance of planet from its host star), we can deduce the ratio of planetary radius to stellar radius simply from the amount of flux lost during the transit:

$$\frac{\Delta F}{F} = \frac{R_p^2}{R_*^2} \tag{1.12}$$

where F is the flux from the host star, and ΔF is the flux lost during transit. We can also use features of the transit light curve to determine the inclination of the orbital system, and thus constrain any measurements of planetary mass that come from RV measurements; so the RV technique and transit technique are highly complementary.

1.2.2.5 Transit timing variations

This detection method is secondary to the transit method in that it requires a prior detection of a transiting exoplanet. The method focuses then on multiple-planet systems, where there are not only gravitational interactions between the planets and their host star, but also inter-planet interactions. Transit Timing Variations (TTVs) are detections of variation in the *timing* of a planetary transit. The gravitational interaction of a secondary planet in a system with the primary (already discovered planet) can cause the primary planet to arrive either earlier or later than expected if no other planets existed. An object of planetary mass can be inferred from the extent of the timing variation. This method is particularly useful for finding low mass planets, and typically allows for the detection of lower mass planets than that currently available to RV detection.

1.2.3 The transit method

The transit method provides a key to unlocking exoplanet system parameters that can otherwise remain locked if we just use RV measurements. To obtain system parameters, we first need to make successive measurements of the brightness of the star as the transiting planet passes in front of it. This is known as a transit 'light curve'. In this section I show in detail which parameters are measurable from a transit

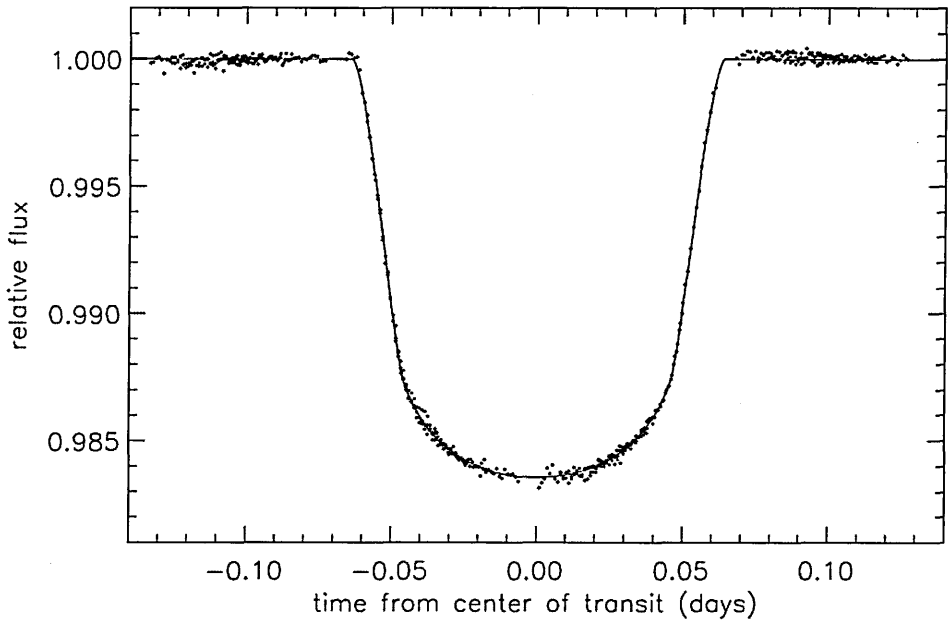

Figure 1.5: The light curve of HD209458b as captured by the Hubble Space Telescope (Brown et al., 2001)

light curve, and how we can go about deducing values for them. Fig. 1.5 shows the light curve of the first known transiting exoplanet, HD 209458 b, taken from Brown et al. (2001) . On initial inspection it is clear the transit shape and profile is more intricate and characterful than a simple box profile. These variations from a box profile can be fitted to, and they can yield a range of parameters. I have already stated that the ratio of planet radius to stellar radius is immediately available from the flux deficit during the transit, and this is the most simple parameter to determine. For the following section I largely follow the discourse of Haswell (2010), which in turn draws upon the work of Sackett (1999) and Mandel and Agol (2002).

1.2.3.1 Transit duration

Another easily identifiable feature of a transit dip is its duration, which we define in terms of contact points between the limb of the planet and the limb of the star. First through fourth contact are defined below, depicted in Fig. 1.6(a). First contact is the point at which the planetary and stellar limbs first co-incide (when seen from the observer's position). Second contact is the point when the whole of the planet's disc just resides within the stellar disc. Third contact is analogous to second contact; it is the last point at which the whole of the planet's disc resides within the stellar disc.

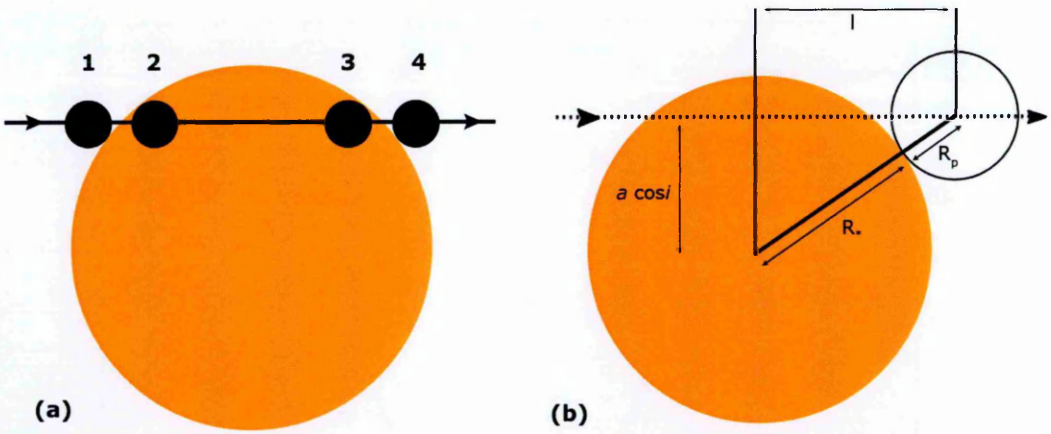

Figure 1.6: (a) The four contact points of a planetary transit. (b) schematic showing the definition of the impact parameter, b and the length, l

Fourth contact is the point at which the planetary and stellar limbs last coincide. It is typical to define the transit duration as occuring between first and fourth contact. Fig. 1.6 (b) again depicts a transiting system face-on, and introduces the definition of the impact parameter, b, which is defined as:

$$b = a \cos i \tag{1.13}$$

where a is again the semi-major axis of the orbit, and i is the inclination. Also defined is l, the projected separation along the planet's path through the sky-plane between the planet's centre and the star's centre. It should be apparent from Fig. 1.6 (b), that, via Pythagoras, the length l can be written:

$$l = \sqrt{(R_* + R_P)^2 - b^2} \tag{1.14}$$

Note that we can also use this expression to show the condition that must hold for a transit to occur. l becomes imaginary if the content within the square root becomes negative, i.e. $b > (R_* + R_p)$. This is equivalent to no transit occuring. For a transit to occur $b < (R_* + R_p)$ must hold true.

If we assume the orbit is circular, then the transit duration, as a fraction of the total period, is equal to the angle, α (in radians) around the orbit travelled by the planet as it moves from first to fourth contact, divided by 2π. The angle α therefore subtends the length of arc (around the orbit) bounded by the positions of first and

fourth contact, the length between these two points being $2l$. As a consequence:

$$\sin\left(\frac{\alpha}{2}\right) = \frac{l}{a} \tag{1.15}$$

$$\therefore \alpha = 2\arcsin\left(\frac{l}{a}\right) \tag{1.16}$$

$$\therefore T_{dur} = P\left(\frac{\alpha}{2\pi}\right) = \frac{P}{\pi}\arcsin\left(\frac{l}{a}\right) = \frac{P}{\pi}\arcsin\left(\frac{\sqrt{(R_* + R_p)^2 - b^2}}{a}\right) \tag{1.17}$$

For the situation where we have sucessfully estimated the host star's spectral type (and thus have a stellar mass and radius estimate) it should be immediately apparent that we can use a measurement of the transit duration to constrain the orbital inclination, i. If we assume $M_* \gg M_p$ then, from capturing successive transits we can obtain a period and thus an estimate of the semi-major axis, a, via Kepler's third law:

$$P^2 = \frac{4\pi a^3}{G(M_* + M_p)} \tag{1.18}$$

We can then, together with our estimate of the ratio of the planet radius to stellar radius (from the transit depth) and our measurement of the transit duration, T_{dur}, obtain a value for b, the impact parameter, and hence obtain a value for the orbital inclination of the system.

1.2.3.2 The shape of the transit light curve

So far the transit duration and flux deficit (transit depth) have been discussed. However, after inspecting Fig. 1.5, one should arrive at the conclusion that there are further features of the transit light curve not yet touched upon. For example, whilst we have defined the duration of the whole of the transit, the duration of the time between contact two and three is also variable, and is certainly not a fixed ratio of the total duration of the transit, but a function of b and R_p. Or, put another way, and given that we have a symmetrical transit; the time between first and second (T_{1-2}), and (via symmetry) third and fourth (T_{3-4}) contact (or 'ingress' and 'egress', respectively) as a fraction of the total transit duration is a function of both b and R_p. Also obvious is that the light curve is not flat between second and third contact. We might expect that, given the area of the stellar disc occluded by the planet remains constant between these two contact points, that the flux lost during transit might

also remain constant. Clearly, this is not the case, and the answer lies with stellar limb darkening. This is an effect we see when looking at our Sun, and therefore we assume to be true when looking at other stars, even though they are apparently point sources with (in most cases) no extent to their stellar disc. When looking at the Sun, it is apparent that the brightness of the Sun's disc drops off towards the edges. This is due to the varying optical depths within the Solar atmosphere from which we see light emitted, which varies as a function of the position on the Solar disc. Due to the nature of spherical geometry, the optical depth for any given physical depth within the star increases towards the limb of the star, and so, any photons we see emanating from the limb of the star have travelled to us from a greater optical depth. The probability of a photon travelling to us from a given optical depth decreases exponentially as the optical depth increases, and thus the limbs of the Sun appear darker. This explains the curvature we see in the light curve of HD 209458 b, as more flux is blocked when the planet passes over the centre of HD 209458 than at the limbs. A mathematical treatment of limb darkening will be given later in subsection 1.2.3.4, but for now, we move onto the question of how we can create a model light curve. I detail the analytical approach of Mandel and Agol (2002), following the discourse in Haswell (2010), which differs from the treatment of Mandel and Agol (2002) in the use of different variable notation. I will stick to the notation of Mandel and Agol (2002) whilst following the derivation of Haswell (2010). This analytical function for the amount of flux lost during a planetary transit is particularly relevant, as its use is employed in the Markov Chain Monte Carlo light curve fitting code (Collier Cameron et al. (2007), Pollacco et al. (2008), Enoch et al. (2010)) employed by the SuperWASP consortium, which is routinely used in this thesis and is discussed in Chapter 4.

We wish to arrive at an analytical function for the amount of flux lost, ΔF, by a star during a planetary transit, as a function of time. To do this, we need to know the area of stellar surface blocked during the different phases of a transit, and we need to know the amount of light emitted from the star, that is subsequently blocked, i.e. we need a 'limb darkening law' that explains how the amount of emitted flux from the star changes over the disc. For this treatment I will derive the area of stellar surface occluded by a transiting planet as a function of time, and discuss the implementation of a mathematical limb darkening 'law'.

1.2.3.3 Fraction of stellar disc occulted as a function of time

There is an analytical solution for the area of star eclipsed by a transiting planet that is a function of R_*, R_p, and z, the separation of centres in the sky plane (normalised to the radius of the star). This involves only the 2-D silhouettes of both the star and its companion in the sky plane. First, we must define the separation of centres as a function of time. For simplicity, circular orbits alone are considered here. The separation of centres is the separation *in the sky-plane*, so we must consider the projection of the inclined orbit onto the sky-plane. We start by defining the orbital phase angle, ϕ:

$$\phi = \frac{\omega t}{2\pi} \tag{1.19}$$

where $\omega = {}^{2\pi}/_{P}$ and is the orbital angular speed.

Figure 1.7 shows the geometry of the orbit once projected onto the sky plane in the lower panel, (b). As the orbit is inclined to the observer, the length $a\cos(\omega t)$ is foreshortened by an amount $\cos(i)$, and so the 'vertical' component of the separation vector $s(t)$ is $\cos(i)a\cos(\omega t)$. Using Pythagoras:

$$s(t) = a\sqrt{\sin^2 \omega t + \cos^2 i \cos^2 \omega t} \tag{1.20}$$

Mandel and Agol (2002) normalise the separation $s(t)$ to the stellar radius, so we define $z(t)$ as:

$$z = \frac{s(t)}{R_*} \tag{1.21}$$

and also the ratio of the planetary radius to stellar radius:

$$p = \frac{R_P}{R_*} \tag{1.22}$$

It is trivial to determine that for cases where the planet does not eclipse any of the star, the elipsed area, A, is zero, and for cases where all of the planet eclipses the star, $A = \pi R_p^2 = \pi p^2 R_*^2$. Using the newly defined z and p, these situations are described by the following inequalties:

$$A = \begin{cases} 0 & \text{if } 1 + p < z \\ \pi p^2 R_*^2 & \text{if } 1 - p \geqslant z \end{cases} \tag{1.23}$$

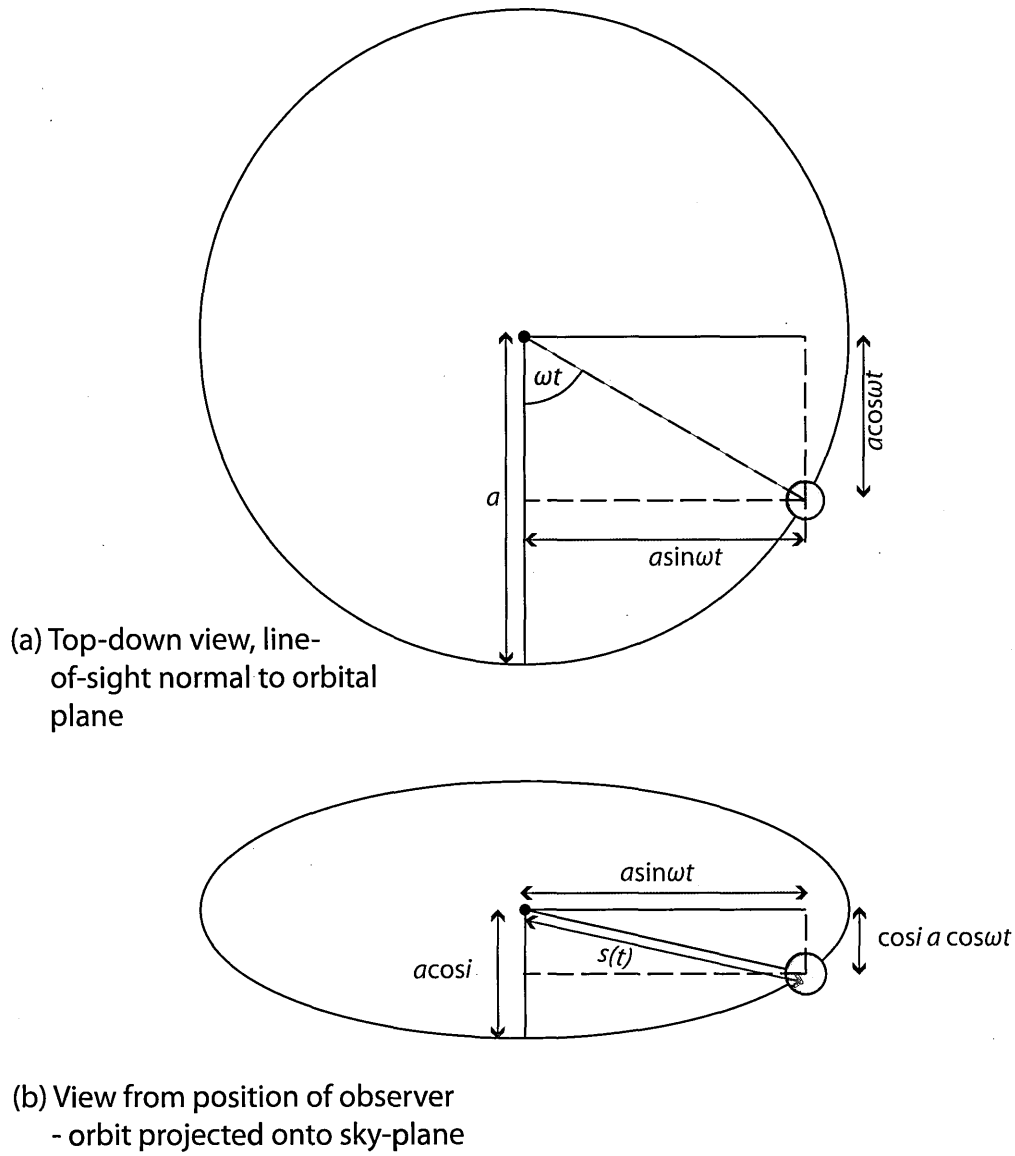

Figure 1.7: Superposition of the planetary orbit onto the sky-plane; the separation of centres is marked $s(t)$

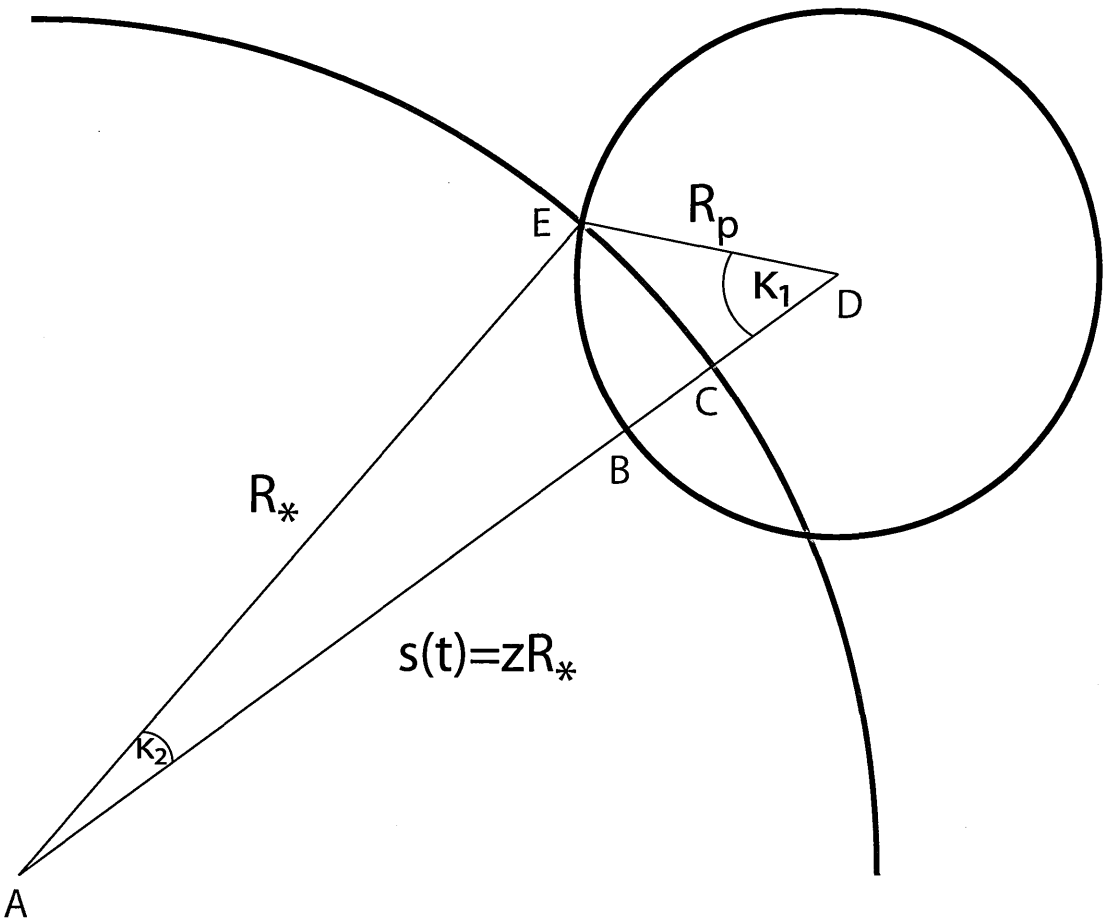

Figure 1.8: Geometry of the partial eclipse scenario, defined by the two radii, the separation of centres (parameterised by z), and two angles κ_1 and κ_2

The more general scenario is that of the eclipsed area during the transition from contact 1 to contact 2, or from contact 3 to contact 4, i.e. during ingress or egress. The geometry of this scenario is displayed in Figure 1.8, recreated from the similar figures in both Haswell (2010) and Mandel and Agol (2002).

Two key new features notable from Figure 1.8 are the two angles κ_1 & κ_2. These both vary as a function of time, and are formed between the 'separation of centres vector', and the lines joining the intersection of the planet's limb with that of the star's limb to both the stellar centre and planetary centre. Relevant vertices and intersections have been labelled in Fig. 1.8 , to make referencing sections of the diagram easier. For example, it is clear that the eclipsed area we wish to determine

is twice the area bounded by set of vertices BCE. The area, A, is also defined by:

$$A = 2 \times (BDE - (ADE - ACE)) \tag{1.24}$$

Where

$$BDE \equiv \frac{p^2 R_*^2 \kappa_1}{2} \tag{1.25}$$

$$ACE \equiv \frac{R_*^2 \kappa_2}{2} \tag{1.26}$$

and, via the sine rule

$$ADE \equiv \frac{R_* \times \ zR_* \sin \kappa_2}{2} = \frac{zR_*^2}{2} \sin \kappa_2 \tag{1.27}$$

We can use the cosine rule to write the two angles, κ_1 & κ_2 in terms of p & z:

$$\cos \kappa_1 = \frac{p^2 + z^2 - 1}{2zp} \tag{1.28}$$

$$\cos \kappa_2 = \frac{1 + z^2 - p^2}{2z} \tag{1.29}$$

however, we require $\sin \kappa_2$ and not $\cos \kappa_2$ before we can substitute our expressions for BDE, ACE, and ADE into equation 1.24. We can make use of the trig. identity $sin^2\kappa_2 + cos^2\kappa_2 = 1$ to accomplish this, and so finally:

$$\sin \kappa_2 = \frac{\sqrt{4z^2 - (1 + z^2 - p^2)^2}}{2z} \tag{1.30}$$

We can now substitute all of the elements into equation 1.24, finally arriving at:

$$A = R_*^2 \left(p^2 \kappa_1 + \kappa_2 - \frac{\sqrt{4z^2 - (1 + z^2 - p^2)^2}}{2} \right) \tag{1.31}$$

The full set of inequalities and their respective analytical equations for the eclipsed area are:

$$A = \begin{cases} 0 & \text{if } 1+p < z \\ R_*^2 \left(p^2\kappa_1 + \kappa_2 - \frac{\sqrt{4z^2-(1+z^2-p^2)^2}}{2} \right) & \text{if } 1-p < z \leqslant 1+p \\ \pi p^2 R_*^2 & \text{if } 1-p \geqslant z \end{cases} \tag{1.32}$$

With knowledge of the eclipsed area for all times in the orbit of a planetary body, we can now introduce a simple integral to deduce to the quantity of flux blocked by the planet throughout one orbit. Assuming a radially symmetric limb darkening law, we set the intensity of the stellar disc as a function of the radial distance r from the centre of the star, $I(r)$. The total change in the amount of flux observed is therefore:

$$\triangle F = \int^{A_{eclipsed}} I(r)\, dA \tag{1.33}$$

1.2.3.4 Limb darkening

As mentioned in the previous section, a star's luminosity is not constant across the disc, but is in fact darker towards the edges. This effect is known as limb darkening, and the physical reasons for it are briefly discussed in the previous subsection. Crucial to the understanding of the effect is the notion of optical depth. As the radial distance from the centre of the stellar disc increases, the 'optical depth' from which an emitted photon incident at the observer has to come increases for a given physical depth in the star. This is simply due to the nature of spherical geometry. At a given frequency, ν, the optical depth, τ_ν, for a photon emitted at position x travelling towards an observer at infinity is:

$$\tau_\nu = \int_{\mathrm{x}}^{\infty} \rho(s)\kappa_\nu \mathrm{d}s \tag{1.34}$$

where s is the position along the path between x and the observer (at infinity), $\rho(s)$ is the density as a function of position along the path, and κ_ν is the optical opacity. The probability that a photon emitted at x travels the full length of the path without being absorped or scattered is $\mathrm{e}^{-\tau_\nu}$, and so the emergent intensity at the end of the

path, I, is:

$$I = I_0 \mathrm{e}^{-\tau_\nu} \tag{1.35}$$

where I_0 is the initial, emitted intensity at the start of the path length. For photons emitted at a depth below the photosphere, h, the path length is approximately:

$$s \approx \frac{h}{\cos\gamma} \tag{1.36}$$

where γ is the angle between the direction the emergent photon must travel to reach the observer and the outwardly directed radius vector. This relation is approximate as it assumes that $h \ll R_*$, and therefore uses the plane-parallel approximation. Obviously, the geometry is actually spherical, so s is slightly overestimated, but the photosphere of the Sun is very slim compared to the total radius of the Sun, so the approximation is valid.

To incorporate the effects of limb darkening in the shape of the transit light curve, the above elements must be tied together to determine a functional form for $I(r)$ in eq. 1.33. We can make a crucial simplification at this point, known as the 'small planet approximation', which allows us to assume that, for a given position on the limb-darkened stellar disc, the occulted area is of uniform brightness. As we move across the disc, that brightness level is then modified according to the 'limb darkening law', which determines how stellar brightness changes as a function of the distance from the centre of the disc. A few types of limb darkening law are often adopted, including the simple linear law, the logarithmic law and the quadratic and cubic laws. These can be seen in Table 1.2. Also included is the four coefficient limb darkening law of Claret (2000), which is used in conjunction with the analytical approach of Mandel and Agol (2002) to generate the light curves used in the SuperWASP consortium MCMC model fitting code (more on this in Chapter 4). In this table $\mu = \cos\gamma$, where γ is as defined previously in this subsection. It is worth noting that this isn't the same as the separation of centres, r in the $I(r)$ function, as the impact parameter sets the minimum separation distance between the centre of the two discs. Unless $b = 0$, $\mu \neq 1$ at any point throughout the transit. In fact $\mu = \cos\gamma = (1 - r^2)^{1/2}$ (Mandel and Agol, 2002).The key implication here is that the choice of limb darkening law and the coefficients chosen for that law will affect the maximum observed transit depth. This is because the limb darkening laws are only equivalent to one another when $\mu = 1$, and therefore all terms involving μ vanish to zero.

Limb Darkening Law (Claret, 2000)	$\frac{I(\mu)}{I(1)}$
Linear	$1 - u_{lin}(1-\mu)$
Logarithmic	$1 - u_{log}(1-\mu) - v_{log}\ln(\mu)$
Quadratic	$1 - u_q(1-\mu) - v_q(1-\mu)^2$
Cubic	$1 - u_c(1-\mu) - v_c(1-\mu)^3$
Non-linear four co-efficient	$1 - u_{Cl}(1-\mu^{1/2}) - v_{Cl}(1-\mu) - w_{Cl}(1-\mu^{3/2}) - x_{Cl}(1-\mu^2)$

Table 1.2: Five commonly used limb darkening laws. Here $\mu = \cos\gamma$, where γ is as defined previously in Eq. 1.36.

To determine the coefficients used in the limb darkening law of choice, knowledge of the star's effective temperature and the pass band in which the observations are made are required. As previously mentioned, optical depth is a function of frequency, and thus so is limb darkening. Due to the fact that higher energy photons scatter more than lower energy photos in the Rayleigh scattering process, less limb darkening is observed when observing a transit in a redder pass band. The resultant light curves are visibly flatter than similar observations in bluer pass bands between contacts two and three.

1.2.3.5 Determining system parameters from the transit light curve

The transit light curve alone cannot determine the scale of the observed system. From the depth of the transit, the duration of the transit, the impact parameter, and the period of the system, the ratio (a/R_*) may be determined (Seager and Mallén-Ornelas, 2003). Two independent checks on the ratio (R_p/R_*) exist, via the transit depth, and by the duration of ingress/egress. Therefore the ratios $R_p : R_* : a$ are known, but not the scale of the system. The scale of the system can be set however by determining a through Kepler's third law, which requires external knowledge of the total mass of the system. In practice, as $m_p \ll m_*$, the total mass of the system $m_{tot} \approx m_*$. In order to have detected a transiting planet, knowledge of the system's period is implicit, and so together with the system's mass and period, the orbital separation may be derived. Alternatively, the stellar density may be determined before the mass of the star is known:

$$\rho_* = \frac{m_*}{R_*^3} = \frac{4\pi^2}{GP^2}\left(\frac{a}{R_*}\right)^3 \tag{1.37}$$

as the ratio (a/R_*) is determinable from the shape of the light curve. The host star mass is still required to set the scale of the system, however. Typically, this can be done either photometrically, by linking the colour index of the star to an effective temperature and therefore radius, then estimating the mass from the radius from a main-sequence mass-radius power law approximation; or it can be done spectroscopically through isochrone fitting. A new approach has also recently been adopted for incorporation in the SuperWASP model fitting software, that uses prior knowledge of the star's effective temperature, metallicity, and density (measurable from the light curve) to determine its mass. More on this is included in Chapter 4.

Chapter 2

PIRATE Hardware & Software

2.1 The PIRATE hardware and its evolution

2.1.1 Primary hardware

PIRATE was first installed on site at the OAM in March 2008. Initially it was housed within a small motorised run-off shed, pictured in Fig. 2.1. The run-off shed was an imperfect solution and PIRATE was later moved to the top of the main observatory building at the OAM in August 2009; its current location, where it sits within a Baader Planetarium 3.5m All-Sky 'clamshell' dome, pictured in Fig. 2.2. Throughout its relatively short life, PIRATE has received numerous hardware upgrades. It is now (as of November 2011) quite different from its original installation phase, and so we must define separate phases of operation, so that the reader is aware of the different hardware (and consequently different hardware characteristics) used to take certain data.

The initial phase of operation involved the use of a Celestron 14" C14 Schmidt-Cassegrain reflector in combination with the Paramount ME German Equatorial Mount (GEM). An Optec TCF-S microfocuser was paired with the SBIGSTL-1001E for imaging. In August 2009 the move to the new Baader dome was completed, but the hardware was otherwise kept the same. August 2010 saw another significant upgrade as the Celestron C14 was replaced with a larger 17" Optical Tube Assembly (OTA), the PlaneWave CDK17, which brought with it its companion 'Hedrick' focusing unit; and so the C14 and TCF-S were removed from the facility at this point, with the STL-1001E still in place. September 2011 saw the last significant upgrade,

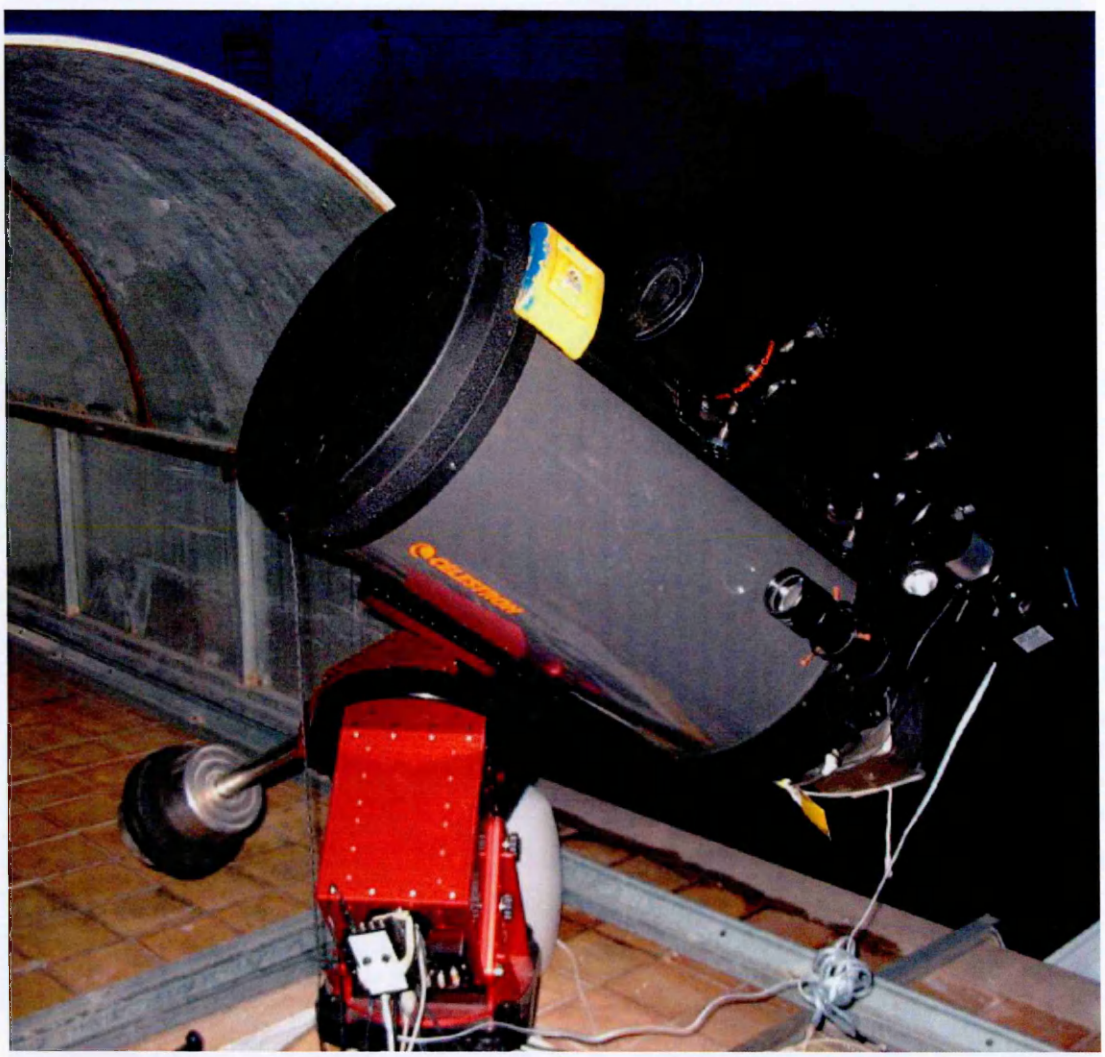

Figure 2.1: The original PIRATE installation, picture taken March 2008. The building on which PIRATE now resides can be seen in the background.

Figure 2.2: PIRATE's new dome enclosure, the Baader Planetarium All-Sky 3.5m dome.

involving the removal of the STL-1001E in favour of the SBIG STX-16803 imager. We define the initial phase of operation 'PIRATE Mark 1', and designate the current (as of November 2011) set-up 'PIRATE Mark 2'. The intermediary phase (where the STL-1001E was coupled with the CDK17) is referred to as 'PIRATE Mark 1.5'(see Fig. 2.3). This phase receives a non-integer moniker as it was always intended to be a transitionary phase. Table 2.1 details these phases of operation, where the two-part Mark 1.5 phase (3.5m dome installation followed by the introduction of the CDK17) is kept as one column, despite these events occurring a year apart. The table also details the vital statistics of the imaging system as it has evolved, with the most notable improvements being the increase in the Field of View (FOV) from $(21.6' \times 21.6')$ to $(43.2' \times 43.2')$ coupled with an decrease in the plate scale from $1.26''\text{pixel}^{-1}$ to $0.634''\text{pixel}^{-1}$ (when operating the STX-16803 at binning 1).

2.1.2 Ancillary hardware

In addition to the primary hardware, there is a selection of support hardware on-site used mostly for monitoring purposes. Weather is monitored by a weather station and Boltwood cloud sensor (to the side of the dome) and by both internal and external dome-mounted weather systems (rain, internal humidity, and internal temperature). The dome has its own firmware, which interfaces with the proprietary dome driver, as well as with the Boltwood cloud sensor, allowing for shutdown conditions to be communicated directly to the dome, bypassing the control PCs.

There are two control PCs used for the operation of the facility; the primary control PC is responsible for the operation of the main telescope components (mount, camera, focuser). The second control PC is used solely for all of the weather software required to run the weather systems.

In addition to the weather systems and control PCs, four D-Link Webcams (three internal, one external) provide live video and audio feeds, as well as IR lights for night viewing. A dehumidifier is present in the dome and set to turn on only when the dome is closed for energy conservation purposes.

2.2 PIRATE software

The PIRATE hardware is controlled by a collection of Windows-based software (MaxIm DL for imaging and camera control, TheSky6 for mount control, and FocusMax for

	PIRATE Mark 1	PIRATE Mark 1.5	PIRATE Mark 2
OTA	*Celestron C14*	*PlaneWave CDK17*	*PlaneWave CDK17*
Focal Length (m)	*3.91*	*2.93*	*2.93*
Aperture Size (m)	*0.35*	*0.425*	*0.425*
F Number	*f/11*	*f/6.8*	*f/6.8*
Main Camera	*SBIG STL-1001E*	*SBIG STL-1001E*	*SBIG-STX16803*
Pixel Size (μ)	*24*	24	*9*
Array Size	1024 × 1024	1024 × 1024	4096 × 4096
Plate Scale ($''/_{pixel}$)	*1.26*	*1.69*	*0.634*
Field of View	21.6′ × 21.6′	28.8′ × 28.8′	43.2′ × 43.2′
Filters	*B, V, R, I, H_α, SII, OII*	*B, V, R, I, H_α, SII, OII*	*R, G, B, Clear, H_{alpha}*
Mount	*Paramount ME GEM*	*Paramount ME GEM*	*Paramount ME GEM*
Focuser	*Optec TCF-S*	*PlaneWave Hedrick Focuser*	*PlaneWave Hedrick Focuser*
Guider Scope	*Celestron 80mm refractor*	*Celestron 80mm refractor*	*Celestron 80mm refractor*
Guider Focal Length (m)	0.6	0.6	0.6
Guider Camera (G.C.)	*SBIG ST-402ME*	*SBIG ST-402ME*	*SBIG ST-402ME*
G.C. Array Size	765 × 510	765 × 510	765 × 510
G.C. Pixel Size (μ)	*9*	*9*	*9*
G.C Plate Scale ($''/_{pixel}$)	3.09	3.09	3.09
Dome	*Custom-built run-off shed*	*Baader Planetarium All-Sky 3.5m*	*Baader Planetarium All-Sky 3.5m*

Table 2.1: Phases of Operation for the PIRATE facility, detailing the evolution of the PIRATE facility. For simplicity, the PIRATE Mark 1.5 phase is convolved into a single column, though it actually involves two separate signification upgrades (to the dome and then to the OTA) that took place a year apart.

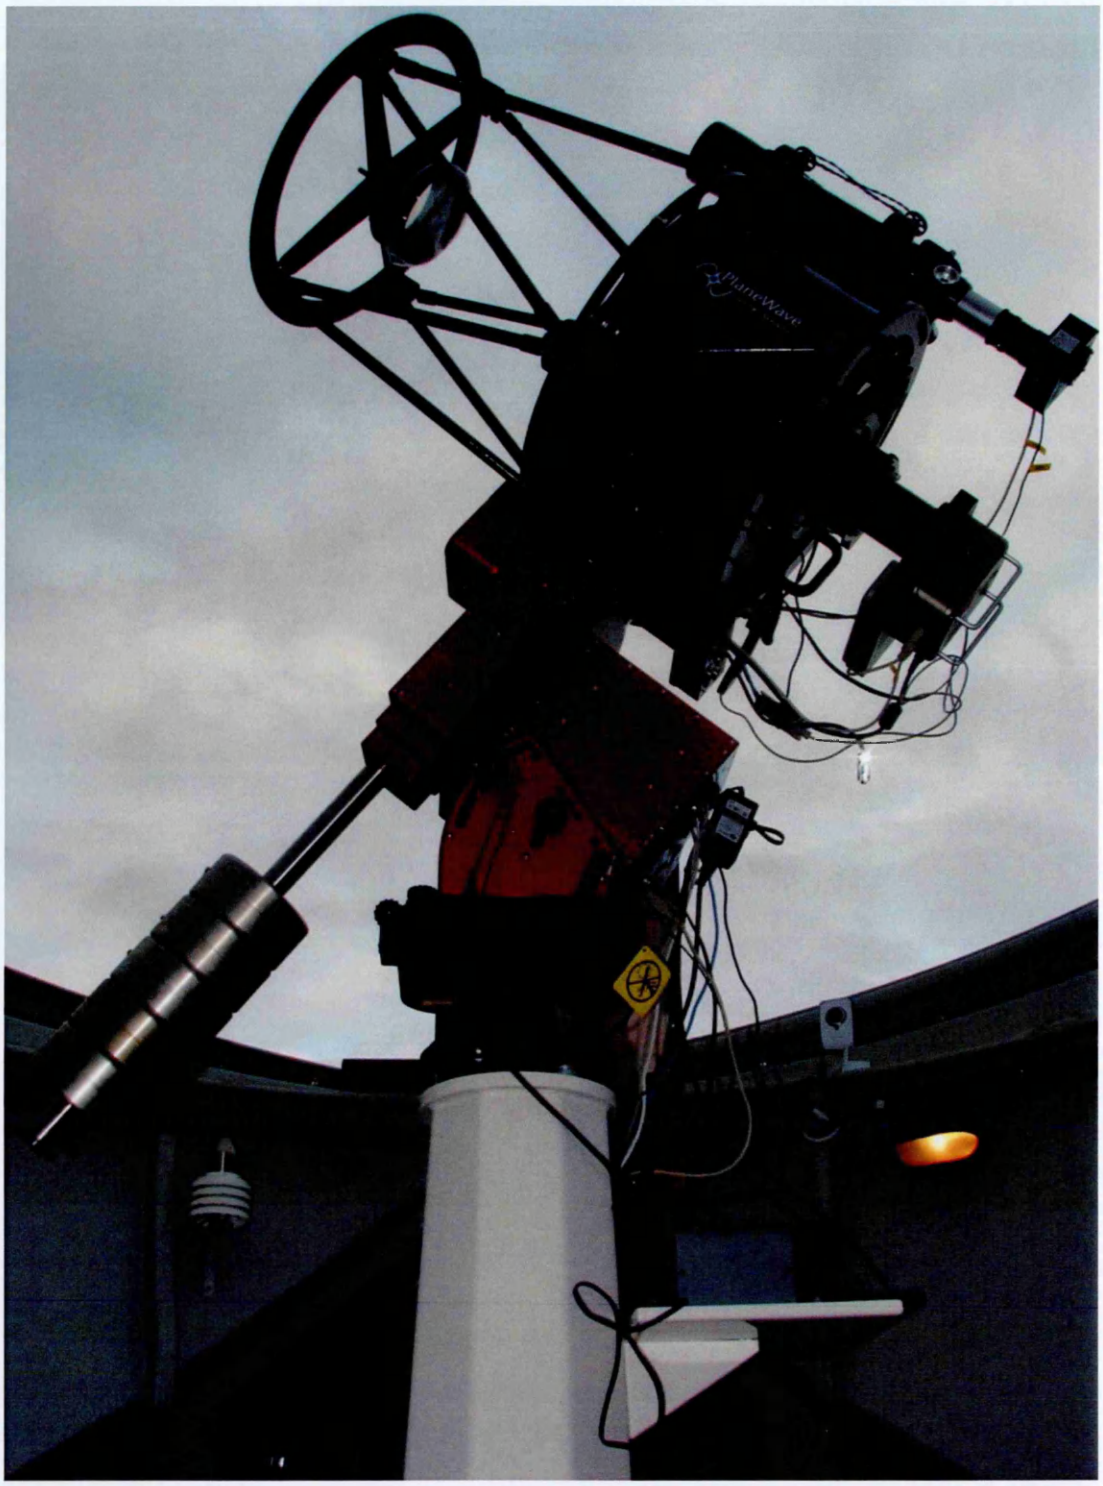

Figure 2.3: PIRATE Mark 1.5 in the 3.5m All-Sky dome

the control of the microfocuser), which are each controlled by ACP, an observatory control suite that automates the command of each of the component software packages through the ASCOM-standard driver layer. These programs are run on a 2.13 GHz CPU Windows XP PC with 2 Gbyte of RAM. The site receives its Internet connection via WIMAX; connection speeds are good, with typical upload speeds from the OAM to the UK of ~250 kbyte s^{-1} and operating latencies of around 60-100 ms. We use Dimension 4 from Thinking Man Software to synchronize the control PC's clock to UTC via a Network Time Protocol server every 10 s.

2.2.1 MaxIm DL

MaxIm DL is commercially available Windows software that has some limited aperture photometry capability and image procressing and calibration procedures. However, the software is only used for these purposes by undergraduate students processing PIRATE data. Within the PIRATE facility MaxIm DL is used for camera control. It is used to connect to and cool the cameras, and is also used for shutter control and chip readout. It also exists within Windows as a scriptable Windows component, meaning that it can be externally controlled by third party software, most notably ACP.

2.2.2 TheSky6

TheSky software is used for mount control and is responsible for all of the pointing and movement functionality of the telescope facility. The current version of TheSky is TheSkyX, but we have chosen to stay with TheSky6 for continuity reasons. TheSky6 currently performs two very important performance related functions alongside simply communicating the requisite number of turns to the Right Ascension and declination gears: modelling the pointing inaccuracy of the mount, and removing periodic error.

2.2.2.1 The pointing model

TheSky6 makes use of the TPoint pointing model software (as used in many major telescope facilities e.g. Keck, VLT) to ensure that when a command is sent to the Paramount ME to move to a particularly Right Ascension (α) & Declination (δ), the mount successfully carries out this request and the OTA ends up pointing at the desired target. The software records the pointing inaccuracy across the whole of the

sky, in both pointing offset size and direction, by asking the user to move to bright catalogue stars (i.e. of known α & δ in J2000 epoch) and centre them in the FOV of the telescope, recording the telescope position as understood by the mount (i.e. using the initial synchronisation point and the number of turns of the worm gear in α & δ). The discrepancy between catalogue position and recorded position of mount is stored and modelled by the TPoint software. Consequently, it is important to use stars across a wide range of altitudes and azimuthal positions. The map of the pointing error is then modelled using a set of geometrical terms that describe various misalignments, flexures, and non-perpendicularities in the OTA-Mount system, as a function of h, δ, and ϕ, the hour angle, declination, and site latitude respectively. They break down into five categories: equatorial, alt-azimuth, special (i.e. unique to a particular telescope or type of telescope), polynomial and harmonic. PIRATE's model makes use of only equatorial, special, and harmonic terms. The model terms currently in use (as of April 2012) by the TPoint software in the PIRATE facility are detailed below in Table 2.2.

2.2.2.2 Periodic Error Correction (PEC)

TheSky6 also has the ability to learn any periodic errors in the sidereal tracking rate of the mount. It does this in conjunction with the guider scope, taking the necessary guider impulse corrections (required to nudge the guide star back to the centre of the guider image) and using these corrections to search for any periodic perturbations in the tracking rate. The source of such perturbations to a perfectly constant sidereal tracking rate are typically irregularities (i.e. deviations from radial symmetry) in the worm gear responsible for the drive in the right ascension direction. We used TheSky6's ability to record the periodic error, and fitted a fifth-order polynomial with a peak-to-peak amplitude of $4.8''$.

2.2.3 FocusMax

FocusMax (Weber, 2001) is freeware software used to control the micro-focusing unit in PIRATE that sits between the camera and the OTA. It can also provide a quick and efficient auto-focus, meaning that seeing-limited images should be achievable on most nights. To do this, FocusMax must build a model of how the focus shifts for a given increment in the micro-focuser. The Hedrick focuser employed in PIRATE Mark 1.5 and later has $3.3 \times 10^5 \mu$ of travel, and can make step sizes as small as 1μ.

Model term (coefficient)	Correction (arc seconds)	Explanation	Δh	$\Delta\delta$
IH	88.31	Index error in Hour Angle (zero-point offset to the hour angle)	IH	
ID	216.37	Index error in Declination (zero-point offset to the declination)		ID
NP	-181.92	Offset due to non-perpendicularity between the declination & polar axis	NP $\tan\delta$	
CH	-16.59	Collimation error (i.e. secondary mirror to primary mirror misalignment)	CH $\sec\delta$	
ME	113.77	Polar axis vertical misalignment	ME $\sin h \tan\delta$	ME $\cos h$
MA	-80.92	Polar axis horizontal misalignment	$-$MA $\cos h$ $\tan\delta$	MA $\sin h$
TF	1.97	Tube flexure	TF $\cos\phi$ $\sin h$ $\sec\delta$	TF$(\cos\phi\cos h\sin\delta - \sin\phi\cos\delta)$
FO	-75.23	Fork flexure		FO$\cos h$
HHCH1	-363.80	Harmonic	HHCH1 $\cos h$	
HDCH1	-75.24	"		HDCH1 $\cos h$
HDCD1	-74.56	"		HDCD1 $\cos\delta$
HDSD2	5.93	"		HDSD2 $\sin 2\delta$
HDCD5	7.72	"		HDCD5 $\cos 5\delta$
HDSD5	-12.06	"		HDSD5 $\sin 5\delta$
HDCD2	-26.03	"		HDCD2 $\cos 2\delta$
HDSH4	17.28	"		HDSH4 $\sin 4h$
HHSD2	-10.52	"	HHSD2 $\sin 2\delta$	
HHSH6	-8.18	"	HHSH6 $\sin 6h$	
HHSH5	10.31	"	HHSH5 $\sin 5h$	

Table 2.2: List of co-efficients for each of the geometric terms employed in PIRATE's current pointing model (as of November 2011)

To find best focus, the software's model must know two things. Firstly, it must know which side of the focus it is on, so that it can determine in which direction the camera must be moved. Secondly, it must know the increment size required to move to best focus. These two vital pieces of information are encapsulated within the V-Curve modelling system.

2.2.3.1 V-curves

The focus model is built from many 'V-curves'. These are plots achieved by stepping the microfocuser through a range of values whilst measuring a metric to determine the quality of focus, producing a 'V-shaped' well when the metric is plotted against focuser position as the microfocuser steps into best focus and then out again. The focus quality metric used by the software is the Half-Flux Diameter (HFD), defined as the diameter of the circle that encloses half the flux from the stellar image when centred on the star in question. This metric is used in preference to Full-Width Half-Maximum (FWHM) as it is less sensitive to variations in seeing conditions, and can also make meaningful measurements of unfocused ring-shaped images.

Once a set of V-Curves has been obtained, each one is characterised by the gradient of the left and right-hand slopes, and the PID, the 'Position-Intercept Difference'. The PID is the difference (in microns) between the intercept of the 'left slope' with the x-axis (HFD=0) and the equivalent x-axis intercept for the right slope. The system's focus is described precisely by these three parameters, as temperature fluctuations only have the effect of moving the position of best focus; the left and right-slope gradient and the PID remain unchanged. This allows the software to find best focus in a range of temperature conditions, and means that, once determined, the focus model does not require updating and will function through all seasons. A schematic representation of a synthetic V-Curve is depicted in Fig. 2.4.

2.2.3.2 The auto-focus procedure

Once armed with the model, FocusMax must then employ it to find the best focus with speed and reliability. FocusMax functions well within the PIRATE system, and so good focus is a feature oft-taken for granted during the operation of PIRATE. The auto-focus procedure makes use of a 'Near Focus HFD' concept, a focus position that produces a user-defined HFD value that is near the best focus HFD value, but sufficiently far away as to reside on the linear section of the V-Curve. Once at the Near Focus position, the best focus position is calculated from the V-Curve,

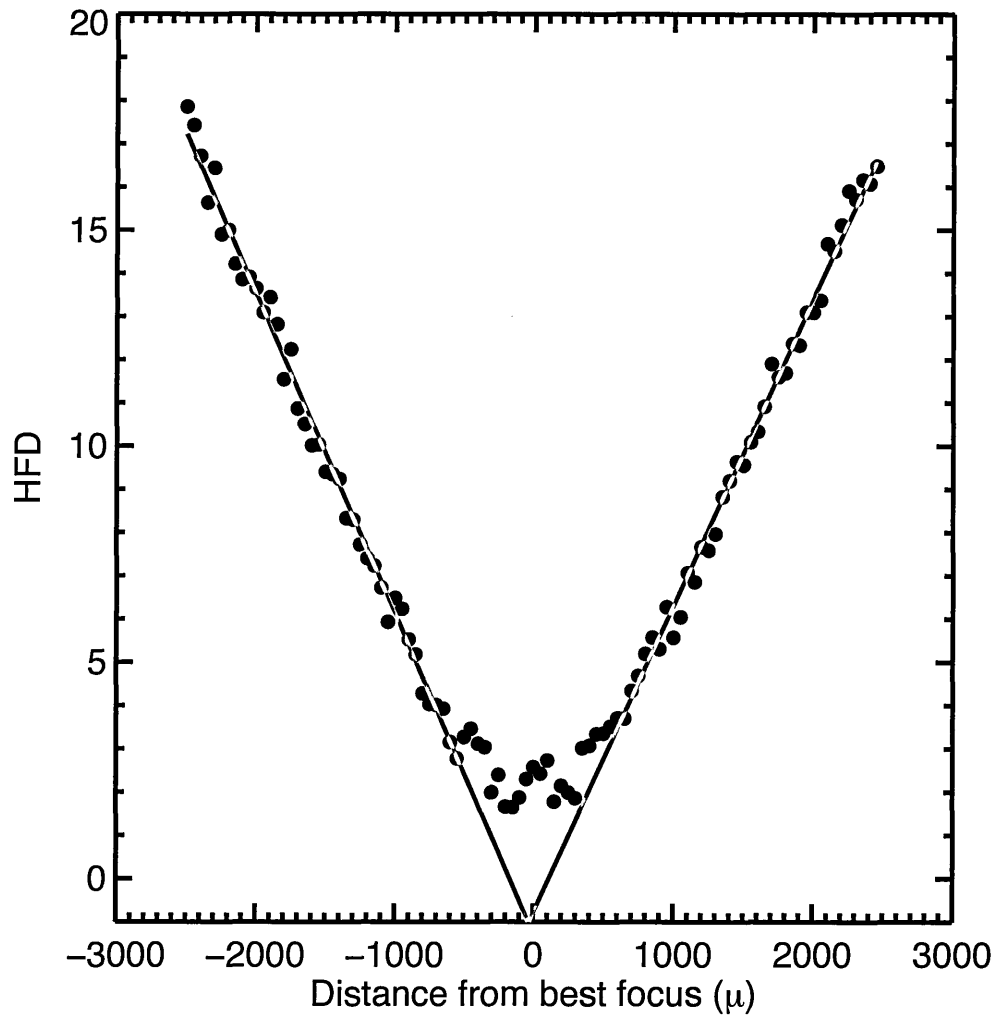

Figure 2.4: Schematic representation of a focus V-curve. This V-curve is synthetic and is not derived from real data. It depicts both the left slope and right slope fit converging with a near zero PID.

and the focuser moves straight to the best focus position. This avoids an iterative process of moving around best focus (at the bottom of the V-Curve 'well'), where determining the direction in which to move is very tricky and subject to fluctuations in the atmospheric seeing conditions. By moving the focuser in one direction towards best focus, any backlash in the focuser motors can also be avoided. As a consequence, the user must specify the direction they wish the focuser to travel in order to reach the best focus position (i.e. rack the focuser in or out). The software will then move to either the left or right slope, which ever is appropriate for the designated direction of travel. The full step-by-step process undertaken by the autofocus procedure is as follows:

1. Take a 3 second pointing exposure at the current location and plate-solve it to determine the current location

2. Search in a $2° \times 2°$ box around the current position for stars of magnitude 4-7, and slew to the nearest. If no suitable stars are found, the box size is iteratively increased until a star is found.

3. Take a full frame exposure and find the target star, setting a subframe around it so that only the target source exists within the subframe. The background level is determined and subtracted, and the HFD determined. For each frame taken, the HFD is calculated and the window frame-size re-evaluated based on the current HFD.

4. In order to verify whether the focuser is on the correct side of the focus, the focuser is racked in the opposite direction to the preferred 'travel-to-focus' direction, and an increase in HFD is verified, otherwise the focuser is moved to the other side of focus.

5. Once on the correct side of focus, the focuser is moved to an HFD of 20 pixels, the 'Start Focus position' (another user-set parameter).

6. A jump is then made from the Start Focus position to the Near Focus position.

7. Five exposures are taken at the calculated Near Focus position, and the average HFD value is determined. In PIRATE, the Near Focus HFD is set to 10, therefore we expect these five HFD measurements at Near Focus to also be around 10 pixels.

8. The best focus position is calculated from the average measured HFD at near focus, and the focuser is moved to best focus.

9. A final exposure is taken and the HFD at best focus recorded in the logs.

2.2.4 ACP

ACP is effectively the brain of the PIRATE facility. It provides the logic and algorithms for automating the control of each of the system's sub-components. Whilst it is possible for a user to instruct TheSky6 to 'coarse' point the telescope, subsequently refine the pointing so the target object is centred, switch to FocusMax to focus the system (in the desired optical passband), then use MaxIm DL to expose the CCD for the desired duration, whilst simultaneously auto-guiding (performed in MaxIm DL); this is a lot of individual work, and particularly time consuming. ACP automates the whole procedure, and even includes a web server for simple target input by web form. All of this is made possible via the ASCOM initiative, which was introduced primarily in the field of amateur astronomy in order to encourage telescope equipment manufacturers to conform to a universal (language-independent) driver standard. ASCOM drivers for a given instrument appear as objects within a Windows environment with a fixed and well-defined set of methods. This means that, no matter which manufacturer of a telescope mount, if their driver is ASCOM compliant it will appear as a Windows COM object and will (as an example) be guaranteed to have the method Telescope.SlewToTarget(), which slews the telescope to given α & δ coordinates. This framework enables the automation of the facility by having each faculty be a COM object that accepts input values, performs operations, and reports output values where necessary. What follows is a description of the operations performed by ACP that underpin the operation of the PIRATE facility, and a description of how the user instructs ACP to perform observations.

2.2.4.1 Web server

ACP runs a web server that allows external connections to the control PC to operate the telescope via a web form, pictured in Fig. 2.5. The web form is particularly user friendly, and deliberately attempts to hide the complex behind-the-scenes operation from the user, utilising a simple 'request and receive' philosophy. For example, if the user required a 60s B filter exposure of M31, accomplishing this would be as simple as selecting 'single image', inputting 'M31' in the target field, asking ACP to resolve the

target name (where it puts the J2000 α & δ coordinates in the requisite fields after searching the Deep Sky Catalog), inputting '60s' and 'B' as the exposure time and filter respectively, ticking 'autofocus' to ensure the image is focused beforehand, and finally hitting 'acquire'. The user would then, assuming the sky conditions were not so adverse as to affect operations, receive a centred, in-focus, and fully auto-guided image in their filespace, accessible for download via the web form. It is this functionality that makes ACP an obvious choice for the educational tasks that PIRATE performs when supporting Open University undergraduate astronomy modules. In the limited time available to an undergraduate student wishing to use the telescope for time-series photometry of an eclipsing binary system, excessive exposure to the individual machinations of each telescope sub-component is undesirable. Instead, we wish to facilitate the delivery of real, high quality data to the student so that they may produce light curves of astrophysically interesting objects, as their learning objectives are typically focused on the physics of the target; not the instrument used to take them.

The web server, when sending a set of commands to the telescope facility, is essentially writing an 'observation plan', a text file containing a pseudo mark-up language specific to ACP that allows for scripting of the telescope, that is subsequently then sent to the main program for compilation and execution. Whilst the web interface is ideal for student use, being able to construct observing plans manually for submission to ACP allows for greater flexibility and control for research users.

2.2.4.2 Observing plans

Observing plans contain a list of instructions, or *directives* that tell the system when and how to observe. Their flexibility allow for a full night's observing to be planned out and scheduled in advance, allowing the user to simply submit the observing plan and collect the data the following day. The details of the specific scheduling and plans used in the SuperWASP candidate follow-up program can be found in Chapter 4. An example observing plan is shown in Fig. 2.1. Officially, all lines prefaced by a hash are 'directives' and all lines without hashes are 'targets'. However, '#CHILL', '#DOMEOPEN', '#DOMECLOSE', '#BIAS', and '#DARK' all behave like targets. Typically (bar the exceptions listed and a few others) directives describe *how* an action should be performed, and targets execute the action. In Fig. 2.1, PIRATE is instructed by the user's plan to do the following:

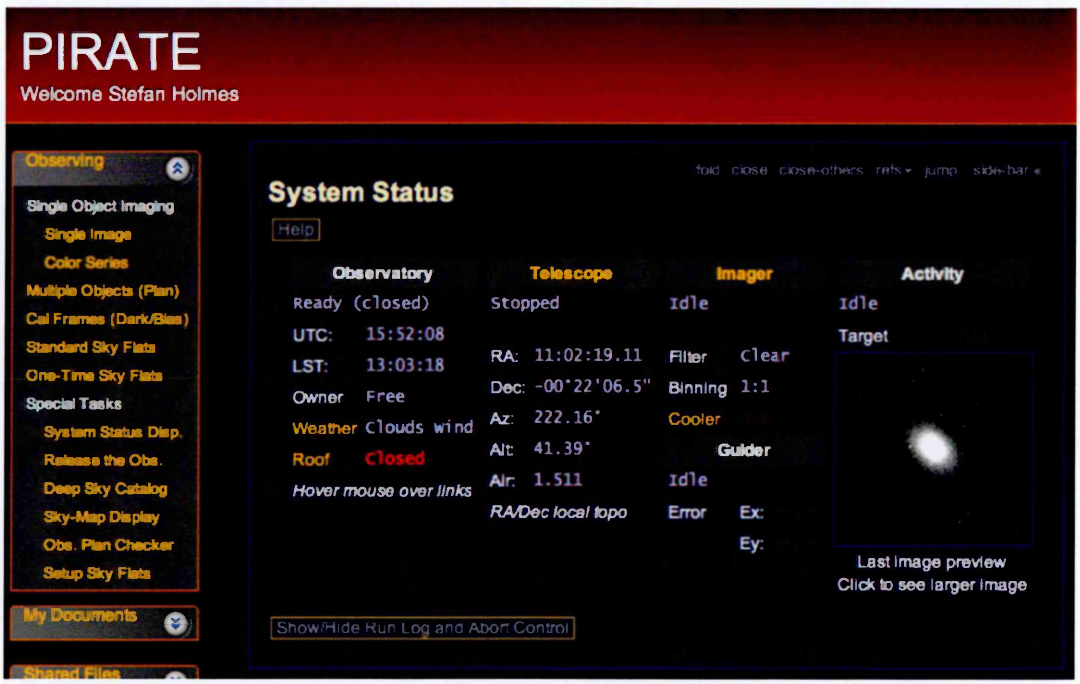

Figure 2.5: The ACP web interface, which runs in a web browser.

1. Wait until the time is 16:45 UTC (ACP will wait up to 24 hours for this criterion to be true)
2. Chill the main imaging camera down to a target temperature of 15 °C
3. Repeat this process until final '#CHILL' directive target temperature is reached (-20 °C)
4. Open the dome
5. Take dusk sky flats (detailed later in this subsection)
6. Close the dome
7. Take 25 bias frames with binning = 1, and 25 bias frames with binning = 2.
8. Take 25 dark frames of 60s each with binning = 1, and the same again with binning = 2.
9. Open the dome
10. Perform an autofocus.

11. Slew to a North-East field of M31, centre the image and take 9 images at binning = 1 at 60s each in the Clear filter, all dithered.

12. Stop this task (imaging M31) at 21:10 UTC

13. Compile and execute another observing plan by the name of '20111009_SWASP-cand_to_SWASPcand.txt'

Whilst this plan, which is a shortened sample of a plan submitted on 9^{th} October 2011 as part of the SuperWASP follow-up program, does not cover all of ACP's plan directives and capabilities, it amply demonstrates the capability for scripting events and running the telescope facility in an autonomous capacity. The following subsection details the mechanisms that allow the scripted image acquisition process to occur.

2.2.4.3 Automatic sky flat acquisition

Instead of taking dome flats, where a screen within the dome or the dome itself is illuminated and the telescope is pointed at a evenly illuminated section of the dome or screen, ACP allows the user to automate the process of taking sky flats, where the telescope is pointed at a 'null point' of the sky during dusk or dawn (15° East of the zenith at dusk, 15° West of the zenith at dawn - Chromey and Hasselbacher, 1996). The null point is where the intensity gradient of the sky is expected to be minimised, to promote even illumination of the chip. If the user was responsible for manually operating the telescope to achieve optimally exposed flat fields in all filters, the user would contend with offsets in the various bandpass widths (i.e. different flux levels due to variation in filter sensitivity), along with a non-linearly decreasing or increasing sky brightness. These two effects combined make determining the correct exposure time for a a given filter at a given time of day particularly tricky. Furthermore, to achieve a better SNR in the master flat, a set of flats taken in the same filter are usually median-combined to make a master flat. The operator may find that even in the time it takes to make a set of ten exposures in a particular filter that the sky brightness has changed significantly, meaning one single exposure time for a set of flats in a given filter may not be ideal. ACP fully automates the process of taking sky flats, leaving the user with only the need to specify how many flats in each filter are required, and to submit the request to take sky flats in good time before either dusk or dawn. The user must specify a target ADU in the ACP settings, which is chosen as 33000 ADU; in the middle of the digital dynamic range. As soon as the

Algorithm 2.1 Shortened version of an observing plan used on the 9th October 2011

```
#WAITUNTIL 1, 16:45
#CHILL 15

#WAITUNTIL 1, 16:50
#CHILL 5

#WAITUNTIL 1, 16:55
#CHILL 0

#WAITUNTIL 1, 17:00
#CHILL -5

#WAITUNTIL 1, 17:05
#CHILL -10

#WAITUNTIL 1, 17:10
#CHILL -15

#WAITUNTIL 1,17:15
#CHILL -20

#DOMEOPEN
#DUSKFLATS
#DOMECLOSE

#INTERVAL 0,0
#COUNT 25,25
#BINNING 1,2
#BIAS

#INTERVAL 60,60
#COUNT 10,10
#BINNING 1,2
#DARK

#DOMEOPEN

#autofocus
#dither 4
#count 9
#interval 60
#binning 1
#filter Clear
M31_NE 00:44:34.905 +41:47:09.82

#QUITAT 21:10

#CHAIN 20111009_SWASPcand_to_SWASPcand.txt
```

automatic sky flat request is made to ACP, it will sample the sky brightness in the middle of the image with the telescope pointing at the null point, taking a continuous sequence of 0.1s exposures at one minute intervals. Typically, a period of waiting is required for the sun to reach the required elevation angle, and for the sky brightness to consequently fall to the target value (assuming dusk sky flats). Once the sky brightness reaches the target value (33k ADU), the subsequent exposure times within the filter set are automatically scaled upwards to account for the ever-decreasing sky brightness (this process is obviously reversed for dawn flats). The process is repeated for each filter set.

The user must specify the filters and the number of exposures to take in a pre-prepared 'flat plan' to submit to ACP. PIRATE mostly makes use of dusk flats, dawn flats are very rarely taken, due to the demands on the user's time and the increased risk of exposure to the sun in case of a mechanical fault. The order of the filters must be optimised for the decreasing light levels, so the most sensitive passbands (i.e. the Clear filter) must be at the end of the sequence, and vice versa. If an unsuitable order is selected, the automatic exposure time scaling will result in very long exposure times, which increases the overall time spent taking sky flats. Alternatively, dusk time is wasted whilst waiting for the sky brightness to drop for flat fields in more sensitive passbands. Given that there is a small window of opportunity where the sky brightness is 'appropriate', it is desirable to optimise the order of flats in order to reduce the overall time spent taking flats. Furthermore, along with the minimum exposure time, a maximum exposure time of 10s is set for PIRATE; if the maximum exposure time is hit for a given filter, that filter set is skipped and no flat fields will be taken in that filter set. The optimum order of filter sets for dusk flats in clear-sky conditions is I, R, B, V, Clear. If narrowband flats are required, these are taken before the I band flats. The order of the filters is largely in order of decreasing wavelength due to the overwhelming blue colour of the dusk sky, barring the switching of B and V due to the increased bandwidth of the V passband with respect to the B passband. In cloudy conditions, the order is set to R, I, B, V, Clear, due to the increased sky brightness in I because of the emissivity of water molecules in that passband. For PIRATE, 10 flat field images are usually taken in each filter, in order to produce a good, strong SNR master flat field (SNR $>$ 1000) when median-combined whilst keeping the total duration of flat field exposures within the window of opportunity provided by dusk.

2.2.4.4 Initial pointing, pointing updates

ACP works in conjunction with the aforementioned TPoint model (applied by TheSky6 when the telescope is instructed to move to a particular part of the sky) to produce high quality pointing via initial pointing procedures and pointing updates. These work like so:

- Once the autofocus procedure has completed, ACP sends a command to TheSky6 to slew to particular coordinates. The input coordinates are adjusted accordingly by the TPoint model in order to correct for systematic pointing errors in the mount
- A pointing exposure is started, lasting 4s. This is plate-solved by the Pinpoint software in conjunction with the Guide Star Catalogue (Lasker et al., 1988) to determine the true centre of the image, and the pointing error (both angular offset and direction) determined.
- A corrective slew is applied to centre the target star in the middle of the frame, typically to within a few arc seconds, completing the initial pointing procedure.
- The maximum pointing error for PIRATE is set to $3''$. After each full exposure is taken, the image is plate solved and pointing error assessed. If deemed to be $> 3''$ then a pointing update is applied following the procedures already described. This helps constrain any drift in the overall pointing of the system throughout an observation run, the effect of which can be seen in Fig. 2.6, taken from Holmes et al. (2011)

2.2.4.5 Autoguiding

Throughout an exposure, ACP engages and disengages Maxim DL's autoguiding feature at the appropriate time. At its most basic level, the principle of autoguiding whilst exposing the main imager and having the mount track at the sidereal rate (and counteract periodic error with the PEC) is very simple. The guider camera is continuously exposed at the specified interval, the image downloaded, and the position of the elected guide star determined. Any deviation from the previous position of the guide star is evaluated, and an appropriate corrective pulse is sent via the cable from the guider camera to the paramount to 'jolt' it back into position. This cycle is repeated throughout the duration of an exposure with the main camera. However, if the user was just using Maxim DL to guide, the user would need to personally select

the guide star and choose a desired exposure time. This is obviously inconvenient if attempting to make automated observations, so ACP takes control of the procedure and adds extra logic and automation in order to achieve consistent, well-performing guiding.

ACP performs the following tasks to achieve 'autoguiding':

- Before each main camera image is taken, a 5s guider image is taken with the ST-402 camera mounted on the guide scope, and is checked for potential guide stars. ACP is responsible for selecting the guide star in the guider image. The brightest star that does not have a companion within 20 pixels, is not within 10 pixels of the edge of the guide chip, and is not saturated is selected.

- The SNR of the selected guide star is measured, and the guiding interval is adjusted to be as short as possible whilst still achieving the minimum set SNR (for PIRATE this is set to an SNR of 3). A lower limit of 0.5s to the guider interval is also imposed, so as not to overwhelm the mount with guide pulses. Typically, dependent on the stars available in the guider scope's field of view, the guider interval is anywhere from 0.5s $\rightarrow\sim$ 2s

- Before the main exposure is started, the guiding is tested by monitoring the guide star position offsets from the initial position with the guiding turned on. A maximum error of 1 guider pixel in both x and y directions is set. If either of the errors in the x and y positions exceed this maximum value, the whole process is repeated with another guide star.

- If there are no required pointing changes (i.e. change of target or pointing update) throughout an image sequence, then autoguiding is left on, including between successive exposures. However, if a change of target is required or a pointing update is necessary (after evaluating the last image's plate solution), then autoguiding stops and the whole process is repeated before taking the next exposure.

For this system to work well, it must be calibrated. ACP has its own autoguiding calibration script, which tests both the direction and magnitude of a guide star's shift in the guider image in response to 'test' pulses sent to the mount. This behaviour is learnt, and the appropriate values automatically entered into the Maxim DL set-up window. These parameters allow for the inevitable rotation of the guider imager with respect to the RA & Dec. axes, and the variation in distance travelled (by the guide star in the image) in either direction due to a single corrective pulse.

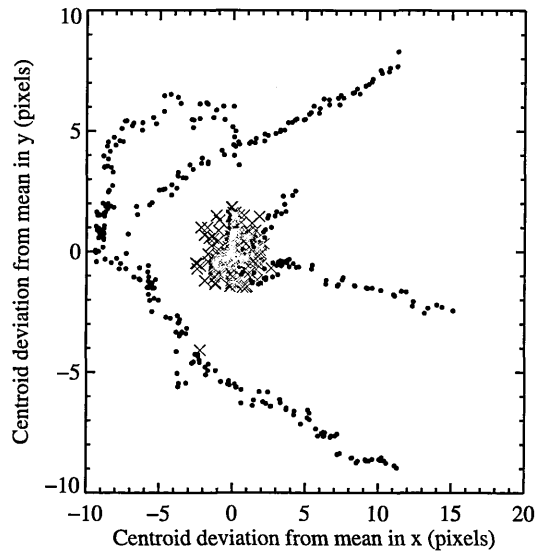

Figure 2.6: Position variations of a given star during a night with no specified maximum pointing error (black circles) and a night with maximum pointing error set to 3″ (red crosses). This figure is Fig. 5 in Holmes et al. (2011)

2.2.4.6 Weather monitoring and safety shutdowns

Weather monitoring by ACP is accomplished via the third-party 'WeatherWatcher Server' software. The weather server is started automatically upon starting ACP, and ACP remains connected to the server permanently. The server only monitors data from the Boltwood cloud sensor, specifically the data file written to disk by the Clarity II software (see subsection 2.2.5.1). The WeatherWatcher server is therefore constantly monitoring the 'sky - ambient' temperature (to determine cloudiness), the wind speed, the humidity, and the Boltwood's rain sensor. The WeatherWatcher server determines whether any of these weather conditions breach the user-set safety limits, and simply presents to ACP a 'safe' or 'unsafe' flag parameter. ACP therefore routinely checks the WeatherWatcher server for the status of this flag. If at any point the 'unsafe' flag is seen, ACP automatically interrupts any running scripts and executes the weather script. The weather script was modified from the default script (which simply transmits a full shutdown command) so as to only close the dome. The modified script (which can be seen in Alg. A.1 in the Appendix) simply determines the present status of the dome and closes it if currently open. It leaves all other components on, to allow for the swift restart of the system and continuation of the observing schedule if the dome has been shut due to transitory heavy clouds. It is worth noting that execution of the weather script automatically prevents the telescope

from tracking at the sidereal rate, bringing it to a halt.

Whilst ACP only makes use of the Boltwood information for safety monitoring, the other systems and software present on PIRATE control PC 2 (used solely for weather monitoring) are made available to users either directly via a VNC connection to PC 2 or via the weather page at pirate.open.ac.uk. This extra information is incredibly useful for gaining a greater overall understanding of the present meteorological conditions in Costitx.

2.2.4.7 File organisation and indexing

ACP automatically stores files in a structured fashion, according to who is using the facility. For all web-based users, all of the images & log files are stored in a separate 'web docs' folder. Within the web docs folder is a further segregation by type, with all images heading to the 'images' folder, likewise all logs to the 'logs' folder. Within each of these folders is a folder for each user, split by username. Upon entering the user's own folder, each night's data is split into separate folders on a night-by-night basis, with file folders taking the form: 'yyyymmdd'. Within each date folder, and alongside all of the 'light' frames is another folder called 'Calibration', which holds all of the Dark and Bias frames. Parallel to the file folders (in date form) the sky flats are stored by date within their own 'Autoflat' folder (this is no longer true of ACP version 6). The folder structure is handled automatically by ACP, and provides a simple way to keep files organised.

The user is guided to prepend files with the following convention:

$$[\{\text{letter}\}\{exp\}\_\{\text{number}\}]$$

Here, the core of the file prefix is 'exp'; all science images for processing by the pipeline must follow this convention. The letter before 'exp' denotes a particular batch of images (in sequence), attributed to a particular target, which is denoted by the number at the end of the prefix. For example, if an observation run of a particular target (e.g. the first target of the night) must be split in two (perhaps to observe a target of opportunity in the intervening gap), then the first batch will have the prefix 'aexp_1'. Upon resumption of observations of this target, the second batch should have the prefix 'bexp_1'. Once observations of that target have ceased, new observations of the second target commence with the prefix 'aexp_2'; and so on. This file name convention is not imposed by ACP, but is required by the PIRATE

	$T_{sky-amb} \leq -25°C$	$-25°C < T_{sky-amb} \leq -10°C$	$-10°C < T_{sky-amb}$
Classification	Clear	Cloudy	Very Cloudy
Safety Status	Safe	Safe	Unsafe due to possibility of rain clouds (weather safety script initiates)

Table 2.3: The Boltwood 'cloudiness index', $T_{sky-amb}$ and its split into three categories, as employed by ACP for weather safety monitoring purposes

pipeline, as it helps with file organisation and maintains timestamp sequence through the alphanumeric sorting of filenames.

2.2.5 Weather software

PIRATE's weather information is fed to a separate control PC, designated PC02, that monitors meteorological data from the various weather sensors attached to the PIRATE facility. The software and their corresponding instruments are detailed separately below, and a snapshot of the programs running can be seen in Fig. 2.7

2.2.5.1 The Boltwood cloud sensor II & Clarity II software

The Boltwood Cloud Sensor, designed by Boltwood Systems Corporation, measures the amount of cloud directly above the observatory by recording the flux from the sky in the $8\mu - 14\mu$ infrared band, using this to compute the 'sky temperature', and comparing this with the ambient temperature at the observatory level. Clouds of all types are particularly emissive across the range of this band, whereas the upper atmosphere is not (due to being significantly cooler), and so a high detected flux by the sensor is indicative of clouds (at a higher temperature than the upper atmosphere). The difference, $T_{sky-amb} = T_{sky} - T_{amb}$ is an indicator of the level of cloudiness of the night sky, therefore. The Clarity II software splits the quantity $T_{sky-amb}$ into three categories: 'Clear', 'Cloudy', and 'Very Cloudy'. The boundaries between each category are user-set. The values currently in place (as of December 2011) for the PIRATE facility are detailed in Table 2.3.

In situations where the cloud level is deemed 'unsafe', the Weather Watcher server (running on PC01 but reading the Clarity II output log from PC02) employed by ACP sends an unsafe flag to ACP (which is constantly monitoring this output flag for any

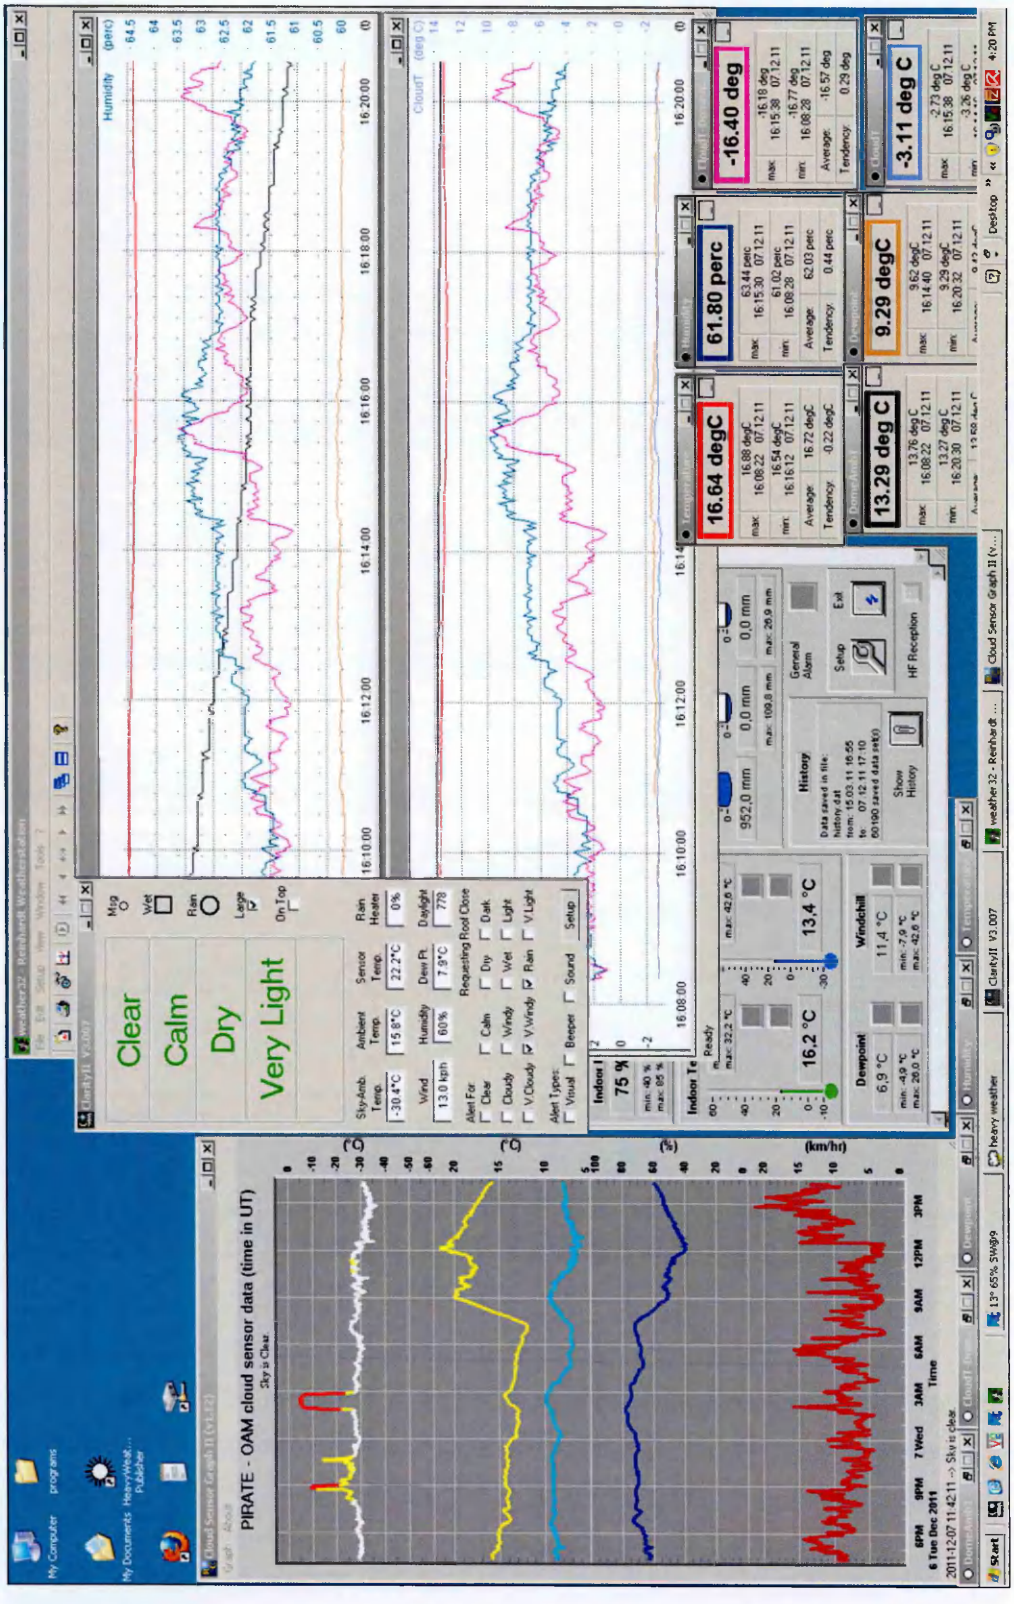

Figure 2.7: PC02, and the weather software running on it.

Category	Unsafe Trigger Point
Wind Speed	$\geq$ 40km h^{-1}
Relative Humidity	$\geq$ 83%
$T_{sky-amb}$	> -10°C

Table 2.4: The three triggers used by the Weather Watcher server for ACP to signal that conditions are unsafe to ACP.

change), and ACP initiates the weather shutdown script. The Boltwood sensor also has a 'hard' connection to the dome's control hardware, also initiating a dome close if Clarity II detects unsafe conditions. This acts as an extra safety feature as it bypasses the control PCs entirely.

Alongside measuring the Sky-Ambient temperature difference, the Boltwood also makes measurements of windspeeds through a stationary anemometer, measures the relative humidity, determines if it is raining, and has a daylight sensor to distinguish between day & night. These readings are also read by ACP's Weather Watcher server, with humidity constraints & wind speed measurements used to send unsafe flags when conditions breach the respective thresholds. The conditions that trigger an unsafe flag are detailed in Table 2.4.

2.2.5.2 Weather32 / Reinhardt weatherstation

The Reinhardt Weatherstation is the dome's own weather station. It contains equipment similar to the Boltwood for making measurements of $T_{sky-amb}$. There is also a humidity monitor and thermometer within the dome, as well a rain sensor outside, consiting of two spoked metal plates offset by a millimetre or so, conductive contact between which is made when water droplets hit the sensor, sending a hard signal to the dome electronics to shut in the event of rain. All of these parameters are recorded by the Weather32 software, though these are typically only used for reference and have no influence over other PIRATE hardware. Typically we see systematic offsets between the values recorded with this instrument and the Boltwood, most notably in the recorded Relative Humidity (RH). High levels of humidity can be a problem in Mallorca, and so effective shut-down at dangerous (and ultimately corrosive) levels of humidity is a necessity. We have ACP shutting down at a Boltwood-determined RH of 83% as this corresponds to anywhere between 90-93% when measured by the Reinhardt system. We consider a RH above 90% to be dangerous and at the limit of operation, hence the humidity safety setting in the Weather Watcher server is set to 83% to correspond to the higher Reinhardt values.

2.2.5.3 Heavy Weather

The Heavy Weather software reads data from another weather station on site that sits beside the PIRATE facility. It provides information on rain levels (recorded in mm), relative humidity, relative pressure, and has a directional anenometer. It provides a useful check on the validity of all other weather information provided by the PIRATE facility.

2.3 Hardware characterisation

2.3.1 Camera characteristics

This subsection contains material taken verbatim from Holmes et al. (2011).

Whilst the imager currently used in the PIRATE facility, the STX-16803, has already been detailed as defining the transition from PIRATE Mark 1.5 to PIRATE Mark 2, for this body of work we largely neglect it; as no commissioning work was done with the camera at the start of writing this thesis. We therefore restrict the data included in this thesis to that generated from the original camera (of Mark 1 & 1.5), the STL-1001E.

To characterise the CCD we determined the gain, linearity, read noise and the dark current (which was measured at a range of temperatures). The gain and the linearity measurements were made using the same set of dome flats for each. A set of dome flats of increasing exposure time were taken in batches of increasing order to measure the median counts of a 100 $\times$ 100 pixel subframe in the middle of each image. At the maximum exposure time for the sequence a new batch of images was started again at the shortest exposure time. This allowed us to monitor for any drifts in the light level. For the short exposure times (of order a few seconds), we attempted to generate a shutter correction map using the methodology of Zissell (2000), anticipating that we might see position-dependent corrections to the exposure time (due to the shutter travel) of order 10^{-3}s. However, this produced a null result, and further investigation into the shutter mechanism of the STL-1001E confirmed a rotating 'shutter wheel' that should be devoid of shutter travel effects. We therefore do not apply any shutter correction to the short exposure flat fields. We measured the bias level from contemporaneous bias frames and subsequently subtracted the mean pedestal level of 107 ADU from each flat. The subframe was chosen to be the centre of the vignetting function, where the image is at its flattest. To assess the

T (°C)	Dark Current ($e^{-}s^{-1}$)	$\sigma_{\text{Dark Current}}(e^{-}s^{-1})$
-25	0.01	0.02
-20	0.02	0.02
-15	0.05	0.02
-10	0.08	0.02
-5	0.27	0.03
0	0.51	0.03
5	1.13	0.03
10	2.25	0.03
15	5.20	0.04

Table 2.5: Measured dark current for a range of chip temperatures

linearity of the CCD, we plot median subframe counts against exposure time (Fig. 2.8 (a)), and fit a linear trend to the same ADU range. The residuals of this fit can be seen in Fig. 2.8 (b). We note a deviation from linearity at the top end of the dynamic range of $< 1\%$, and $< 2\%$ at the bottom end. We measure the gain using the following relationship: $\sigma^2_{\text{ADU}} = (1/g)\,\langle N_{\text{ADU}}\rangle$ Howell (2006), where σ_{ADU} is the standard deviation of the sub-frame counts, and $\langle N_{ADU}\rangle$ is the mean recorded counts in the subframe. To derive this expression, one assumes the statistical relationship $\sigma_{e^-} = \sqrt{N_{e^-}}$ holds, where σ_{e^-} and N_{e^-} are the uncertainty in and the number of recorded photoelectrons respectively, which only holds when the process is governed by Poisson statistics. To determine g we therefore plot σ^2_{ADU} against $\langle N_{ADU}\rangle$ (see Fig. 2.8 (c)). We need sufficient photoelectrons for photon-counting (Poisson) noise to be dominant, and so we fit to the range 20k ADU $< \langle N_{ADU}\rangle <$ 40k ADU. We choose an upper limit of 40k ADU to be well clear of the digital counting limit of 65535 ADU (16-bit). We measure a value for the gain of 1.62 ± 0.01 $e^{-}ADU^{-1}$.

To measure the read-out noise, we took 30 bias frames at a temperature of $-20°$, and median combined them. The median-combined bias frame (which we assume to have negligible read noise) was then subtracted from each of the 30 bias frames, and the standard deviation of each of these difference images measured; we averaged this figure to achieve the read noise in ADU (Howell, 2006). We measure then (using the previously determined gain value) a read-out noise of 10.9 ± 1.3 e^{-}. To check for consistency, we also split the bias frames into pairs, subtracting one from the other within each pair, so that $d = b_1 - b_2$, where b_1 & b_2 are the bias frames in each pair, and $\sigma_{read} = \sigma_d/\sqrt{2}$, yielding a read-noise measurement for each pair. These values were consistent with the previous read-noise measurement.

Dark current was assessed for a range of temperatures (from 15°C through to

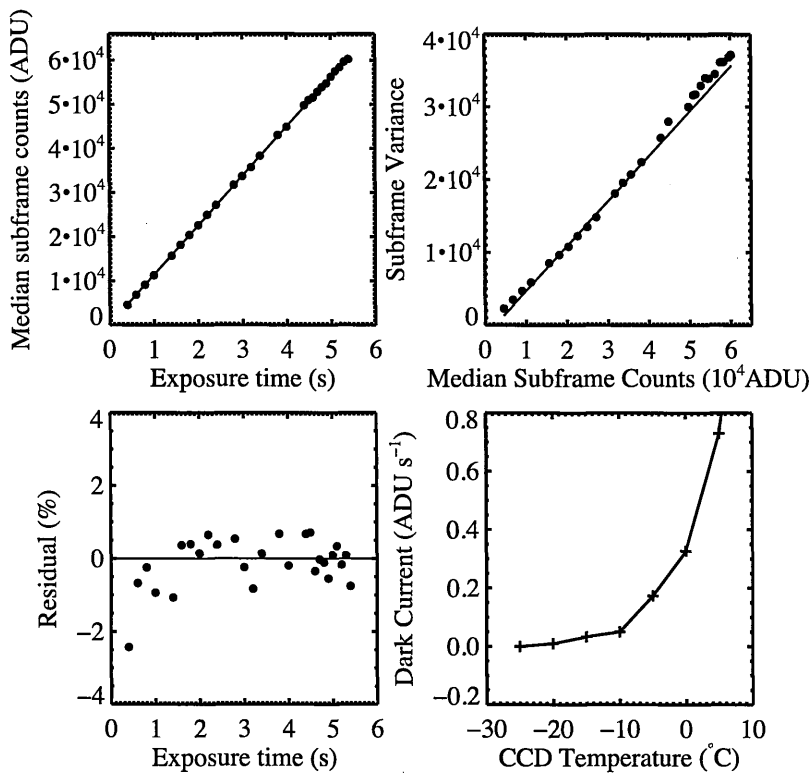

Figure 2.8: *Top left:* Counts as a function of exposure time, with a linear fit through the expected linear region 20,000–40,000 ADU. *Bottom left:* Residuals of the preceding fit. *Top right:* Photon transfer curve, with linear fit to the same 20,000–40,000 ADU region. The deviation from linearity at ~0–10000 ADU is due to the transition from the shot noise to the read-noise regime. *Bottom right:* Dark current as a function of chip temperature.

-25°C in increments of 5°C) by median-combining 10 darks at each temperature, 10 biases at each temperature, subtracting the bias and dividing the residual counts by the exposure time (30s for each dark frame). The measured values are listed in Table 2.5. We measure a dark current of 0.02 ± 0.02 $e^{-}s^{-1}$ and 0.08 ± 0.02 $e^{-}s^{-1}$ at -20°C and -10°C, the typical operating temperatures for winter and summer respectively. We also find a value of 0.51 ± 0.03 $e^{-}s^{-1}$ at 0°C, c.f. the manufacturer's specification of 1 $e^{-}s^{-1}$ at 0°C.

2.3.2 Flat-field inacuracy and optical response calibration

This subsection contains material taken verbatim from Holmes et al. (2011). This work was originally performed on PIRATE Mark 1 data. Subsection 2.3.3 repeats the following methods, but for PIRATE Mark 1.5 data, drawing a comparison between the previously published results from PIRATE Mark 1 and the new data from PIRATE Mark 1.5

The systematic errors of flat fields as optical response calibrators are often overlooked in the application of differential photometry, as modern auto-guiding systems are sufficiently good at locking the stellar point-spread functions to fixed positions on the chip for the duration of a time series. As the observer is only interested in time-variability, systematic offsets in the magnitude zero-points of any two stars in the field can be ignored. PIRATE's GEM executes a 'pier-flip' as the tracked target crosses the Meridian. This moves the OTA to the other side of the pier; and inverts the image of the stellar field with respect to pre-flip frames. This operation effectively moves each stellar point spread function to an entirely different part of the focal plane. To maintain continuity in the flux ratio between two stars across the pier flip, the vignetting must be determined perfectly. In reality, most flat fields suffer from inaccuracies at the 10^{-2} level (Manfroid, 1995). We therefore expect and indeed see varying flux ratios between two objects in the field across the pier-flip. We term this effect a 'Light Curve Discontinuity' (LCD). Several mechanisms contribute to the flat field systematic errors:

- Uneven illumination from the source. This can occur in sky flats, even when they image the sky null point (Chromey and Hasselbacher, 1996) as PIRATE's automatic procedure does. It will also occur in dome flats.

- Physical changes in the hardware between taking flat fields and science images that modify the response function of the system. In the case of PIRATE Mark

1, it is suspected that primary mirror 'flop', which occurs despite the presence of mirror locks, and/or flexure in the micro-focuser assembly changes the vignetting function. It may also lead to the misalignment of the focal plane with the CCD surface, meaning consistent focus across the field of view is not achievable. In this scenario, the FWHM of the stellar profiles will be a function of chip position, and the measuring apertures used for photometry would see different levels of flux leakage due to the changed stellar profile after a pier-flip.

- Scattered light in the optical tube assembly. Off-axis, unfocused stray light due to insufficient baffling can contribute to the flat field exposure. Scattered light is additive. If included in the flat field, this additive light is incorrectly used in the multiplicative flat field calibration.

We have been unable to develop a wholly effective procedure for self-consistently calibrating the LCD effect without investing much of the night in taking calibration observations. We thoroughly explored these possible approaches:

We used data from the night of 23/07/2009, during which the moon had 3.2% illumination, so we might expect sky background gradients to be limited. The data were processed in the usual manner, and investigated separately as two groups of 'pre-flip' and 'post-flip' frames. The frames in each group were median-combined, and a 6^{th} order Legendre polynomial was fit to the sky background in each of the two (pre and post) resultant frames using the IRAF task 'IMSURFIT'. Note that the pre-flip frames were rotated 180°. We denote the background fits by $A_{pre}(x, y)$ and $A_{post}(x, y)$, where x and y are image co-ordinates corresponding to positions on the sky (not pixel co-ordinates, due to the aforementioned rotation of the pre-flip frames). A simple ratio of the pre and post-flip background fits reveals any discrepancy in the sky background for a given star position. We denote this ratio as $A_{map}(x, y) = \frac{A_{pre}(x,y)}{A_{post}(x,y)}$. We show A_{map} in Fig. 2.9 (top) . We note that it has structure predominantly in the x direction, and displays a peak-to-peak variation of $\sim$ 7%. The (non-differential) light curves of each of the N stars (here $N = 517$), which we denote by $F_i(t)$, where $i = 1...N$, were median-combined to create an approximation (to first order) of the sky transparency, $\tilde{F}$. We assume each star remains at a fixed position in the image (x_i, y_i) for the duration of the observing run. For all of the light curves (including the transparency function), the fluxes were averaged over time for pre-flip frames and also for post-flip frames, producing a single pre-flip and post-flip flux value for each light curve. We therefore use $F_{pre,i}$ and $F_{post,i}$ to refer to these time-averaged values.

We define the LCD of a star to be:

$$\Delta F_i = \frac{F_{pre,i}/F_{post,i}}{\tilde{F}_{pre}/\tilde{F}_{post}} \tag{2.1}$$

We now have the ability to assess the applicability of the background 'map'; if the shape of the sky background (after flat-fielding) provides an estimation of any residual error in the flat-fielding process, then it should correlate well with the observed LCDs. To compare the map with the LCDs, we fit a 2^{nd} order polynomial through least squares regression to the irregularly gridded ΔF_i values. The result of this fit can be seen in Fig. 2.9 (bottom). As can be seen, far from correlating well, the pattern of the LCDs also exhibits greatest deviation in the x direction, but with opposing orientation. The structure of the map in Fig. 2.9 (top) suggests a fixed and constant light source co-moving with the optical tube assembly, inducing the same background structure in both pre and post-flip frames, which is then amplified when the pre-flip frames are rotated. One would expect any transformation of the vignetting function to be recorded in the sky background, as this gives a continuous indication of the vignetting function. However, the presence of any scattered light in either the science frames or flat fields renders the true vignetting function for each side of the pier flip unrecoverable.

The scattered light structure present in the science frames could also have been present in the flat field, and thus the LCDs introduced to the data through the flat fielding process alone, but the un-flat fielded data shows the same LCD structure seen in Fig. 2.9 (bottom). Given the dark night (one night after a new moon) for this data set, it is apparent that scattered (additive) light is likely always present, and greatly reduces the accuracy of standard flat-fielding and sky-flat procedures; the pier flip required by the GEM simply highlights this inaccuracy.

The data of 23/07/2009 exhibit strong structure in the LCD map, and are hence a good demonstration of the problem. A more typical structure can be seen in Figs. 2.10 (top) and 2.10 (bottom), from the night of 23/11/2009, WASP-12 field. The polynomial fit to the LCDs exhibits a smaller peak-to-peak amplitude of $\sim$2%, and there is significant scatter in the residuals of this fit to ΔF_i.

We conclude that background fitting, sky flats, and even flat fielding are all ineffective at removing the percent-level LCD effect seen in PIRATE Mk. 1. Instead, there are two routes to achieving correct photometric calibration across the pier-flip.

The first of these involves creating a photometric super-flat. This is discussed in detail in Manfroid (1995), Selman (2004), Boyle (2007), Grauer et al. (2008), and

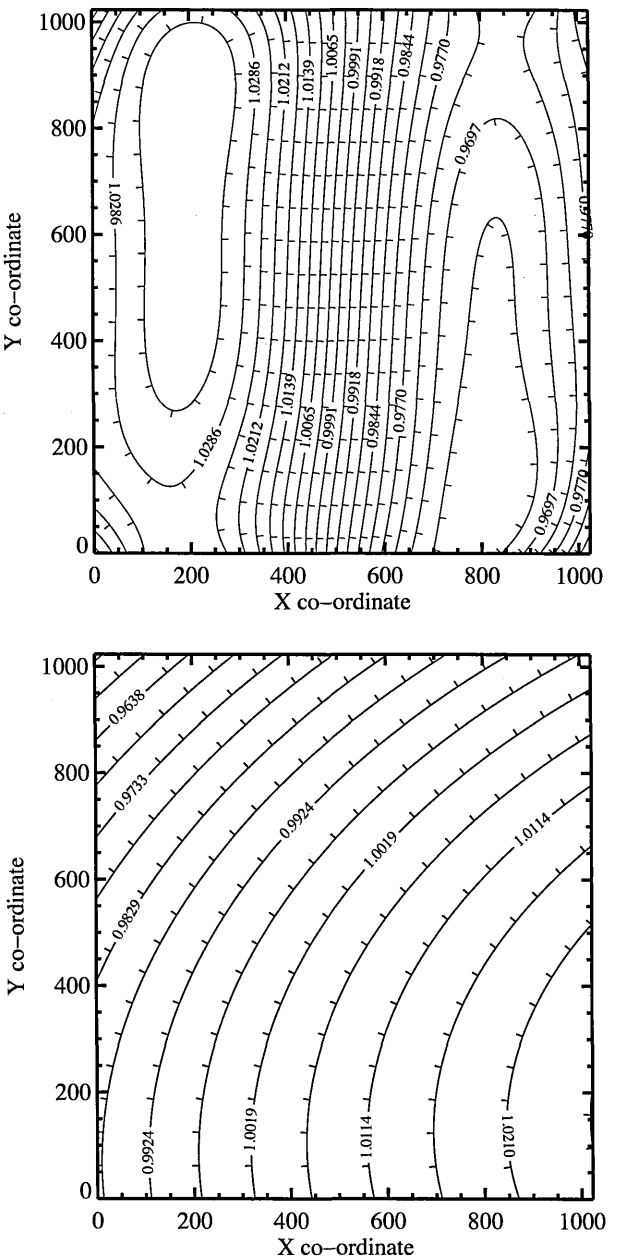

Figure 2.9: *Top:* Sky background map produced (for data taken 23/07/09) by fitting to the sky backround for all pre-flip frames and all post-flip frames, then taking the ratio of these two fits. The structure of this map suggests that there may have been a fixed scattered light source with respect to the OTA. *Bottom:* Second-order polynomial fit to the flux deficits, ΔF_i.

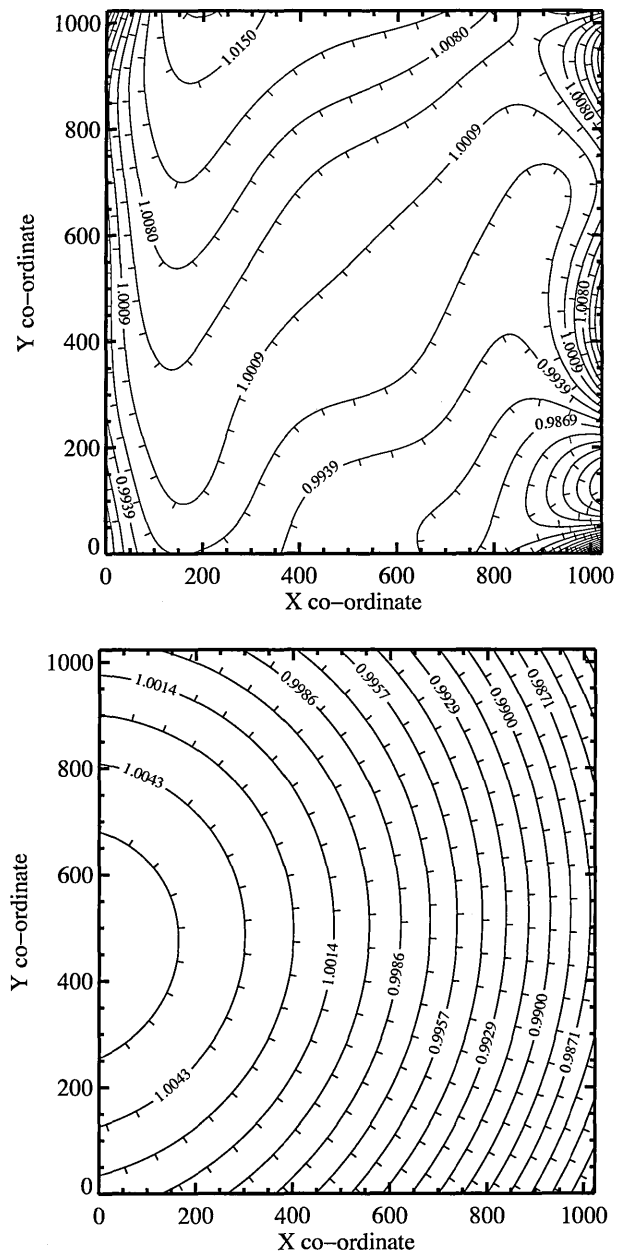

Figure 2.10: *Top:* The same as Fig. 2.9 (top), but for 23/11/2009. *Bottom:* The same as Fig. 2.9 (bottom), but for 23/11/2009.

Regnault et al. (2009). At its most simple, this involves observing a set of standard stars at different positions in the field of view, and determining the position-dependent response to the standards in order to build the large-scale, low-frequency optical response of the system into the final flat field. Sky flats or dome flats are observed in tandem, with polynomial fits to their large scale variation made in order to 'flatten' them. This retains the small scale, high frequency intra-pixel response, which can then be combined with the standard star-determined low frequency response function to create a 'photometric super-flat'. In constructing the large-scale component, it is preferable to make multiple dithered observations of a cluster containing many standards that span the field of view in order to reduce the number of frames needed. As we expect the response to be different from one side of the pier to another, this involves observing the same cluster twice, once in the eastern sky, and again later in the western sky.

The second method involves lending a degree of freedom to the normalisation of each light curve section (pre & post-flip), shifting the flux levels up and down to minimize the χ^2 of a model fit. This method has its limitations in the extent of fore-knowledge required to work successfully. In the case of SuperWASP follow-up, the Markov-Chain Monte Carlo fitting routine of Collier Cameron et al. (2007) can take the light curves from each side of the pier as separate input light curves and include the normalisation factor as a free parameter in the fitting procedure. However, this algorithm works with 'prior knowledge' of the light curve, in that its initial parameters from which the MCMC routine iterates are taken from the SuperWASP light curves. This prior expectation of the light curve's model parameters (including, crucially, the transit ephemeris) assists with the cross-flip normalisation. If there is no prior knowledge of the light curve, adjusting flux levels is almost certainly perilous. Take, for example, the worst case scenario of a pier-flip occurring mid-way through ingress or egress. In such a scenario it would prove difficult to determine if the flux deficit across the pier-flip is instrumental or astrophysical. Great caution is called for in interpreting such light curves.

2.3.3 Investigation of PIRATE Mark 1.5 data

Since the publication of the work on the flat field accuracy and optical response of PIRATE Mark 1 in Holmes et al. (2011), PIRATE was upgraded to PIRATE Mark 1.5, via a change of OTA from the Celestron C14 to the PlaneWave CDK17. The inconsistent optical performance of the C14 was suspected to be the dominant cause

of LCDs in PIRATE follow-up data, so the switch to an upgraded OTA provided a perfect opportunity to compare the two OTA units. We investigated whether the LCD phenomenon is still present in Mark 1.5 data, and to what extent. In summary, LCDs are still present, but they are significantly reduced in magnitude. To demonstrate the effect of changing the OTA, the analyses of the previous subsection 2.3.2 are repeated here, but for Mark 1.5 data. We focus on the observation run of 27/07/2011, a SuperWASP candidate follow-up observation (see Chapter 5), as it has a sufficient number of frames taken from both sides of the pier included in one run (94 eastern sky, 66 western sky). Fig. 2.11 shows the same maps as in Figs. 2.9 & 2.10, but repeated for the night of 27/07/2011. Fig. 2.11 shows that the sky background map, formed from the ratio of sky background between pre- and post-flip frames, shows no semblance of similarity to the 2$^{\text{nd}}$ order polynomial fit to the flux deficits, ΔF_i, as per PIRATE Mark 1. However, the magnitude of the flux deficits / LCDs is greatly reduced from those seen in the Mark 1 data, with the vast majority of the frame showing $0.995 \lesssim \Delta F_i \lesssim 1.005$. Whilst only one night's analysis is included here, the reduction in the scale of the LCD effect across all Mark 1.5 data is perfectly evident from the light curves produced with the Mark 1.5 data. The scale of the LCD phenomenon has been reduced from a 1-5% effect or more with the C14 to a $< 1\%$ effect with the CDK17; with the effect often being reduced to $\ll 1\%$ in ideal cases.

So what might be the cause of this improvement? In the previous section it was postulated that one of the main causes of the LCD effect might be the variation in image focus throughout the focal plane, specifically in the section of focal plane inhabited by the detector. Point spread functions whose full-width-half-maximums vary as a function of position in the focal plane are troublesome, as a measuring aperture of fixed radius used on a star in one part of the image will capture a different percentage of the star's flux to the same aperture used in another part of the image, where the typical FWHMs are different. The difference in the amount of a star's flux captured between two stars is still fine when performing aperture photometry, so long as the ratio of the different amounts captured of the two stars remains fixed for the duration of an observation run. When the OTA is pier-flipped however, the required persistent ratio is instantly violated. The only scenario in which this isn't the case, is when the variation in FWHM across the chip's FOV is a perfectly symmetric radial function about the centre of the image. From initial inspection, it appears that the C14's curved focal plane is significantly bettered by the significantly flatter focal plane of the CDK17.

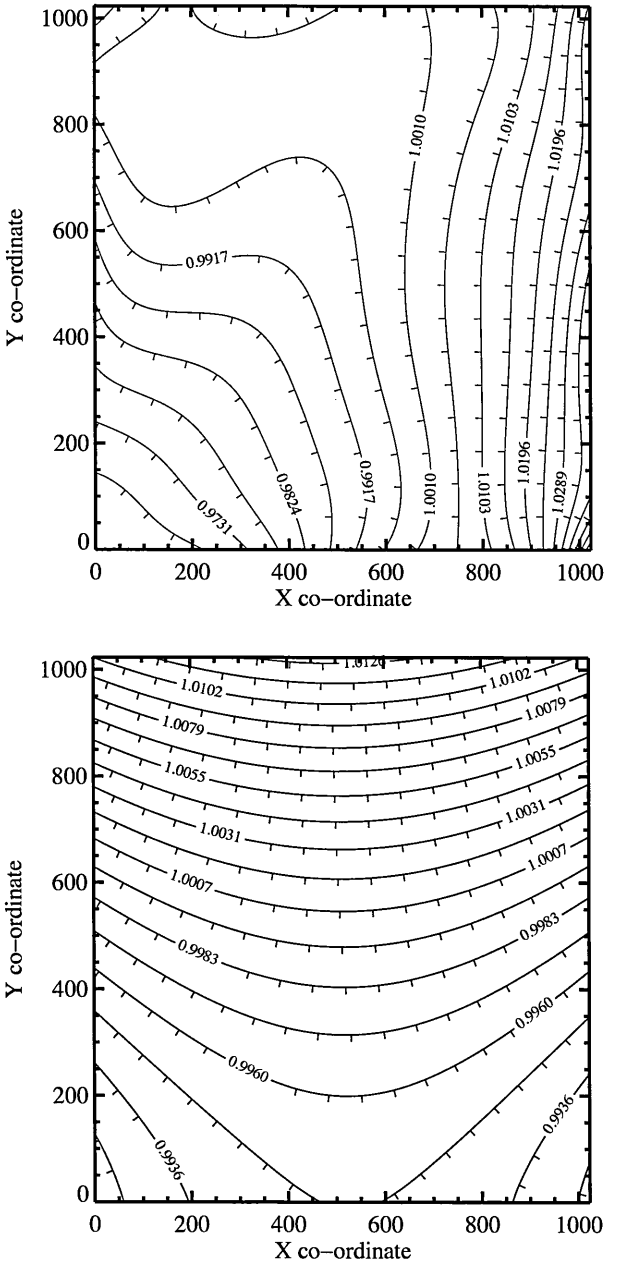

Figure 2.11: As per Figs. 2.9 & 2.10, but for PIRATE Mark 1.5 data (27/07/2011)

To test this, we return again to the observation runs of 23/11/2009 (PIRATE Mark 1), and 27/07/11 (PIRATE Mark 1.5), and attempt to map the variation in PSF FWHM across the extent of the STL-1001E's field of view. All of the observation frames taken in the observation run are included. Data from each frame are normalised and aggregated into one map representative of the entire run; the details of how this is done follows. The IDL routine GAUSS2DFIT is employed to fit an elliptical gaussian function at the locations of the stellar sources as detected in the PIRATE reduction pipeline (see Chapter 3). Each source is assigned a FWHM (in pixels) that is the mean of the minor and major axes in the elliptical gaussian. For each frame in the observation run, a linear fit is made for all stars in the field of the form $\omega = a_1\Delta_{mid}(x,y) + b_1$ where ω represents the individual sources' FWHM, and $\Delta_{mid}(x,y)$ their Pythagorean distance from the centre of the image to the location of each source. The distance to the centre of the image is used as it is rotationally invariant. For each frame, the measured FWHMs of all the sources in that frame are then normalised by dividing through by b_1. This normalisation to the FWHM value at the middle of the image should ensure some homogeneity across all of the frames, and allow them all to contribute to the FWHM map by correcting for any changes in the zero-point of the FWHM (i.e. due to drifts in the quality of atmospheric seeing throughout the run). Once this is done for each frame in the observing run, a 2^{nd} order polynomial is fit through the irregularly gridded normalised FWHM values, to produce our initial FWHM map. These initial maps can be seen in Fig. 2.12. These initial contour maps reveal where in the field of view the smallest values of FWHM (i.e. best focus) reside. Using this information, the map production process is repeated, this time fitting $\omega = a_2\Delta_{foc}(x,y) + b_2$ for each frame, where $\Delta_{foc}(x,y)$ is no longer the pythagorean distance of each star from the location of the centre of the image, but the distance of pixel (x,y) from the location of the FWHM minimum in pixel coordinates. The FWHMs for each source are normalised by dividing through by b_2, and all the frames' data are amalgamated and the contour map produced again. Thanks to locating the minimum with the first iteration maps, the second set of maps now show the fractional increase in measured FWHM with increasing distance from the measured FWHM minimum. The second stage maps can be seen in Fig. 2.13.

As is readily apparent, the difference between the variation in FWHM across the FOV for the C14 and for the CDK17 is dramatic. The width of the fitted gaussians to the point sources within the C14 images can vary by as much as 50% across the frame. For the CDK17, the variation is reduced to less than 10%. Recall also that the FOV of the Mark 1.5 system is actually larger than that of Mark 1, so the variation would

be even less if the FOV was rescaled to that of PIRATE Mark 1. It is very likely that the improved optics of the CDK17 are responsible for the significant reduction in the LCD effect between PIRATE Mark1 & 1.5. The Corrected Dall-Kirkham optics of the CDK17 were designed to offer astrophotographers pin-point stars out to the very edges of their colour composite images. It is evident that this consistency of focus throughout the image is also highly desirable when attempting to generate precision time-series photometry when coupled with a pier-flipping German equatorial mount.

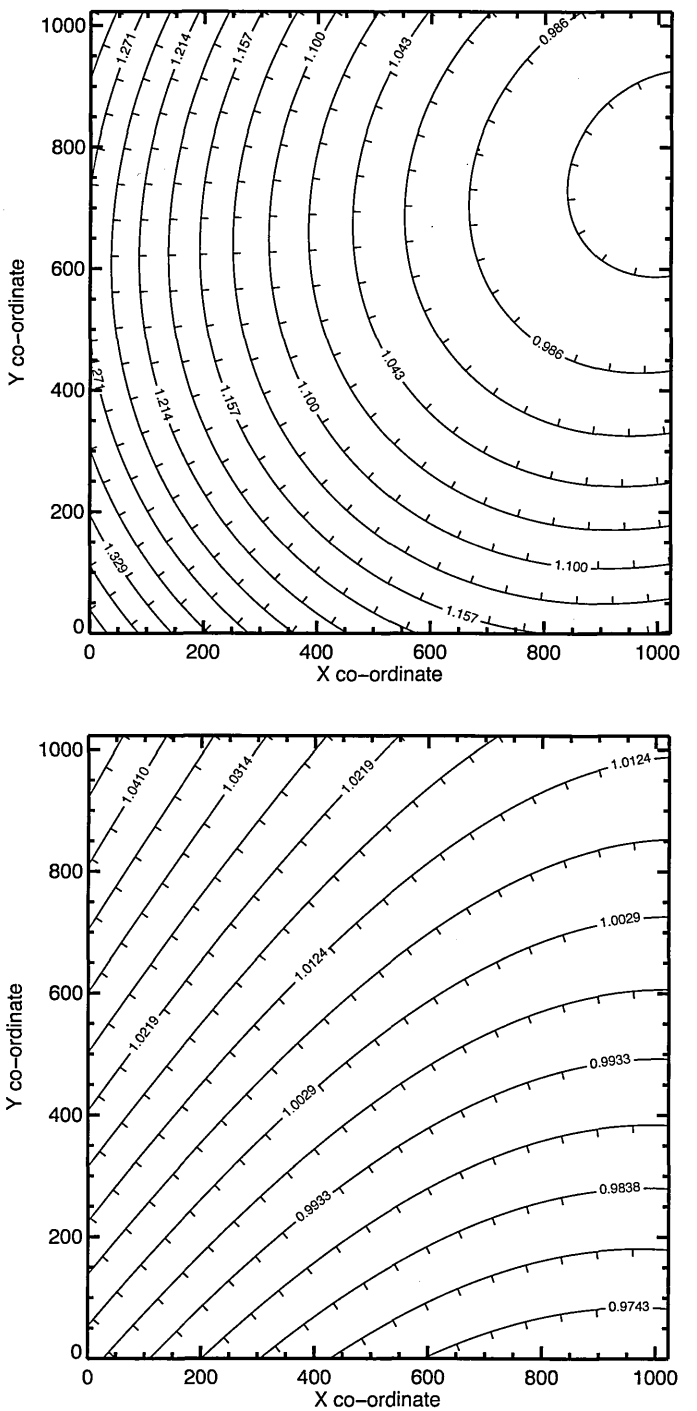

Figure 2.12: First stage FWHM maps for the night of 23/11/2009 (top), and 27/07/2011 (bottom). The FWHM values are normalised to the estimated FWHM value at the centre of the image (x = 512, y = 512).

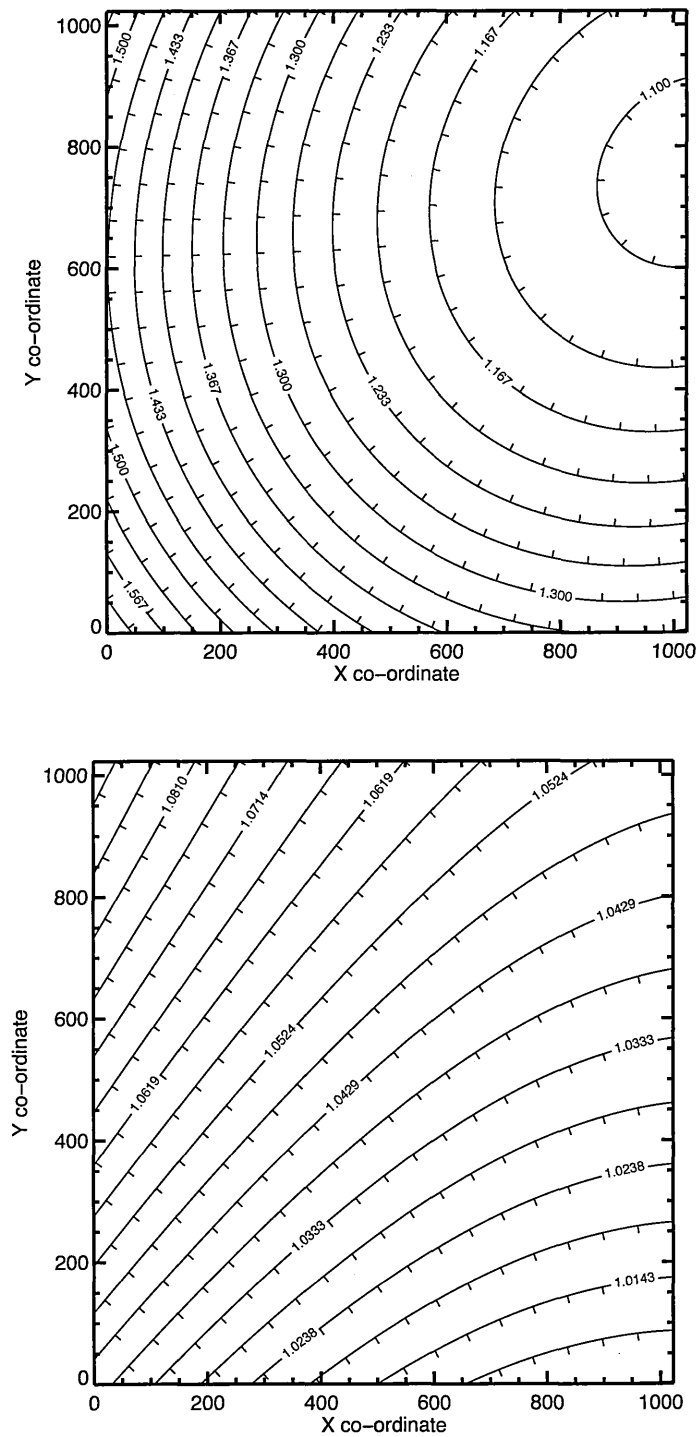

Figure 2.13: Second stage FWHM maps for the night of 23/11/2009 (top), and 27/07/2011 (bottom). The FWHM values are normalised to the estimated minimum FWHM value in the first stage maps.

Chapter 3

PIRATE Light Curve Generation

3.1 Photometric precision and sources of uncertainty

Photometry is a fundamental tool in astronomy. At its core, it is the process of measuring the total emitted flux of astronomical sources in a given part of the electromagnetic spectrum, during a finite time period. All of the time series data in this thesis are produced with aperture photometry. Aperture photometry is the process of defining a section of the detector, in this case the CCD chip, that surrounds the recorded flux from an astronomical source, and first determining the total quantity of flux within the region, and secondly determining what proportion of the included flux is attributable to the astronomical source, and what proportion is attributable to other factors (sky background, dark current, read noise, scintillation noise). Doing so allows us not only to determine the flux due to the astronomical source, and therefore learn about the source, but also understand the background noise contribution to our measurement; determining the signal-to-noise ratio (SNR) in the process.

3.1.1 Making a photometric measurement

A CCD camera is made up of an array of pixels, each of which is exposed to photons which are in turn recorded as photoelectrons during an exposure. It is not necessary to explain how this process occurs; for the purpose of this section it is simply necessary to understand that the light from a star, when focused on the CCD chip, forms a 2D distribution known as the 'Point Spread Function' (PSF). Photons originating from a point source incident on the chip are dispersed due to travelling through the Earth's

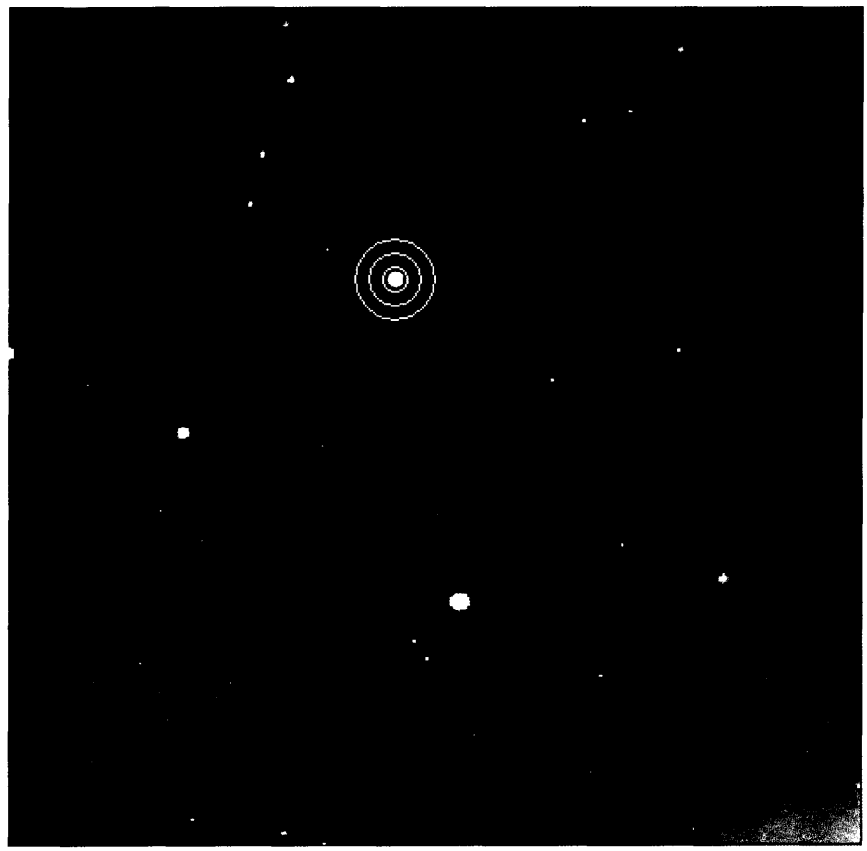

Figure 3.1: Measuring aperture and annulus centroided on a foreground star in the region of M31, taken with PIRATE Mark 2.

atmosphere, and due to general imperfection in the instrument. Atmospheric seeing is the term used to describe the process by which turbulence in the Earth's atmosphere continually changes the refractive index of the air through which the photons pass. This, combined with the imperfect nature of all instruments employed in any process of measurement spread the photons out into a peaked 2D distribution. The extent of the spread of this function (typically a measure of how good either the focus or the atmospheric seeing is) is defined by the Full Width Half Maximum (FWHM).

In attempting to measure the signal from an astronomical source we must take account of the dispersal of the photons into a PSF. We do this by using a circular measuring aperture as part of the process of 'aperture photometry', to define a region of the CCD in which we wish to count the recorded photoelectrons. This is the inner yellow circle in Fig. 3.1. Counting the recorded photoelectrons within this region allows us to ascertain the recorded flux from the star, so long as we first know what the gain (conversion rate between detected photoelectrons and ADUs) of the chip is, and secondly what the expected contribution to the recorded photoelectrons within

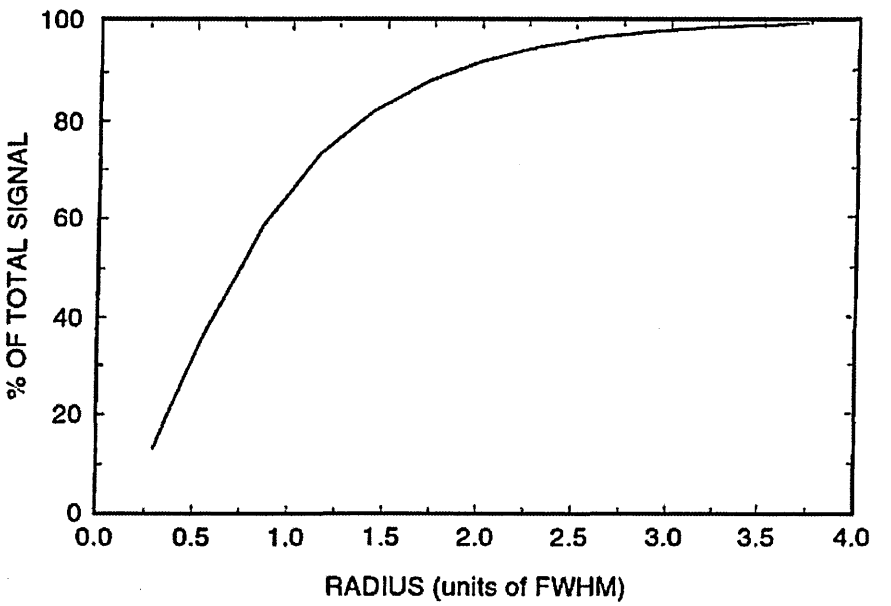

Figure 3.2: Fig. 5.6 from Howell (2006), showing the fraction of a star's signal residing within an increasing set of radii in units of PSF FWHM

the aperture is from sources other than the star; the *noise* contribution.

One such contribution is the contribution from the night sky itself, the so-called 'sky background', as no sky is perfectly dark. The second and third circles (outwards from the star's centre) in Fig. 3.1 are known as the 'outer annulus' and serve the purpose of sampling the sky background. The sky background value is determined through counting the photoelectrons recorded between the second and third apertures. No measurement is made between circles one and two, as this is simply a gap out to the outer annulus. The purpose of the outer annulus is to determine the sky background level away from the contamination of the star's light, but suitably close enough to account for localised variation in the sky background level.

At its most basic, the process of aperture photometry involves determining the flux within the central measuring aperture, determining the flux within the outer annulus (sky background), and subtracting the sky background estimation from the total flux within the central measuring aperture. Other factors that contribute to the uncertainty in the flux measurements must also be accounted for; more on these in the following subsection.

In order to determine a suitable radius for the measuring aperture, gap width and annulus size, we must look at the variation in flux originating from the star as a function of radial distance from the centre of the PSF. Fig. 3.2, which is Fig. 5.6 in Howell (2006), shows the fraction of the total flux from the star encircled by an ever

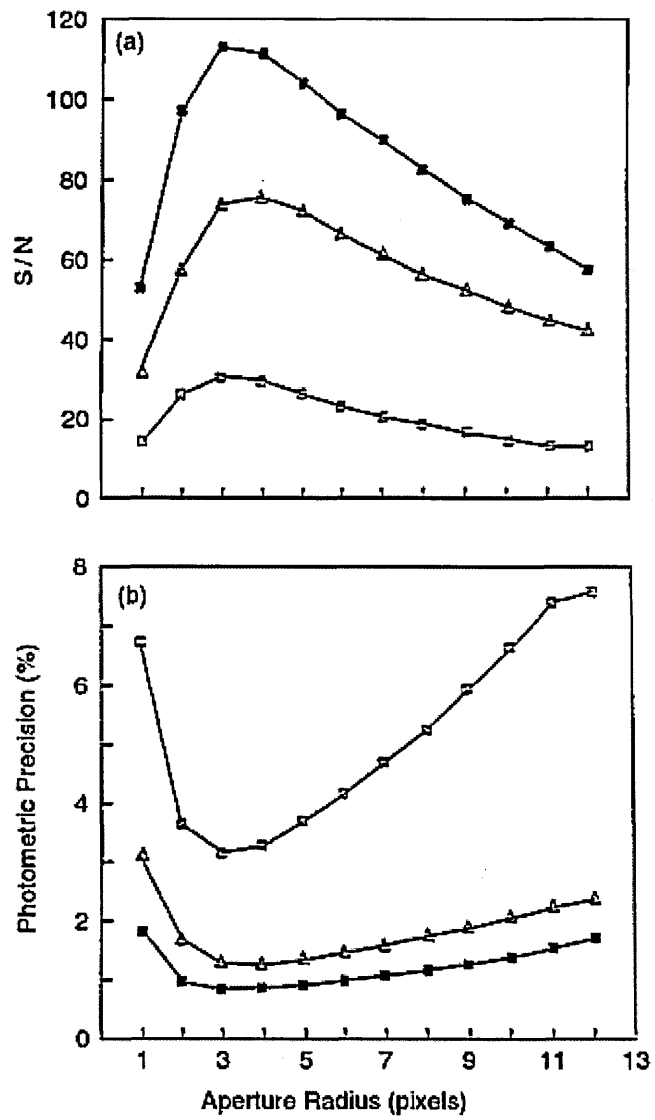

Figure 3.3: Figure 6 from Howell (1989), showing SNR and resultant photometric precision for a range of measuring aperture sizes. The plate scale of the system used to produce this data was $0.4''\text{pixel}^{-1}$, and the seeing was near $1.2''$, or 3 pixels.

increasing radius in units of the star's FWHM. As can be seen, most of the star's light is encompassed by an aperture of radius 4 × FWHM. One might look at this plot then, and determine that the measuring aperture should always be at least four times the FWHM of the star, in order to ensure that all of the star's light is captured within the inner measuring aperture, and none of it contaminates the sky background estimation from the outer annulus. In fact, as the measuring aperture is increased in size beyond $\sim 1 \times$ FWHM, the SNR drops off due to the increasing contribution from sky background flux as a proportion of the total flux within the measuring aperture. This can be seen in Fig. 3.3, which is Figure 6 from Howell (1989).The peak SNR is achieved at about $\sim 1 \times$ FWHM; though there are also difficulties associated with adopting such a small measuring aperture, which are discussed later in this chapter. The effect of increasing the aperture size on the SNR is further evaluated in the next section detailing the CCD equation.

3.1.2 The CCD equation

The CCD equation is an equation for evaluating the SNR of a photometric measurement. We will jump straight to the complete CCD equation as introduced in Merline and Howell (1995), and an explanation of each term in this equation follows. The equation provides a formula for the signal to noise ratio of either a single pixel, or a group of pixels forming the area enclosed within a measuring aperture. This particular version of the CCD equation is an extended version, and includes extra noise terms that allow for CCD observations in which there are high sky background levels, particularly faint sources, poor spatial sampling or large gain values:

$$\mathrm{SNR} = \frac{N_*}{\sqrt{N_* + n_{\mathrm{pix}}\left(1 + \frac{n_{\mathrm{pix}}}{n_B}\right)\left(N_S + N_D + N_R^2 + G^2\sigma_f^2\right)}} \tag{3.1}$$

Here N_* is the number of photoelectrons collected from the source, n_{pix} is the total number of pixels within the measuring aperture (if investigating the SNR of just one pixel, $n_{\mathrm{pix}} = 1$), n_B is the number of pixels used to estimate the sky background level, N_S is the number of photoelectrons per pixel from the sky background, N_D the number of thermally induced photoelectrons per pixel (dark current), N_R the read noise per pixel, G the gain, and σ_f the uncertainty introduced by the A/D converter and the discrete nature of Analogue-to-Digital-Units (ADU).

This particular version of the CCD equation extends the pre-existing version, $\frac{S}{N} = \frac{N_*}{\sqrt{N_* + n_{pix}(N_S + N_D + N_R^2)}}$ Howell (1989), with the introduction of the term $\left(1 + \frac{n_{pix}}{n_B}\right)$, which allows for added uncertainty introduced through the use of insufficient pixels to estimate the sky background levels, and the term $G^2\sigma_f^2$, which accounts for 'digitisation noise'. Digitisation noise is the uncertainty introduced by converting analogue voltages into digital, discrete ADUs. After dividing the analogue signal into digitised 'packets' or 'counts', any remaining charge is not accounted for. Note that the uncertainty term σ_f is therefore combined with the gain - so this noise term vanishes in anything but high-gain systems. PIRATE Mark 1 and Mark 1.5 use the STL-1001E with a measured gain of $1.62 \pm 0.01 e^- ADU^{-1}$, so this term is essentially irrelevant for PIRATE data.

For 'bright objects', which this thesis focuses almost entirely on (see Chapters 4, 5 and 6), by which we effectively mean those which are not 'photon-starved' and just above background level, the equation effectively collapses to:

$$\frac{S}{N} = \frac{N_*}{\sqrt{N_*}} = \sqrt{N_*} \tag{3.2}$$

a simple poissonian SNR relation.

Due to the low gain of the STL-1001E, the minimal dark current and low read noise (see Chapter 2) , the dark current, read noise and digitisation noise terms can be neglected from this equation for PIRATE follow-up observations of bright objects. The justification for doing so is presented in

3.1.3 Correlated noise

The CCD equation deals entirely with contributions to the uncertainty in the flux that are random in nature. Whilst some noise components in the CCD equation may follow a non-gaussian distribution (e.g. photon noise is poissonian), all components will, via the central limit theorem, combine to produce a gaussian distribution. All of these components are 'white noise' contributions. There is, however, another colour of noise present in the PIRATE light curves (and, to varying extent, all light curves from all instruments), which is 'red noise', also known as correlated noise. PIRATE light curves apparently show some time-correlation. For a given deviation of a single data point from the light curve's median value, the probability that the following data point will deviate by a similar amount is greater than it being governed purely

by a gaussian distribution about the median. In effect, there is a 'communication' between the noise properties of consecutive data points.

We can attempt to quantify this using the methodology of Winn et al. (2009). To do so, we take the light curves (with known rms, σ_{lc}) output from the PIRATE data reduction pipeline (discussed in the next subsection) for a given night's observations and average the data points into M bins, each of N data points, and measure the rms, σ_N of the binned values. If the light curves display only white, gaussian noise properties, we should expect that:

$$\sigma_N = N^{-1/2} \left[M/\left(M-1\right) \right]^{1/2} \sigma_{lc} \tag{3.3}$$

If correlated noise is indeed present, as it appears to be, then in fact $\sigma_N = \beta\sigma_{lc}$ (where $\beta > 1$).

3.2 The PIRATE data reduction pipeline

Early on in the project it became obvious that the data reduction process needed to be automated to decrease the amount of time taken to reduce a data set. It was expected that large quantities of data per night would be obtained when the SuperWASP follow-up observation programme was fully fledged; so the ability to extract the necessary results from the nightly observations quickly and efficiently was paramount. Prior to this, solely the IRAF IMRED (for image calibration) and DAOPHOT (Stetson, 1987) package was used (specifically the aperture photometry tasks) to perform aperture photometry on all detected sources in every frame in an observation run. This can be a cumbersome process that can take anywhere from an hour to a few hours for a given observation run. In order to script some of this procedure and provide an element of automation, IDL was employed to work in conjunction with IRAF to produce a reduction 'pipeline'. Thanks to the NASA IDL Library (Landsman, 1993) adding a large toolset of FITS image processing tools to IDL's functionality, it was seen as a good choice as a scripting 'engine' for the pipeline. As IDL is not known for its speed (as it is not a compiled language but interpretative), other alternatives might have produced a faster-operating pipeline, such as PyRAF (Greenfield and White, 2000), which replaces IRAF's Command Language (CL) with one based on Python, an object orientated scripted language. However, this author's lack of familiarity with Python, coupled with the wide use of IDL within the department, made IDL

the obvious choice. Aside from the routines in the respective IRAF packages and the NASA IDL library routines, the pipeline is entirely this author's code. An explanation of the pipeline routines can be found in the Appendix.

3.2.1 Pipeline operation

The reduction pipeline is essentially split in two halves. The first of these is responsible for interfacing with IRAF (with some elements of user interactivity required), and is tasked with housekeeping issues such as file organisation, calibration and performing statistical tests on fits files. This half also performs aperture photometry via the IRAF DAOPHOT package. The second half of the pipeline has no use for IRAF, and works on an ingested 'cube' of data that contain all of the salient information from the IRAF photometry tasks, operates on these data, and produces the final light curves for all of the objects in the field.

3.2.1.1 Calibrating the science images

Upon being told which folder to operate within, the pipeline searches for FITS files of different types (science frame, flat field frame, dark frame, bias frame) within the folder, not by the IRAF header entry 'IMTYPE', but by the conventions employed by ACP in file naming. The user is asked to either interactively go through each science frame in sequence, answering 'yes' or 'no' to the question of whether the file should be included or excluded in the analysis, or provide an input list of 'bad' frames that the user wishes to pre-exclude. Statistical tests to perform this task would ideally replace this section of the pipeline to increase the level of automation in the future, but they must be robust. Due to the way in which IRAF works, if a frame is badly aligned or has no sources, IRAF will stop mid-process if unable to determine the extent of a particular image's displacement from that of the master frame, the operation of the pipeline cannot be recovered from this point, and the user must start from the beginning again. Ways in which this procedure can be improved are discussed in Chapter 7.

Based on the content of the folder, a script for IRAF is prepared and executed that calibrates all frames in the usual manner using the IMRED package. For example, if no flat field frames are present in the folder, the script will not include flat-field calibration in the process (it is up to the user to identify suitable replacement flat

fields if no flat fields were taken on the night). For images taken with the STL-1001E camera, dark frame subtraction is not included in the process, as the dark current shows no noticeable structure across the chip, so the step is removed so as not to introduce further noise during the calibration process. Once the files have been calibrated, all images taken with the OTA West of the pier (i.e. images of the Eastern sky) are rotated 180 degrees. This ensures the orientation of the star field is consistent across all of the images in the observation run. The next stage involves performing the aperture photometry.

3.2.1.2 Performing the aperture photometry

For this stage of the reduction, a 'master' frame must be selected. The frame used as the master frame should be the 'best' of the run, i.e. it should be taken under the best seeing conditions, and contain the most point sources. The pipeline suggests to the user a candidate for the master frame via a short statistical test, choosing the frame which has the highest standard deviation from the quartile of frames that have the lowest mean ADU value. This ensures a frame with a low sky background and plenty of stellar sources is selected as the master frame. The purpose of the master frame is to establish which stars are included in the catalogue of detected sources upon which aperture photometry will be performed across all of the frames. The input catalogue is established by first detecting all sources above a certain detection threshold in the master frame, then performing aperture photometry about these positions. All sources for which the aperture photometry is successful are included in the catalogue.

For this step the user must supply a value for the stellar FWHM (in pixels) representative of the master frame's seeing. This value, in conjunction with an input source detection threshold given in multiples of the sky background uncertainty (6σ is typically used) is fed into the DAOPHOT task 'DAOFIND' to detect the stellar sources above the provided detection threshold. Once the preliminary list of sources has been obtained, aperture photometry is performed, and the final input catalogue consists of all sources for which successful photometry is performed. Typically, the sources that fail to make it into the final input catalog are at the far extremities of the frame, where the aperture and/or sky background annulus is clipped by the edge of the frame. The positions of the measuring apertures are then linearly transformed from their position in the master frame to each respective frame in the observing run using the IRAF task IMCENTROID. Once a coordinate input list has been generated for each frame, aperture photometry is performed for all sources in the input catalogue

across all frames. The user may specify the radius of the measuring aperture, in multiples of the FWHM in the master frame. Typically, a radius of 3 × FWHM has been employed for PIRATE Mark 1 data, whereas 2 × FWHM has been employed for PIRATE Mark 1.5 data. We can employ a smaller aperture with PIRATE Mark 1.5 thanks to the CDK17 superior optics that yield a 'flatter' image, with FWHMs of the stellar PSFs that are significantly more stable across the field of view than for the C14. Reducing the aperture size yields an increase in S/N (as the ratio of star light to sky background per pixel increases with reducing radial distance from the PSF centre - see subsection 3.1.1). However, a smaller aperture is more sensitive to changes in PSF FWHM across an image. If the PSF shape is a function of position in the focal plane, then for any given aperture size, the fraction of the star's light captured by the measuring aperture is also a function of position on the chip. Systematic effects in the resulting photometry arising from any guiding or tracking errors will therefore be more prominent with smaller apertures, as stars change their profiles as they move in and out of different regions of the image. We therefore adopt a wider aperture for PIRATE Mark 1, due to its non-flat focal plane and increased guiding error. The superior optics of the CDK17, combined with a tuned and improve guiding system, allow us to use smaller apertures for the PIRATE Mark 1.5 data, and receive a consequent boost in aperture signal.

The gap width (in pixels) between the measuring aperture and the outer annulus, as well as the width of the outer annulus are also available as inputs. Typically values of 3 & 10 are used for both.

Uncertainties in the photometric measurements are taken directly from the IRAF PHOT tool estimation for the uncertainty in the flux, which differs from the denominator in the CCD equation established in section 3.1.2:

$$\sigma_{flux} = \sqrt{N_* + n_{\text{pix}} \left(1 + \frac{n_{\text{pix}}}{n_B}\right) \sigma_{N_S}^2} \tag{3.4}$$

Where the notation is the same as that used for eq. 3.1. Note that IRAF does not use its sky background estimation in this equation, but instead the variance of the sky background level; under the assumption that the two are equal via equation 3.2. This relationship is an estimation as it holds only for situations where the SNR $\gg 1$, due to the read noise, dark current and digitisation noise terms being neglected from the full CCD equation. However, as previously explained this is a perfectly fine approximation

for PIRATE data as this thesis deals only with the follow-up of bright objects; and doesn't deal with faint sources (for which these terms could not be neglected).

3.2.1.3 Producing the final light curves

The result of performing aperture photometry on the input catalogue across all of the science frames is a sequence of files, one for each image, that contain the UTC (Coordinated Universal Time - see Chapter 6, subsection 6.1.2.5) timestamp, airmass, nominal 'Star IDs', x-coordinates, y-coordinates, fluxes, instrumental magnitudes, and uncertainties in the instrumental magnitudes for all the stars in the input catalog. These files are ingested by the pipeline, and collated into a data 'cube'. The cube is transposed so each 'slice' features the time sequence of measurements for each star. At this point, the UTC time stamp of mid-exposure is converted to Heliocentric Julian Date after extracting the target's right ascension and declination from the master frame FITS header.

A further stage of frame rejection (beyond the manual input at the start of the pipeline procedure) is introduced here, based on the number of stars with unsuccessful photometric measurements (as reported by the DAOPHOT 'phot' task) in a given frame. The user can set a 'failure threshold' that all frames must pass. The failure rate is defined as thus:

$$F = \frac{n_{\text{fails}}}{n_{\text{stars}}} \tag{3.5}$$

where n_{fails} is the number of registered photometry failures for a given image/frame, and n_{stars} is the number of detected sources in the input catalogue. A frame i is removed from the cube if F_i for that frame exceeds the specified threshold. Typically, a value for the failure threshold of 0.1 works well to weed out any genuinely bad frames from the data set; typically those where transient cloud has dramatically reduced the number sources in the frame.

After the second frame-cut, any stars that continue to suffer any photometry failures throughout the entire duration of the observing run are exluded from the catalogue of stars. This prevents stars that drift (via changes in sky transparency) into or out of either the saturation limit of the chip (digital saturation occurs at 65535 ADU) or the magnitude limit at the faint end from being included in the analysis. Any differential photometry analysis that includes a star that hits such a limit during the observing sequence will be invalidated; which is why any star remaining that suffers a fail throughout the duration of the run is culled. This star 'cull' step makes the

frame failure threshold step very important. If one frame in particular is contributing a high number of photometry failures, then there will be a significant reduction in the number of stars included in the final analysis. Keeping a stringent value for the failure threshold ensures that plenty of stars remain in the analysis until the end.

Once the data cube has been trimmed for both frames and stars, the final light curves are produced using an 'optimal ensemble compilation' technique described in Section 3.3. In short, the pipeline attempts to find the optimal ensemble of comparison stars to perform differential photometry against for each star in the input catalogue. Each star's final light curve is produced from a ratio of its own flux versus the combined flux of its own unique comparison ensemble. Further detail is provided in Section 3.3. With the final light curves produced for each remaining star in the data set, the final component `investigateblends` is run. This module is focused on the SuperWASP follow-up photometry programme, as it compiles the light curves of all detected sources in close proximity (within the largest SuperWASP aperture size) to the intended target star. More of this module can be seen in Chapter 5, with the output from this module shown in Figs. 5.4 and 5.5.

3.3 Optimal ensemble compilation

3.3.1 How the pipeline computes final light curves

The `ensemble` module within the pipeline attempts to compile an 'optimal comparison ensemble' that is unique to each star. It uses a method that approximates that used in Burke et al. (2006), where for each star, all the corresponding comparison light curves (using all stars within the data set) are sorted according to a light curve 'figure of merit'. Each comparison light curve is normalised to unity, and these light curves' RMS values (i.e. the standard deviation about the mean) are used as the figure of merit. Starting with the best-rated light curve (i.e. lowest RMS) for a given star, each subsequent light curve in the sorted list is iteratively included in the ensemble, using inverse variance weights, where the weights derive from the IRAF photometric uncertainties. Note that the ensemble of comparison stars never contains the flux from the target star, as the target star is excluded from the list of light curves. The following refers to the production of the final light curve for one star only. The ensemble flux, E_j, for a given frame j, containing comparison stars i =1, ..., N, is compiled as follows:

$$E_j = \sum_{i=1}^{N} \omega_{ij} F_{ij} / \sum_{i=1}^{N} \omega_{ij} \tag{3.6}$$

where F_{ij} are the stars' individual fluxes, and the individual weights are given by:

$$\omega_{ij} = \frac{1}{\sigma_{ij}^2} \tag{3.7}$$

The uncertainty in the ensemble is given by the standard deviation of the weighted average:

$$\sigma_{E_j} = \left[\sum_{i=1}^{N} \omega_{ij} \right]^{-1/2} \tag{3.8}$$

With each new star added to the ensemble, the differential target star light curve $d_* = (d_{*,1}, ..., d_{*,j...})$, where $d_{*,j} = {}^{F_{*,j}}/_{E_j}$ $(j = 1, ..., N)$ is calculated from the recorded flux of the target star, $F_{*,j}$, and calculated ensemble flux, E_j. The light curve d_* is then normalised by dividing through by its median value, and the RMS is computed. The final light curve, $d_{*,\mathrm{fin}}$, makes use of the ensemble that produces the lowest RMS; each star therefore has its own unique comparison ensemble.

3.3.2 Optimal ensemble compilation performance

To assess the performance of the ensemble compilation routine, we can look at a typical field of stars, in this case the field surrounding the star WASP-12, and look at the RMS versus apparent magnitude for all stars in the field with light curves output by the pipeline. Rather than convert instrumental magnitudes to absolute magnitudes through standard star observations (which were not taken alongside the WASP-12 observations, and are therefore unavailable), the USNO-B1 catalogue (Monet et al., 2003) is used to determine apparent magnitude. To do so, routines from the NASA IDL library (Landsman, 1993) were used to first convert source locations (in pixel coordinates) to equatorial coordinates (via the plate solve parameters in the FITS headers). These equatorial coordinates were then used to query the USNO-B1 catalogue for R band apparent magnitudes, using a $3''$ search radius, selecting the brightest source within the search radius. This provides some tolerance for differences between the determined equatorial coordinates (using the WCS parameters) and the catalogue coordinates. It does however introduce some source confusion, which will lead to some spurious outliers in the RMS versus apparent magnitude plot.

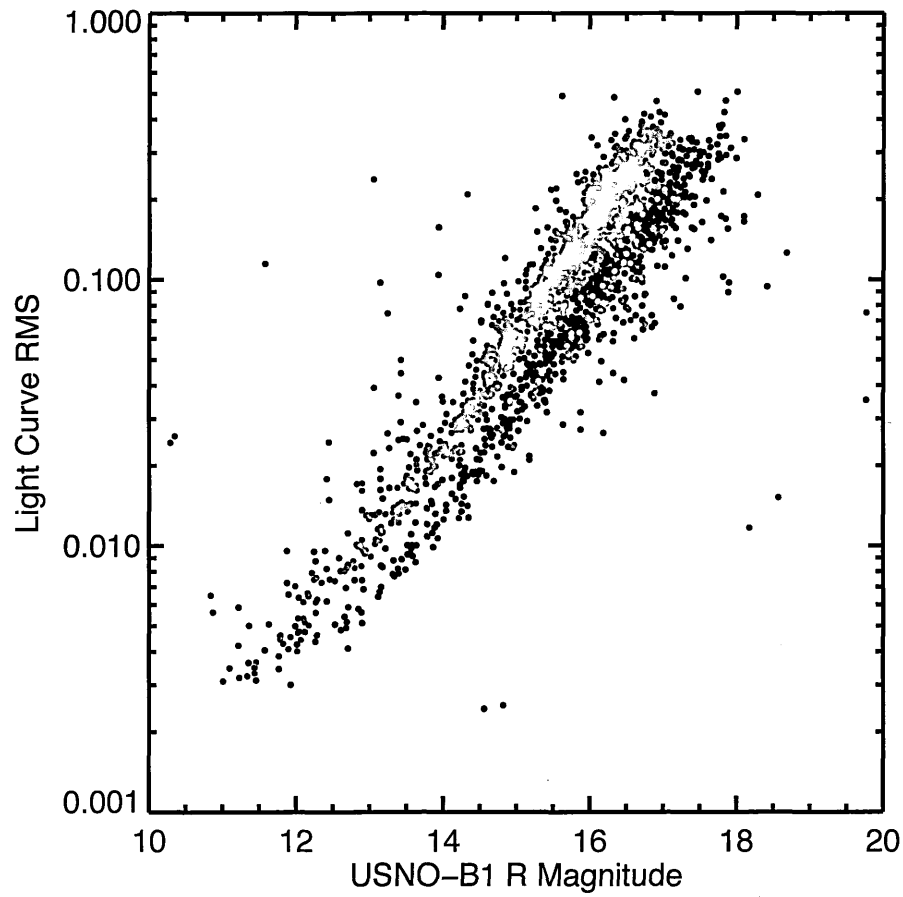

Figure 3.4: Light curve RMS versus USNO-B1 R magnitudes for the field of WASP-12 for: (black) PIRATE Mark 1, on the night of 23/11/2009, and (red) PIRATE Mark 1.5 on the night of 11/01/2011

Fig. 3.4 shows RMS versus R band magnitude for the field surrounding WASP-12 for both PIRATE Mark 1 data (23/11/2009 - black) and Mark 1.5 data (13/01/2011 - red). The data were taken under differing conditions where seeing was approximately 2.6″ on the night of 23/11/2009, and 3.4″ on 13/01/2011. It is apparent from both plots that millimag precision is achievable for all stars with magnitude $R \lesssim 13$. Outlying stars with RMS > 0.01mag and $R \lesssim 13$ are the result of source confusion by the catalogue-matching algorithm leading to faint stars being mistakenly ascribed low R magnitudes.

In Fig. 3.5 I have shown the data of 11/01/2011, this time with the instrumental magnitudes plotted against light curve rms, with the contributing elements of the CCD equation overlain (excluding the digitisation noise term). For all objects with

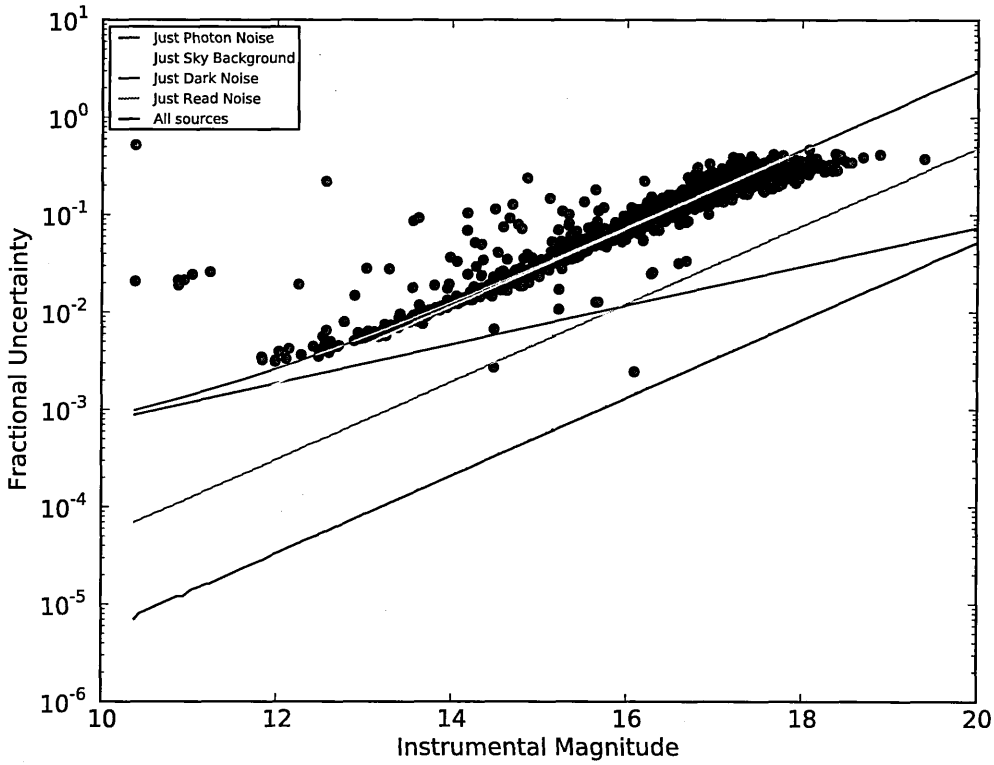

Figure 3.5: Breakdown of the contribution from each of the terms in the CCD equation (Eq. 3.1) overlain on the data of 11/01/2011

instrumental magnitude $\lesssim 13$ the data is photon noise limited, whereas the dominating contribution for all objects with instrumental magnitude $\gtrsim 13$ comes from the uncertainty in the sky background estimation (sky noise). Read noise's contribution to the noise budget is an order of magntiude less than the sky noise, and the contribution from dark noise is an order of magnitude less than the read noise. Fig. 3.5 shows that the optimal ensemble compilation routine performs well for all stars in the field, with most light curves' rms measurement in line with the theoretical lower limit provided by the CCD equation.

3.3.3 Scintillation noise

Not included in the error budget so far is scintillation noise. Scintillation noise (see e.g. Kornilov et al., 2012) is caused by turbulance in the atmosphere, which changes the refractive index of the atmosphere as light waves pass through it. This can result in phase distortions from the interplay of constructive and destructive interference at ground level. Scintillation noise is a noise familiar to us all, as it is responsible for

giving stars their 'twinkle'. It is hard to measure, and its extent is hard to infer from pre-existing data. Fluctuations in the flux from a given star due to scintillation noise can occur on very short time scales, changing by a factor of 2 over a time scale of minutes.

Attempting to understand the extent of scintillation noise in PIRATE data can be achieved by employing the scaling law of Gilliland et al. (1993):

$$\sigma_{\text{scint}} = 0.09 D^{-2/3} X^{1.75} \frac{e^{(-h/h_0)}}{\sqrt{2t_{\text{int}}}} \tag{3.9}$$

where D is the diameter of the telescope in cm, X is the airmass, h is the height of the observatory above sea level (m), h_0 is a scale factor for the atmosphere, and t_{int} is the integration time of the exposure. Fig. 3.6 details the expected value of σ_{scint} for a range of exposure times and airmasses for both PIRATE Mark 1 and 1.5. The scaling law suggests that even in the worst observing scenario (30s exposure time, PIRATE Mark 1 aperture size, airmass = 2) scintillation noise contributes less than 0.2% to the error budget.

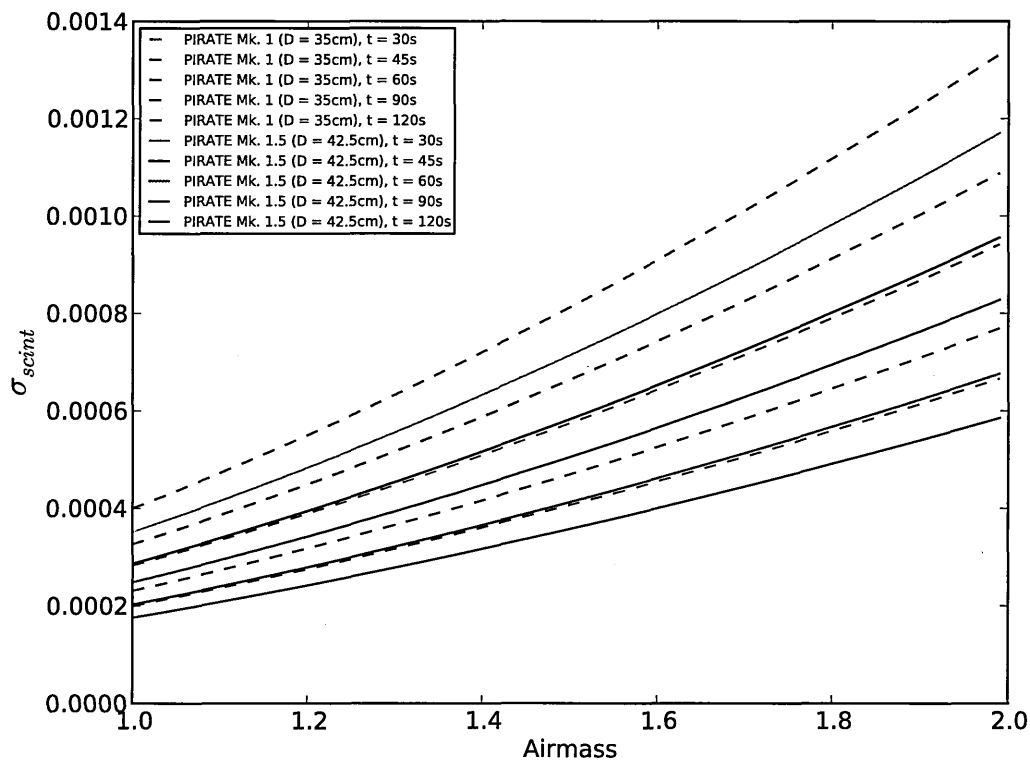

Figure 3.6: Scintillation noise as a function of airmass (as given by the scaling law of Gilliland et al. (1993)) for a range of exposure times (30s, 45s, 60s, 90s, 120s) for both PIRATE Mk. 1 and PIRATE Mk. 1.5

Chapter 4

SuperWASP candidate follow-up observations

4.1 The SuperWASP experiment

This section contains a detailed overview of the SuperWASP experiment. The level of content should not indicate involvement in the set-up and development of the experiment by this author. This author was involved in some 'eyes-on' candidate selection and prioritisation, but the systems and operation of the experiment were firmly in place before the start of this PhD. The detail, however, is provided in order to set the follow-up observations of SuperWASP planetary candidates in sufficient context.

4.1.1 Overview

The Wide Angle Search for Planets (WASP) consortium was established in 2000 with an aim to produce a ground-based, wide-field transiting exoplanet survey. After promising initial results with the prototype instrument, WASP0 (Kane et al., 2004), in 2000 and 2001, funding was secured to produce the SuperWASP (North) instrument (Pollacco et al., 2006), which saw first light in November 2003 in La Palma. A second instrument to cover the Southern hemisphere, SuperWASP South (SuperWASP-S) was added in 2005, at the South African Astronomical Observatory (SAAO).

The main science aim of the SuperWASP instruments (North and South) is to search for bright transiting exoplanet systems suitable for spectroscopic follow-up observations (used to confirm or deny the planetary nature of the system). It is a 'wide and shallow' survey, as opposed to so-called 'pencil-beam' surveys, such as the Optical Gravitational Lensing Experiment (OGLE) (Udalski et al., 2002), extending from $V \sim 7.0 - 15$, with better than 1% precision for stars with $V \sim 7 - 11.5$. As a result, the survey produces light curves of millions of objects, each with thousands of data points. To date, there are about 70 WASP planets discovered with SuperWASP-N and SuperWASP-S, encompassing a wide range of masses, from the sub-Saturnian WASP-29b (Hellier et al., 2010) to the $10M_J$ WASP-18b (Hellier et al., 2009). To date, SuperWASP has been the most successful ground-based transiting exoplanet survey.

4.1.2 Hardware

The hardware consists of 8 Canon 200 mm telephoto lenses (f/1.8) mounted on an Optical Mechanics equatorial fork mount, in a battery formation. Attached to each of these telephoto lenses is an Andor DW436, with back-illuminated CCDs manufactured by e2v that consist of 2048×2048 13.5μ pixels. With a gain of ~ 2, linearity is better than 1% across the dynamic range, and readout noise is decent at $\sim 8{-}10$ e^- across all cameras. Each detector is connected via a PCI card to one Data Acquisition System PC each, so that the large quantities of data from each detector are segregated and can be dealt with independently.

Combining these detectors with the 200mm f/1,8 lenses yields a plate scale of $13.7''\text{pixel}^{-1}$, and a FOV of ~ 61 deg^2 per camera. Each unit is fitted with a broad-band filter with a passband limited to $\sim 400 - 700$ nm. Exposure times for each camera are fixed at 30s, during which time ~ 0.6 $e^-\text{pixel}^{-1}$ of dark current are accrued, due to the CCD being cooled to -50 °C.

4.1.3 Operation

The overall observational strategy of the SuperWASP telescopes is to observe fields that have a large number of stars, to increase the chances of detecting exoplanets, but to avoid the Galactic plane where the stellar image density is greatest, to mitigate problems of confusion in source identification and to reduce instances of blended eclipsing binaries mimicking planetary transits (Christian et al., 2006). In a typical

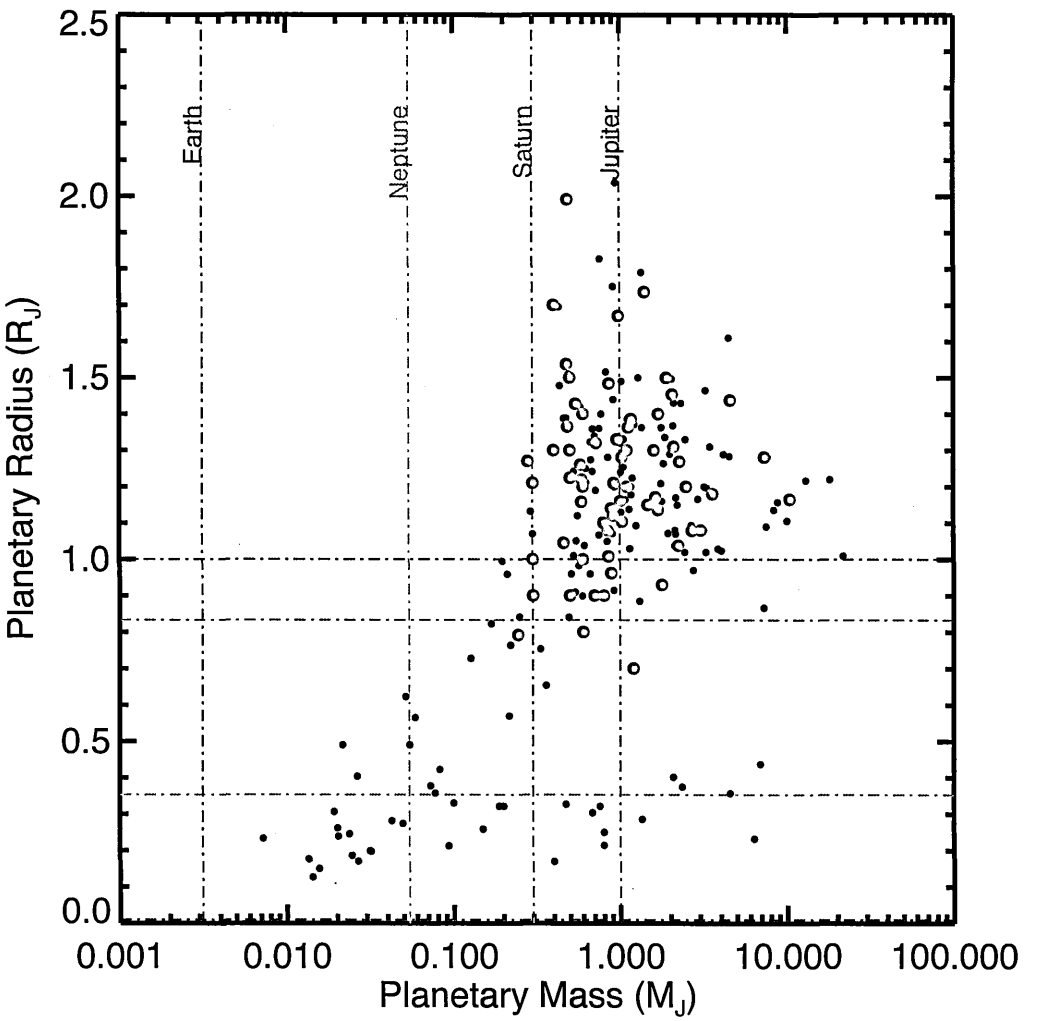

Figure 4.1: Planetary mass-radius plot for all of the WASP planets (large red points) and the transiting exoplanet population as a whole (smaller grey points). Data acquired from exoplanets.org, February 2012(Wright et al., 2011).

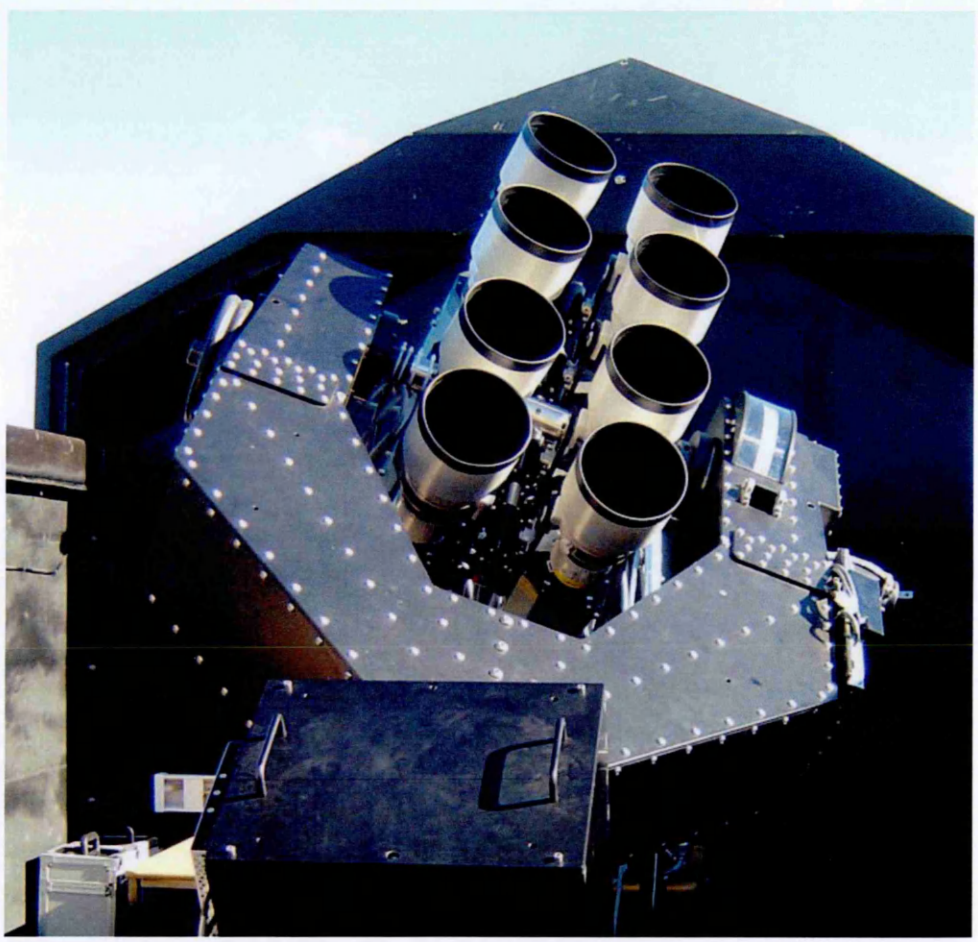

Figure 4.2: The SuperWASP-N hardware, situated at the Observatorio del Roque de los Muchachos on La Palma, showing the arrangement of the eight 200mm f/1.8 telephoto lenses on the equatorial fork mount.

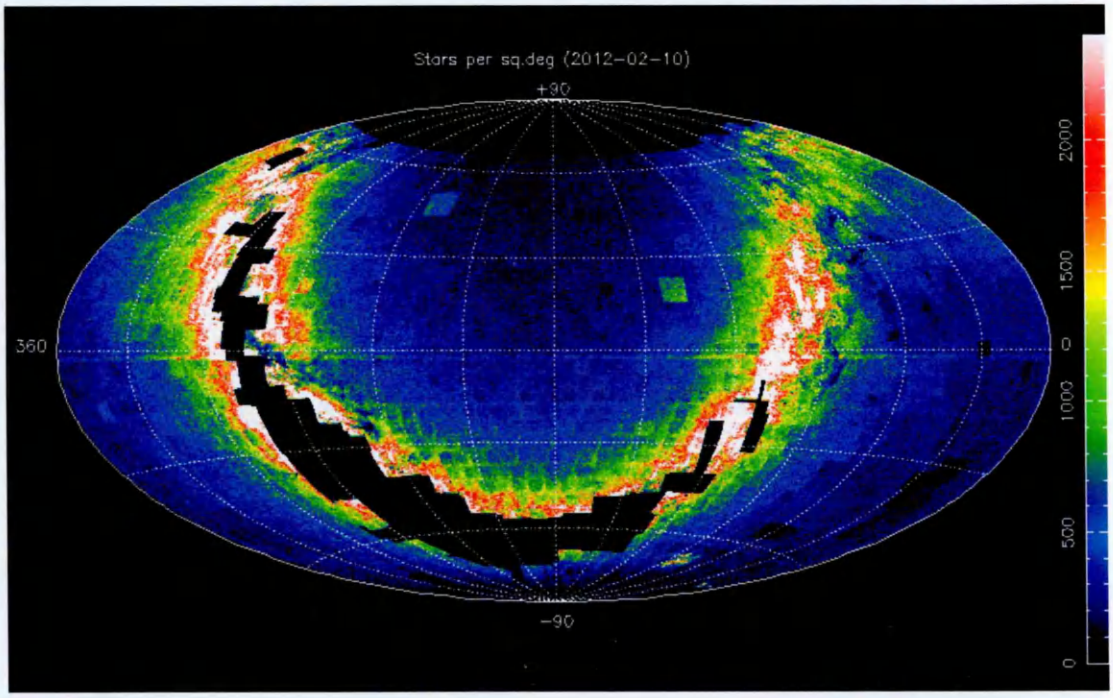

Figure 4.3: Sky map derived from SuperWASP data showing the number of stars recorded per square degree as a function of position on the sky.

SuperWASP field, there are around 20,000 stars brighter than V ~13. A depiction of the field density for a given part of the sky (as taken from SuperWASP data) can be seen in Fig. 4.3. For a given declination angle, the mount steps through RA in 1 hour increments centred on the current Local Sidereal Time (LST). Each individual field exposure takes $\sim$ 1 minute (30s exposure time, 4s read-out time, and time taken to slew to next field). In a night, 8 fields are observed, so each field is repeatedly observed with a cadence of 8 minutes. In the initial 2004 observing season ($\sim$ 150 nights), over 6.7 million stars were detected across all the observed fields. As of February 2012, over 30 million unique objects have been observed, with light curves containing over 320 billion data points, built from over 7 million science frames. It is readily apparent that the systems used to manage such a dataset, and perform meaningful time-series analyses on the data, must be extremely rigorous and robust.

4.1.4 Pipeline

The pipeline used to process the vast quantity of SuperWASP data performs a range of tasks, from identifying and processing calibration frames, calibrating the science frames with the pre-prepared calibration frames, adjusting for various effects com-

mon only to wide-field systems, performing aperture photometry on the thousands of sources per frame, to producing final light curves after a detrending stage, ready for the transit search algorithm to go to work on them. These stages are detailed further in (Pollacco et al., 2006), (Collier Cameron et al., 2006), and (Collier Cameron et al., 2007) though a brief summary is included here.

4.1.4.1 Calibration frames

Bias, darks and sky flats are taken on each night of operation, and subjected to statistical tests to remove obvious anomalous frames. Master calibration frames are generated by iteratively sigma-clipping about the mean pixel values for the batch of calibration frame. The master bias frame is subtracted from the science frames and dark and flat-field frames, and any residual temporal drifts in the bias pedestal level are accounted for using the overscan region of the chip. The master dark frame is scaled to the required exposure time, and subtracted from the flat field and science frames. Any sky gradient in the sky flats is able to be removed, as the flat fields show a radially symmetric vignetting pattern. The flat fields are rotated 180 deg. about the centre of the vignetting pattern, the rotated image is subtracted, and a planar least squares fit to any residual gradient is removed. A shutter map is constructed using the same method applied to the PIRATE flat field (see Chapter 2 - in which it is determined to exhibit no shutter effect variation across the frame), and removed from the flat fields and science frames. Finally, the flat fields are exponentially down-weighted with increasing time from the point of calibration, so that more recent flat field frames dominate in the sigma-clipped mean.

4.1.4.2 Astrometry and catalogue generation

Once calibrated, each image is processed with SExtractor to find all sources 4σ above the sky background level. These detected sources are used compared with the Tycho-2 (Høg et al., 2000) and USNO-B1.0 (Monet et al., 2003) catalogs in conjunction with prior knowledge of the approximate image centre (from the mount's pointing information) to plate solve the science frames (typically with an RMS in the astrometry of 0.2 pixels or 2.74"). Once plate solved, all objects in the USNO-B1.0 catalog brighter than R =15 within the FOV of the frame are fed in as coordinates for performing aperture photometry, by converting from RA & Dec to image co-ordinates in pixel number via the prior plate-solve.

4.1.4.3 Aperture photometry and PPWASP

The known positions of the USNO-B1.0 catalog stars and their catalog magnitudes are used to create a magnitude-dependent exclusion mask around all sources. The remaining pixels are then iteratively fit with a 2-D quadratic function to remove the sky background. On the second iteration, outliers are removed and flagged as bad pixels for inclusion in a bad pixel mask. Automatic image rejection occurs at this point, with images being removed if they suffer more than 50% of the image being clipped as outliers, or if the χ^2 value of the fit is too high - typical in a source-less high sky background image - indicative of cloud cover. The fit is subtracted from the image, and aperture photometry with aperture radii of 2.5, 3.5, and 4.5 pixels is performed at each catalog position. The sky background is calculated from an outer annulus between 13 and 17 pixels radius from the aperture centre. After performing the aperture photometry, and before being handed over to the PPWASP (PP: 'Post Pipeline') module, each object is labelled with its airmass, catalogue magnitude, and heliocentric time, as both airmass and heliocentric time vary significantly over the FOV of the instrument; so this is done on a per-object basis. Primary and secondary extinction coefficients are determined from a least-squares fit to the variation of raw magntiude with air mass through the night, and the instrumental color response and system zero point are determined using 100 bright, non-varying stars as secondary standards, once all stars have been corrected to a constant air mass. This process is repeated 3 times for each of the measured fluxes from each aperture radius.

4.1.5 Candidate selection

The SuperWASP archive employs various algorithms that perform three main functions:

- Remove the effects of correlated 'red' noise from the light curves
- Instigate a 'Box Least Squares' transit search to look for transit-like periodic dips in the light curves, and provide initial best parameter estimates for the potential star-planet system
- Perform Markov Chain Monte Carlo (MCMC) simulations to provide more detailed parameter estimations by fitting model light curves to the candidates surviving initial rejection thresholds.

4.1.5.1 Removal of systematic trends

Pont et al. (2006) showed that correlated 'red noise' had a large impact on the planetary yield from initial survey data. Early estimations of planetary yield from TEP surveys (e.g. Horne, 2001) incorporated only white, uncorrelated noise in their estimations of planetary yield from transit surveys, but the reality was very different. Light curves in transit surveys, and indeed PIRATE data, actually exhibit 'pink' noise - a combination of both independent, random noise (white) and co-variant, systematic noise (red) (see Fig. 4.4). The additional contribution of correlated noise poses two problems. Firstly, the quantity of noise in a survey light curve is underestimated if white noise alone is taken into account. This is especially true for bright stars, where the dominant white noise source (the photon noise) is greatly reduced as a fraction of the total signal. If no red noise were present, one would expect to easily detect shallow transits, thanks to the ever decreasing white noise component with increasing stellar brightness. Secondly, correlated noise can often act on timescales akin to that of an exoplanet transit duration (a few hours). Spurious transit-like signals can therefore be generated by the systematic reduction in detected flux from a star for a period of time typical of a planetary transit. Such spurious signals can then, in turn, be detected (i.e. as false positives) by a transit-hunting algorithm, especially if the causes of such systematic noise happen to be periodic.

To minimise the effects of systematic, correlated noise on the transit-hunting capabilities of the SuperWASP project, the SysRem algorithm of Tamuz et al. (2005) is employed to remove a set of four basis functions, each denoted a_j where $\{j = 1, \ldots M\}$ and M is the number of measurements in a light curve, taken from N stars, where $\{i = 1, \ldots N\}$. Each basis function a_j has its own coefficient, c_i. As the algorithm was generalised from attempts to determine atmospheric extinction coefficients in the absence of colour data, a simple basis function a_j to consider is the variation of airmass with each successive measurement in a light curve, the corresponding c_i is therefore the extinction coefficient. The aim of the algorithm is to minimise the global expression:

$$S^2 = \sum_{ij} \frac{(r_{ij} - c_i a_j)^2}{\sigma_{ij}^2}$$

where r_{ij} are residuals and represent the average-subtracted stellar magnitudes of the light curves in the dataset. Thus c_i represents the slope of the linear fit to a set of a residuals, given a_j as a proposed basis function. Note the summation over both indices in the general expression above. In order to practically minimise S^2 the

algorithm must iterate over both indices until convergence. Further detail on how this is achieved is contained within Tamuz et al. (2005).

The effect of red noise on the SuperWASP dataset in particular is discussed in detail in Smith (2006). Fig. 4.5 shows Fig. 1 from Smith (2006), which shows the before-and-after effect of decorrelation with SysRem on a SuperWASP dataset. The analysis in this figure is underpinned by the fact that, were the noise in a SuperWASP light curve composed entirely of white, purely random noise, the rms scatter of a boxcar-smoothed (i.e. moving averaged) light curve, σ_b, would be:

$$\sigma_b = {}^{\sigma_{lc}}/_{\sqrt{n}}$$

where σ_{lc} is the rms scatter of the full, non-averaged light curve, and n is the number of data points averaged in each boxcar. In Fig. 4.5, the green data points are the rms values of non-averaged light curves against V mag, the red data points are the rms values of boxcar-smoothed light curves, and the blue data points are the rms values of the full, non-averaged light curves divided by $\sqrt{n}$. The top part of the figure shows the data before decorrelation with SysRem, the bottom part of the figure shows the rms values after decorrelation with SysRem. Three things are immediately apparent from this figure. Firstly, the improvement made from the SysRem decorrelation is apparent. Secondly, whilst SysRem clearly has a positive effect, if it were able to fully remove all of the correlated noise from the light curves, the red data points would overlay the blue in the lower plot. This is clearly not the case, so red noise still remains. The third point of note is the flattening of the red curve towards the brighter magnitude values, at an rms of $\sim$ 2.5mmag. This demonstrates the proportionally greater contribution of red noise to the overall noise profile of bright stars versus faint stars, and provides a lower noise limit due to red noise applied to all stars in a field. Despite the improvement of the data brought about by the SysRem algorithm, the remaining red noise does have an impact on the detection rate in the SuperWASP data, Smith (2006) concludes. Modelling in that paper reveals an expected linear increase with number of nights observed in the number of planets with transit SNR > 10. In other words, the effect of residual red noise in the light curves is to increase the required observation time to detect planetary transits.

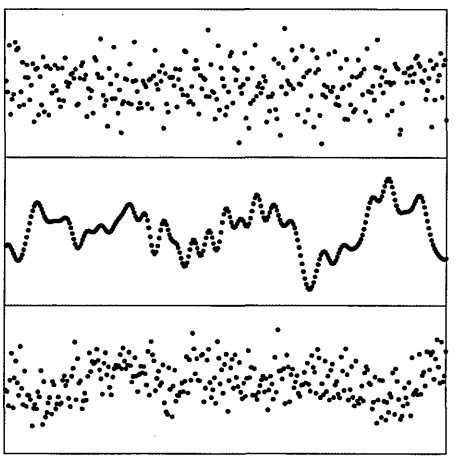

Figure 4.4: Fig. 2 from Pont et al. (2006). Illustration of light curve noise components. The upper light curve suffers from white (uncorrelated) noise only. The middle light curve contains only the red (correlated) noise. The bottom light curve suffers from pink (white and red combined) noise, and is representative of a typical transit survey light curve.

4.1.5.2 The HUNTER algorithm

After the light curves produced by the SuperWASP data reduction pipeline have been produced and decorrelated after the initial coarse decorrelation and subsequent decorrelation with SysREM, the hunt for periodic transit signal begins. The method by which this is done is described in detail in Collier Cameron et al. (2006). It is not necessary to repeat such detail here, and thus a brief summary is included in this thesis.

A Box Least Squares (BLS) algorithm akin to that described in Kovács et al. (2002) is employed for the transit search. At its most basic, this involves modelling the transit light curve with a box function described only by the flux level for points out of transit, the flux level for points in-transit (i.e. transit depth), the transit width, the period and the transit epoch. A coarse search grid of frequencies and transit epochs is set up and the data is phase folded according to each value of P and T_0 that constitute the search grid. For each position within the search grid, the best fitting in-transit flux value and the fraction of the total period spent in-transit are minimised through χ^2 minimisation. For each frequency, the best fitting transit depth is stored, along with the difference in χ^2 between the best fit box model and that of a constant flux model ($\Delta\chi^2$). A periodogram may be generated at this point, plotting $\Delta\chi^2$ against frequency; essentially a power spectrum for the range of frequencies tested. Also stored is the goodness of fit of a constant flux model to the

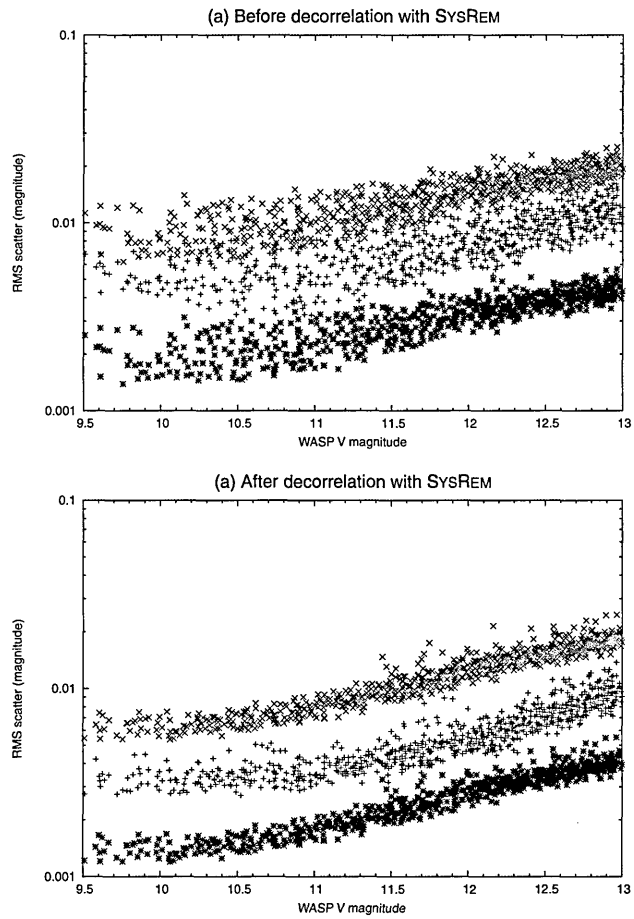

Figure 4.5: Fig. 1 from Smith (2006). Plots of light curve rms scatter against WASP V magnitude for a given field. The upper graph comprises data from light curves without decorrelation from SysRem; the lower graph contains data *after* decorrelation with with SysRem. Both plots contain the rms scatter of the whole light curve (green), of a boxcar-smoothed moving average with box size 2.5h (usually 20 data points - red), and of the full, non-smoothed light curve, divided by the root of the number of data points in a boxcar (i.e. $\sqrt{20}$ - blue).

determined out-of-transit data points. This is used later in the search for potential ellipsoidal variation in the light curve. Ellipsoidal variation arises in the case of a close binary system where gravitational distortion of each star upon the other breaks their spherical symmetry. The cross-sectional area of each star visible to an observer is thus a function of phase, which can be detected in the out of transit points by looking for sinusoidal variation.

With the initial grid search complete, a series of cuts are made to reduce the list of candidate systems. Firstly, all stars for which the post-fit $\chi^2 > 3.5N$, where N is the number of observations, are removed as variable stars. Secondly, any of the best fit solutions that exhibit gaps greater than 2.5 times the expected transit duration (after phase folding) are also removed. Further cuts are then made on two new light curve statistics.

The first of these is the signal-to-red-noise ratio, $S_{\rm red}$. This is defined in Collier Cameron et al. (2006) as:

$$S_{\rm red} = \frac{\Delta F \sqrt{N_{\rm t}}}{\sigma L^b} \tag{4.1}$$

where ΔF is the transit depth, N_t is the number of transits observed, σ is the RMS scatter of the unbinned light curve, L is the average number of data points spanning a single transit, and b is the power-law index that quantifies the covariance structure of the correlated noise. For light curves unaffected by correlated noise, $b = -1/2$.

The second statistic is the anti-transit ratio, $\left(\Delta\chi^2/\Delta\chi^2_-\right)$. This is defined as the ratio of the strongest peak in the aforementioned periodogram attributable to a periodic dimming, to the strongest peak attributable to a periodic brightening. For both statistics the thresholds are set to $S_{\rm red} < -5$ and $\left(\Delta\chi^2/\Delta\chi^2_-\right) > 1.5$. This selection of cuts typically removes 95.5-97.5% of all stars from a sample.

The remaining candidates are subject to a further BLS search, this time with a finer grid. The five most significant peaks in the periodogram are selected for each candidate. The transit parameters are then refined with a smoothed box like function; a function composed of hyperbolic tangents. This function (see Collier Cameron et al. (2006) for further detail) has the added benefit of being differentiable with respect to key transit parameters, and thus the iterative Newton-Raphson method may be used to refine transit parameters around the key periodogram peaks. At this point, the extent of any possible ellipsoidal variation in the out of transit light curve is determined, and a set of proposal parameters can be generated for further iteration by MCMC.

4.1.5.3 Parameter refinement by MCMC

The Markov-Chain Monte Carlo (see e.g. Tegmark et al. (2004)) method for parameter deduction and refinement has (over the last decade or so) become a vital tool, not just in astronomy, but in science in general. The perpetual increase in processor speed has dramatically reduced the run-times of such algorithms. Such 'simulations' are used to solve multivariate parameter-fitting problems, and serve two purposes: to optimize the fit of a proposed model to a data set; and secondly to 'wander' through all of the multi-dimensional parameter space in a bid to map the posterior probability distrubutions of the fitted model parameters. This second point is especially vital in situations where inter-parameter correlations may exist that lead to a range of degenerate solutions, where a large range of possible values exist for the correlated parameters. MCMC allows for these correlations to be mapped and understood.

The current iteration of the SuperWASP consortium MCMC code is that detailed in Enoch et al. (2010), which is an extension of the version described in Pollacco et al. (2008), itself an extension of that in Collier Cameron et al. (2007). Pollacco et al. (2008) describes the introduction of simultaneously modelling photometric and spectroscopic radial velocity measurements; but no radial velocity measurements are included in the work contained in this thesis, so this extension is not used. Instead the core photometric modelling process is described here, along with the change to the handling of stellar mass introduced in Enoch et al. (2010).

In the initial version of the code, six proposal parameters were employed: T_0, the transit epoch; P, the orbital period; ΔF, the transit depth; t_T, the transit duration; b, the impact parameter; and M_*, the stellar mass. In the updated version, stellar mass was removed from the set of proposal parameters and replaced with $T_{\rm eff}$, the effective temperature of the star; and [Fe/H], the metallicity of the star. More on this change later. Of the initial set of six proposal parameters, T_0, P, ΔF, and $t_{\rm T}$ are measurable directly from the light curve; indeed, the values used to initialise the respective proposal parameters are taken directly from those obtained in the BLS analysis, along with their 1σ uncertainties. Of the remaining two parameters, b is initialised as 0.5, and given an uncertainty $\sigma_b = 0.05$; and M_* is determined from the 2MASS $J - H$ colours for the host star. A linear calibration between $J - H$ colour and $T_{\rm eff}$ was determined from a sample of 100, 000 Tycho-2 F, G, and K dwarf stars. This calibration is used to convert $J - H$ into $T_{\rm eff}$. From there, $T_{\rm eff}$ is converted to stellar radius using a polynomial fit to the temperature-radius relation for main-sequence stars. The stellar mass is then estimated using the main-sequence mass-radius relationship ${}^{M_*}/_{M_\odot} \simeq \left({}^{R_*}/_{R_\odot}\right)^{1.25}$.

With the proposal parameters and their 1σ uncertainties initialised the simulation is allowed to evolve each parameter by taking a random walk through multi-dimensional parameter space, constrained according to the Metropolis-Hastings rule. This is performed as such:

1. Starting from the initialised values of the proposal parameters, a trial state is adopted for each proposal parameter, which we generalise through notation to be x_i, referencing the i^{th} step of any of the proposal parameters. A new trial state is adopted as thus:

$$x_i = x_{i-1} + \sigma_x G(0,1) f$$

   where σ_x is the 1σ uncertainty of the proposal parameter, $G(0,1)$ is a random Gaussian deviate with mean zero and a standard deviation 1. f is scale factor that acts as an adaptive step-size controller, and can be modified throughout the duration of the MCMC simulation to affect its performance.

2. For the trial state composed of the six new values for the proposal parameters, the physical parameters R_*, R_p, a, and $\cos i$ are derived from ΔF, t_{T}, b, and P using the generalised equations of Seager and Mallén-Ornelas (2003) detailed in Appendix A of Collier Cameron et al. (2007). The separation of centres of the planet and star as a function of t_j, the time of observation, can then be derived in terms of R_*, b, a, and ϕ_j, the orbital phase angle at time t_j, where $\phi_j = 2\pi\left(t_j - T_0\right)/P$. The flux deficit at all observed orbital phases can then be computed using the algorithm of Mandel and Agol (2002), to produce the model light curve for the current trial state. The appropriate limb darkening coefficients for the four coefficient nonlinear limb-darkening model of Claret (2000) (see Table 1.2) are adopted according to the photometric band of the observation. For SuperWASP data, the unfiltered photometry approximates to the R band.

3. For each trial state the fitting statistic Q_i is calculated for the data fit to the light curve model produced in the previous step. In the absence of any applied Bayesian priors, the fitting statistic is simply a slightly modified χ^2 statistic:

$$\chi_i^2 = \sum_j \frac{\left(m_j - \mu_j - \Delta m\right)^2}{\sigma_j^2} \tag{4.2}$$

where m_j are the observed light curve magnitudes, μ_j are the generated model values, σ_j are the measured magnitude uncertainties, and Δm is the modification from the standard χ^2 statistic and adjusts for the optimally averaged magnitude level of the light curve - in other words the zero-point offset of the observed magnitudes, which is determined before the MCMC chain is initialised.

4. For the new trial state, Q_i is compared to Q_{i-1}. If $Q_i < Q_{i-1}$ the new trial state is accepted as the new (current) set of proposal parameters. If $Q_i > Q_{i-1}$, the new trial state is accepted as the new set of proposal parameters with probability $\exp(-\Delta Q/2)$, where $\Delta Q = Q_i - Q_{i-1}$. This is the Metropolis-Hastings rule. It ensures that the global parameter space is more successfully explored by sometimes accepting 'worse' trial states. This allows for a global solution to be found and allows the algorithm to step out of a localised minimum it may otherwise get trapped in.

5. The extent of parameter space traversed by the algorithm around a solution represents the posterior probability distribution for each parameter in the set, from which parameter uncertainties can be derived from determining the 15.87^{th} and the 84.13^{th} percentiles of the posterior probability distribution.

The previously mentioned step-size controller f is adjusted after every 100 steps to ensure the acceptance rate of all new trial states is equal to 0.25, the optimal value. All six proposal parameters usually converge to the optimal values in around 500 steps. This initial convergence is referred to as the 'burn-in' phase, and once achieved, the chain is cut and the algorithm is allowed to explore again after being re-initialised at parameter values of the last step in the burn-in phase chain. At the point of re-initialisation, the parameter uncertainties are adjusted to account for correlated noise in the light curve by rescaling them so that the reduced χ^2 statistic, i.e. (χ^2/number of degrees of freedom) is equal to unity. From this point on, the uncertainties themselves are allowed to evolve along with the step-size controller, being recalculated every 100 steps directly from the chain. Finally, the step in the chain with the best value of Q_i is used to refine the solution further using the AMOEBA downhill simplex algorithm Press et al. (1992).

The version described above was changed in 2009 in order to remove stellar mass as a proposal parameter; thus making the stellar mass a parameter derived from other proposal parameters. The modification is based on the work of Torres et al. (2010) which uses a calibration to determine accurate stellar radii and masses from measured

values for $T_{\rm eff}$, $\log g$, and [Fe/H]. Given that the stellar density can be measured from the transit light curve, a substitution can be made for $\log g$, replacing it with $\log \rho_*$. In this iteration of the MCMC code, $T_{\rm eff}$ and [Fe/H] are jump parameters, replacing M_*. More detail can be found in Enoch et al. (2010).

4.2 PIRATE follow-up of candidate planets

4.2.1 Why follow-up?

Whilst much can be done with the SuperWASP data set to attempt to find the candidates most likely to exhibit periodic dimmings indicative of a transiting exoplanet, the majority of detected candidates will in fact be imposters, or 'astrophysical false positives'. The nature of such false positives is discussed in detail in the following chapter; but put simply these astrophysical false positives are usually stellar binaries in configurations that mimic the periodic, shallow dimming that transit search experiments are looking for. These can either be binary stars whose orbital inclination is such that the stellar limbs 'graze', producing a shallow, 'V-shaped' transit, or they may be deep-eclipsing binary stars whose eclipse depth is diluted by third light. Lastly, they may also be transiting dwarf stars whose radii do not significantly differ from those of Hot Jupiters.

Due to the capabilities of the SuperWASP hardware, it is often impossible to discern an astrophysical mimic or imposter from a real planetary candidate with SuperWASP data alone. Wide-field ground based transit photometry is limited in each of angular resolution, photometric precision, and temporal resolution. These limiting factors contribute to the inability to remove false-positives from the best candidate lists produced by SuperWASP. Brown (2003) attempted to predict expected detection rates and false-alarm rates for Jovian size transiting planets based on prior knowledge of the binary star population. His most pertinent conclusion for ground based surveys is that astrophysical false positives should occur at a rate 12 times that of true planetary systems (see also Charbonneau et al., 2004). Remarkably, this prediction is very close to the conversion rate of best planetary candidates to confirmed planetary systems seen in the SuperWASP North experiment, which is currently 7%, or a false-alarm to confirmed planet ratio of 13.3 : 1.

In light of this success rate, follow-up photometry with small to medium class telescopes (such as PIRATE) is a necessity. If there was no opportunity to perform

such follow-up photometry, the next step available would be to perform reconnaissance spectroscopy on all of the best candidates. Such observations attempt to measure the reflex motion of the host star at the points of quadrature - the points of peak radial velocity variation - which can be inferred from the orbital period in the photometric light curve. The host star might exhibit a large radial velocity variation - indicating a stellar companion, or it might show a small radial velocity variation - indicating a possible planetary system, or it will show no variation whatsoever, indicating a blend scenario. Without follow-up photometry, over 90% of valuable large telescope ($\gtrsim$ 2m) time would be wasted on astrophysical false positives. Follow-up photometry with small to medium class telescopes can cut down on this wastage.

The ability to detect and eliminate false positives scales with the angular resolution, photometric precision and temporal resolution of the instrument. Follow-up photometric observations of candidate planets with instruments larger than SuperWASP can identify third light imposters contributing to the dilution of a deep-eclipsing binary system to that of something that mimics a planetary transit in the SuperWASP data set. The increased photometric precision and temporal resolution allows for false-positives to be rejected from the shape of the light curve alone, providing the ability to resolve an anomalously 'V-shaped' transit (i.e. low (t_f/t_T)) for a given transit depth.

Follow-up observations with PIRATE can therefore remove many of the imposters before they reach the spectroscopic follow-up phase, boosting the scientific yield of the larger, more expensive facilities.

4.2.2 Estimated rejection yield

With knowledge of the prevalence of astrophysical false positives, an obvious topic to turn to would be the assessment of how many mimics PIRATE could eliminate per year. What follows is a coarse assessment of the possible number of false positives that could be removed from the SuperWASP candidate lists, utilising a year's worth of weather data, knowledge of the daily daylight time and its variation throughout the year, the amount of observation time required to make a definitive classification of an object, and the amount of time in a year yielded to educational pursuits instead of research.

Firstly, we assume that the ephemeris for each observed system is correct and has a sufficiently small uncertainty that ensures the observation window covers the

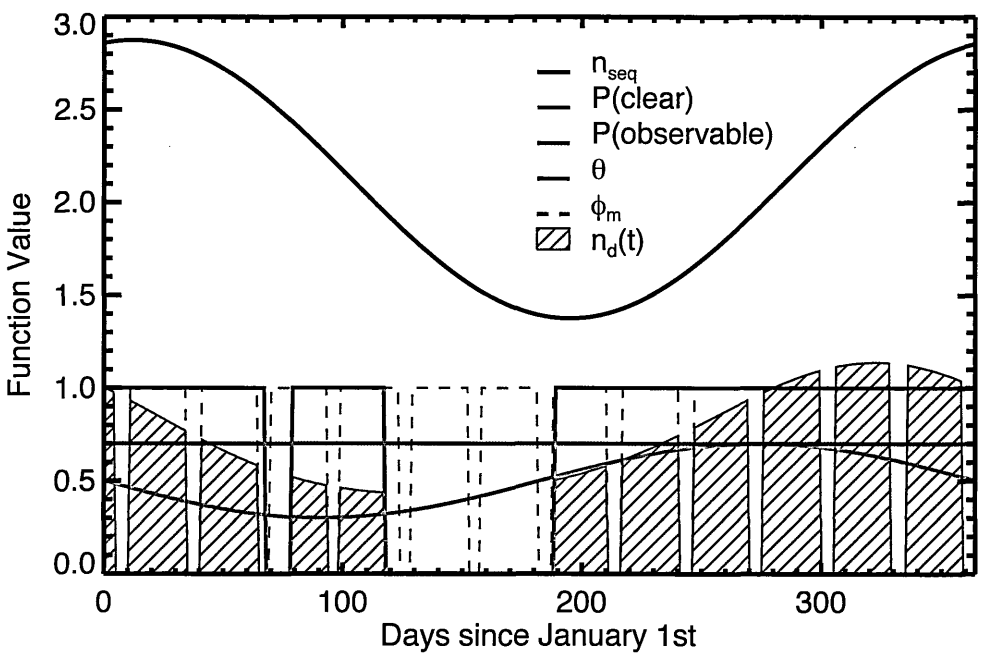

Figure 4.6: Variation of all the functions in Eq. 4.3. The black line is $n_{\rm seq}$, the green line is $P(\text{clear})$, the blue line is $P(\text{observable})$, the red line is θ, the availability function, the dashed purple line is ϕ_m, the moon phase function, and the shaded brown region is the resultant function $n_{\rm d}(t)$; the number of possible candidate observations per day.

event under investigation. If conditions allow for an observation to be made, then a rejection is counted.

To estimate the number of possible rejections of SuperWASP planetary candidates per year, we first define a function that estimates the number of possible candidate observations per day ($n_{\rm d}$) as a function of the day of the year, t:

$$n_{\rm d}(t) = n_{\rm seq} \times P(\text{clear}) \times P(\text{observable}) \times \theta \times \phi_m \tag{4.3}$$

where all the terms on the right side of the equation are also functions of the day of the year, t. Here, $n_{\rm seq}$ is the number of possible separate observation sequences possible on a given night in the year, which is restricted by the number of hours of daylight on a given day. This function is defined later in this subsection. $P(\text{clear})$ is the probability that a given night will be clear and free from clouds. For a given year, the seasonal variations in this quantity average out to about 0.5. $P(\text{observable})$ is the probability that conditions will be observable, and is governed by factors other than cloud, such as humidity and the strength of the wind. θ is the 'availability function' and adopts the value 1 or 0, depending on whether the telescope is available for research purposes. PIRATE is used annually for two separate Open University modules which occur between the same dates each year; for these dates this function is equal to zero. It is equal to unity at all other times. The last function, ϕ_m, is the moon phase function. This function is set equal to zero for all dates where the fraction of the moon illuminated is $> 90\%$, and equal to unity for all other dates.

As mentioned, $n_{\rm seq}$ is the number of possible observation sequences in a given night, which is restricted by the number of hours where the sky is sufficiently dark, and the amount of time required to typically make a successful observation. The following function is employed as an approximation:

$$n_{\rm seq} = \frac{12 + 3\sin\left(\frac{2\pi t}{365} + \frac{(2\pi \times 79)}{365}\right) - n_{\rm twi}}{T_{\rm dur}} \tag{4.4}$$

where t is again the number of days past since January 1$^{\rm st}$ in the calendar year, $n_{\rm twi}$ is the number of hours lost to twilight and preparation time at the beginning of the night just after sunset and at the end of the night before sunrise, $T_{\rm dur}$ is the amount of time taken (in hours) to perform a successful observation, and the phase term $\frac{(2\pi \times 79)}{365}$ provides the offset from the vernal equinox to the 1$^{\rm st}$ January. This function is not perfectly accurate, but it provides a sufficient approximation to the amount of

observation sequences that could be fitted into a given night's 'dark time'. For the purposes of this estimation, $n_{\mathrm{twi}} = 3.5$ hours, and $T_{\mathrm{dur}} = 4$ hours.

As we are treating t as a discrete series and not a continous one, the total number of possible observations of SuperWASP candidates is given by:

$$n_{\mathrm{total}} = \sum_{\mathrm{t=0}}^{365} n_{\mathrm{d}}(t) \tag{4.5}$$

Given the current functions input into the model, this yields a figure of 181 candidates per year. As of November 2011, there were around 400 A-Class candidates in the SuperWASP database (for each of SuperWASP North and South). New candidates are found at a rate of $\sim$150 per year. Clearly, PIRATE would be able to make a significant contribution to removing imposters from this dataset given the predicted number of observable candidates in a year.

4.2.3 Optimal target selection

The SuperWASP database pages provide a useful tool for determining which candidates will undergo a transit event in a given evening for a particular observatory site. This can therefore be used to create a schedule for a night's observations with PIRATE. In choosing a schedule, consideration must be given to:

- The positioning of the object in the sky throughout the night. What are its rise and set times, when does it cross the meridian, and to what elevation does it climb? This affects the minimum achievable airmass through which it can be observed. Furthermore, when is the transit expected to occur throughout its rising and setting? Will the target cross the meridian mid-way through the transit event, causing PIRATE to pier-flip?

- The brightness of the object. Due to the magnitude cut-off limit in the SuperWASP data set of $V = 14$, PIRATE follow-up observations are not troubled by faint SuperWASP candidates. However, an object can be too bright to be practical. Bright objects ($V \lesssim 10$) reduce the duty cycle of the observation (due to requiring less exposure time to avoid saturation) and may not have any nearby comparison stars of comparable brightness.

- The detected transit depth. Does the target exhibit a transit depth measurable with PIRATE in a single observation, or is it too shallow, thus requiring repeat observations? (See Chapter 5).

- The uncertainty on the ephemeris. Candidates detected in SuperWASP data that is a few years old could exhibit 'ephemeris drift' where the uncertainty in their ephemeris accumulates over time, reducing the likelihood that the transit event will occur in the chosen observation window.

- The utility of a PIRATE follow-up observation - has the target already been observed with a larger telescope? Could a different, unobserved target therefore be a more fitting inclusion in a follow-up schedule?

After asking such questions about the targets available on a given night, suitable targets are selected according to the time their expected transit event occurs, and slotted into either 2 or 3 observation 'bins' of around 3-4 hours, depending on the season and the length of dark time for a given night.

Chapter 5

SuperWASP Candidate Follow-up Observations - Results

5.1 Typical false-positive rejection scenarios and follow-up observation outcomes

Follow-up observations of SuperWASP planetary candidates with PIRATE typically fall into one of four categories: resolved blends, 'v-shaped' non-planetary eclipses, plausible planetary transits, and non-detections. Descriptions of these outcomes together with example case-studies of these outcome categories are detailed in sections 5.1.1 - 5.1.4. In section 5.2 a full overview of the follow-up programme's results (up to September 2011) is given, along with notable candidate observations from the programme worth expanding upon.

5.1.1 Resolved blends

SuperWASP's 13.5μ pixels and 200mm focal length for each camera combine to produce a plate scale of $13.7''\text{pixel}^{-1}$. This is many multiples of the typical astronomical seeing for any normal observatory site. For example, 'good' seeing at the OAM is anywhere between 1.5"-2". At La Palma (where SuperWASP-North is located), it will be better than this. In any SuperWASP image, the photons recorded in a pixel may come from many unresolved stellar sources. Aperture photometry is performed on all

locations defined as point sources in the USNO-B1 catalog (Monet et al., 2003), with radii of 2.5 (34.25), 3.5 (47.95), and 4.5 (61.65) pixels (arc seconds). The recorded flux within the measuring aperture is therefore an amalgam of all sources within the extensive apertures, unless it so happens that no other sources reside within the series of apertures. The intrinsic variability that any source may exhibit is therefore diluted by flux from background sources. Transit signals of planetary depth and duration may be mimicked by diluting the transit signal of an eclipsing binary star with other non-varying stars in the aperture. These astrophysical mimics are known as 'Blended Eclipsing Binaries' (BEBs), and are the most common source of false-positives in the SuperWASP dataset.

Given that the PIRATE plate scale, (which has been 1.26"/pixel, 1.69"/pixel, and 0.634"/pixel throughout the three operational phases detailed in Table 2.1) provides significantly better angular resolution than the SuperWASP instrument, most BEBs that are unresolvable by SuperWASP are resolvable by PIRATE. Observing throughout the anticipated transit event with PIRATE can reveal the true nature of the SuperWASP candidate. As an example, the target 1SWASPJ092349.39+503821.4 is detailed in the following section.

5.1.1.1 1SWASPJ092349.39+503821.4

This target was detected by the Hunter algorithm in multiple fields and cameras in data taken by the SuperWASP instrument between November 2007 and January 2009. In total 10466 separate observations of this target were made by SuperWASP. We'll look at the SYSREM-corrected dataset comprising all data points across the aforementioned fields and cameras. The primary peak of the Box Least Squares (BLS) algorithm used to search for a transit signal found a periodic signal with $P = 0.5868876$d and $T_0 = 2454491.6149$, where the epoch, T_0, is the HJD. Crucially, the depth of the box model fit to the SuperWASP data was 5.7 mmag, a relatively shallow transit. The J-H 2MASS colours and V-K Tycho-2 V_t and 2MASS K colours indicate an effective temperature of the primary star between $\sim$ 5200 – 5400 K, suggesting a near Solar-type G9-K0 spectral type. The transit depth would therefore suggest a relatively small planet, with radius $\sim 0.5 R_J$. To expand upon the initial assessment of the candidate's nature from SuperWASP data alone, the MCMC fitting routine of Enoch et al. (2010) was used to fit to the SYSREM-detrended SuperWASP data. The results of this fit can be seen in Table 5.1 and Fig. 5.1. As can be seen, the fit suggests a sub-Saturnian body orbiting a Solar-type primary. However, to be

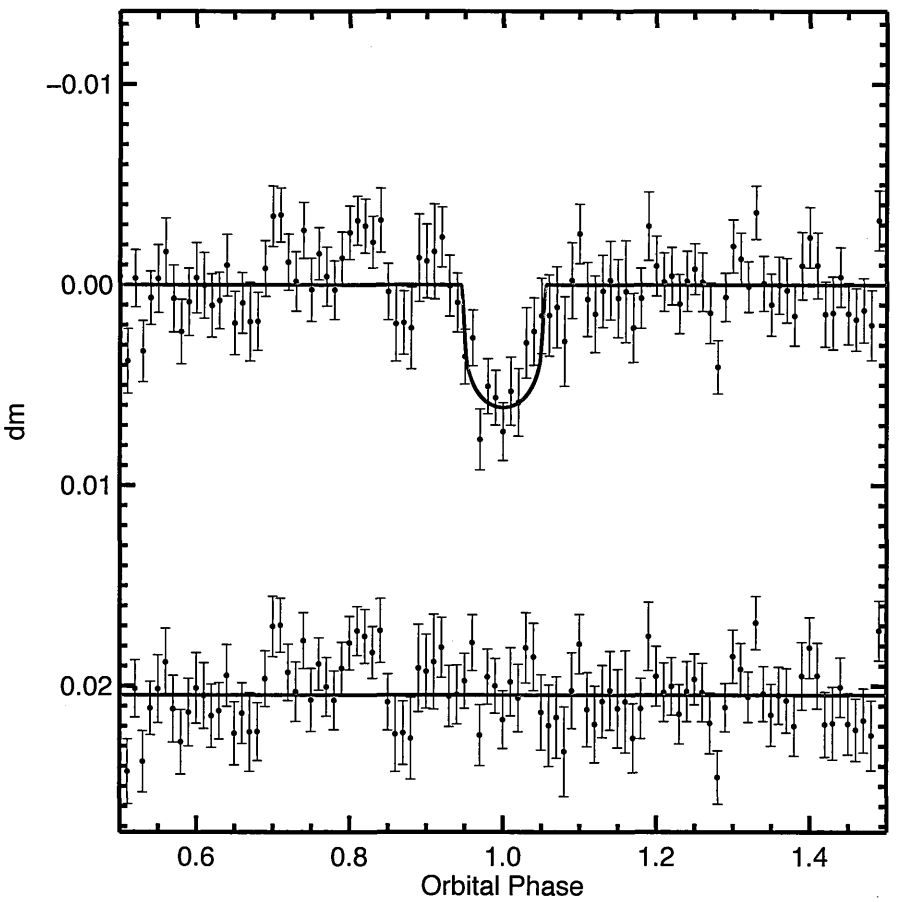

Figure 5.1: All of the SuperWASP (phase-folded) data for candidate 1SWASPJ092349.39+503821.4 after decorrelation with the Tamuz algorithm. The data is binned with bin size = 0.01ϕ. The model overlaid is the best fit model as found by the SuperWASP MCMC model fitting routine. Shown offset are the fit residuals.

confirmed as a planet, the candidate must first endure and pass a series of tests. The first of these is follow-up photometry; hence the target was subjected to PIRATE's gaze on 7th April, 2011.

The target was observed between 22:40:29 UTC (07/04/2011) and 03:35:40 UTC (08/04/2011), using 60s V band exposures. 228 good frames were used in the analysis. These data were passed through the PIRATE pipeline, undergoing the usual calibration and photometry procedures, before being fed into the light curve generation routine in order to produce the final light curves. In Fig. 5.2, a SuperWASP image of the target (left panel) is set next to a PIRATE image of the target (right panel). Two things are immediately evident: the relative size of the pixels, and the

Parameter	Value	1 σ (upper)	1 σ (lower)
Transit Epoch (T_0) (HJD)	4794.44152	0.00355	0.00141
Orbtal Period (days)	0.586873	0.00001	0.000005
$(R_p/R_*)^2$	0.0047	0.0003	0.0003
Transit Width (days)	0.0641	0.0024	0.0028
b (impact parameter)	0.26	0.05	0.03
Stellar mass ($M_\odot$)	0.961	0.025	0.024
Stellar radius ($R_\odot$)	0.966	0.041	0.045
Stellar density ($\rho_\odot$)	1.068	0.149	0.113
Orbital semi-major axis (a/AU)	0.0135	0.0001	0.0001
Orbital inclination (i)(°)	84.97	0.59	0.86
Planetary radius (R_p/R_J)	0.64	0.03	0.03

Table 5.1: Best fit parameters from the MCMC analysis of SuperWASP data for 1SWASPJ092349.39+503821.4

greater number of resolved stellar sources in the PIRATE image. We can clearly see that the candidate is not one single source, as in the SuperWASP image, but actually one bright source with two fainter, proximate neighbours.

Upon studying the output light curves of the main target ('B' in Fig. 5.3) and its faint companions, it is quickly evident that the faint companion SDSSJ092348.94+503835.7 ('C' in Fig. 5.3) undergoes a deep, 'V-shaped' (~25%) eclipse, a little later than the predicted transit midpoint of 2455659.4921. Cleary, this object is an eclipsing binary star. However, when its flux is added to the 'joint' Point Spread Function (PSF) in the SuperWASP image ('A' in Fig. 5.2), the periodic eclipse signal attributable to this object is diluted (mostly) by the bright, non-varying companion 1SWASPJ092349.39+503821.4. The effect of the dilution is that performing aperture photometry on this 'joint' PSF yields a light curve with periodic dip matching the orbital period of the binary star (0.586873 days), but with suppressed depth, approximately by a factor of 60. As a consequence of the PIRATE follow-up observations, the candidate was rejected from the candidate pool, avoiding unnecessary large telescope follow-up reconnaissance spectroscopy.

Resolving blend scenarios with telescopes such as PIRATE is arguably a very productive pursuit. Targets with significant blending (deduced from catalog matching) are logged on the target's Hunter page, along with DSS (Digitized Sky Survey) snapshots of the targets. Targets can then be flagged for follow-up with a specific view to look for blending scenarios. These candidates are the most productive to observe, as the 'usefulness' of the follow-up observation is maximised. Observation of such targets is more likely to yield a convincing rejection; and so biasing scheduling

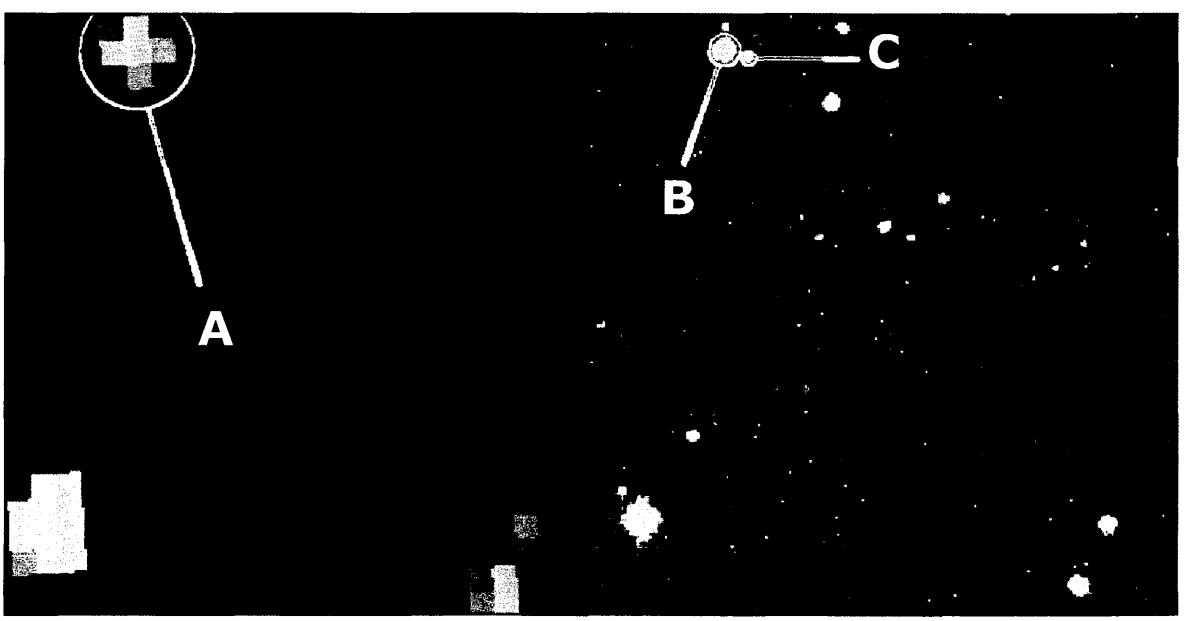

Figure 5.2: 1SWASPJ092349.39+503821.4 as seen by SuperWASP (left panel) and PIRATE (Mark 1.5) on 7th April 2011 (right panel). A: 1SWASPJ092349.39+503821.4 as seen by SuperWASP (plate scale 13.7"/pixel). It is indentifiable as a single source only. B: 1SWASPJ092349.39+503821.4 as seen by PIRATE Mark 1.5 (plate scale 1.69"/pixel). Clearly identifiable are two faint neighbours. C: SDSSJ092348.94+503835.7, a $g = 16.1$ mag eclipsing binary system.

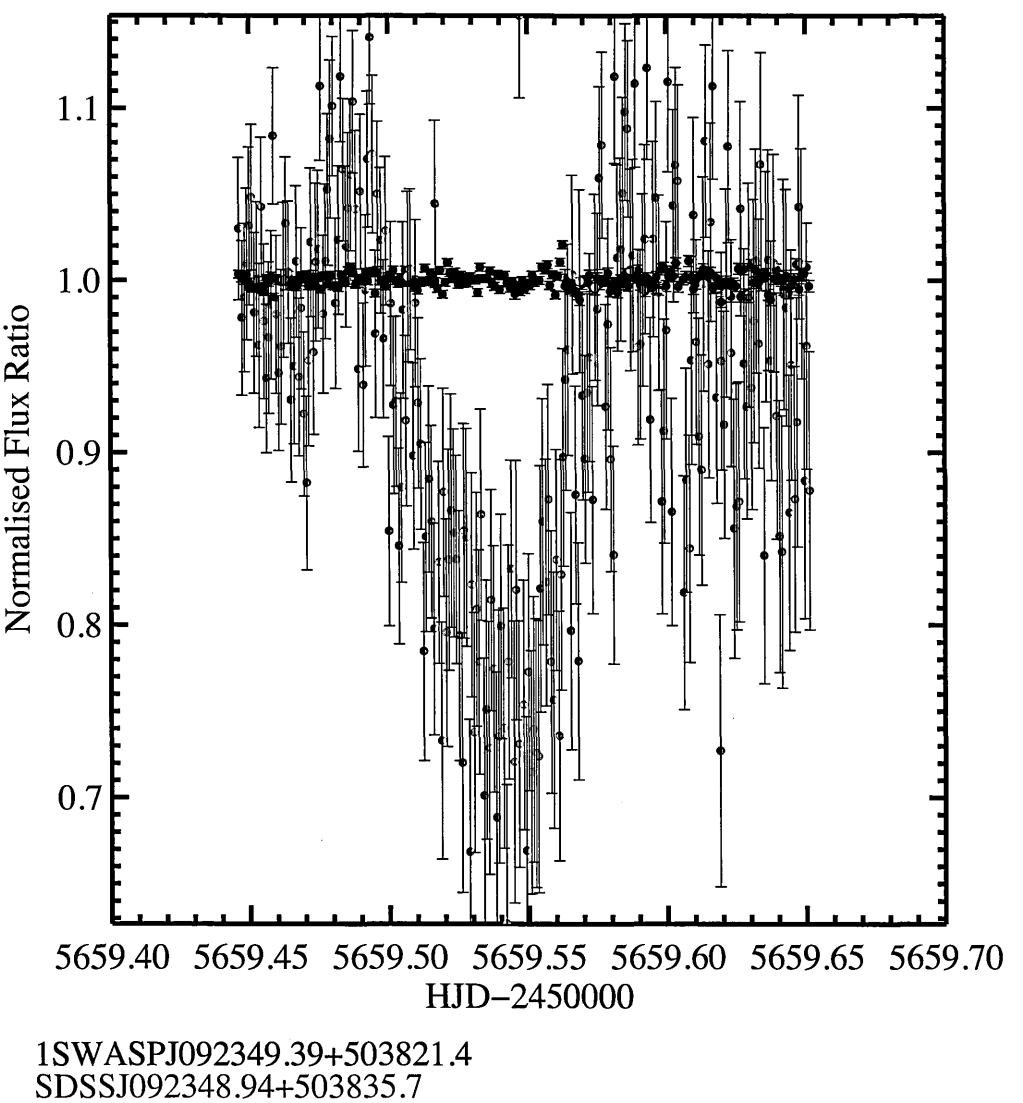

Figure 5.3: The light curve of 1SWASPJ092349.49+503821.4 (black data points - 'B' in Fig. 5.3) and that of its 16th mag companion, SDSSJ092348.94+503835.7 (red data points - 'C' in Fig. 5.3).

towards them allows for a greater fraction of false-positives to be removed from the candidate list per unit observing time.

To assist in the quick and easy identification of blend scenarios, an extra procedure was added to the PIRATE reduction pipeline which, once the final light curves have been generated, collates the light curves of detected sources within the radius of the largest SuperWASP measuring aperture (with radius equal to 4.5 SuperWASP pixels = 61.65") for easy inspection. The output of this procedure '`investigateblends`' can be seen in Figs. 5.4 & 5.5 for the BEB 1SWASPJ232332.12+522539.5, observed on 26/10/2010 in the R band by PIRATE Mark 1.5. As can be seen, the target (Star ID 894) actually has a faint neighbour that wasn't detected as a separate source in the PIRATE image. Fortunately, this faint neighbour was not the blend culprit; but the brighter ($m_V \sim 13.8$) star (Star ID 940 - 1SWASPJ232332.42+522624.2) at a separation of 45" clearly is.

5.1.2 'V-shaped' non-planetary eclipses

These eclipse types are characterised by a brief 'transit floor', with the combined duration of ingress and egress making up the majority of the total transit duration. In simpler terms, the transit shape approximates a letter 'V' as opposed to the typical 'U' shape of a planetary transit. Such transit shapes are attributable to either 'grazing' eclipsing binaries, or unresolved (by PIRATE) blended eclipsing binary stars.

5.1.2.1 Unresolved blended eclipsing binary stars

This situation is similar to BEBs recorded in the SuperWASP data set. Like the SuperWASP blend example, these blends are also unresolved, though this time they remain unresolved even with the superior angular resolution of PIRATE. Such systems may be chance alignments of background eclipsing binary stars with foreground single, non-varying stars, or even gravitationally bound systems, so called 'hierarchical triples', in which a binary system and single star orbit a common barycentre.

As mentioned, both systems (SuperWASP & PIRATE) are unable to resolve such objects into their constituent flux sources. However, the SuperWASP instrument and PIRATE have very different noise profiles. It is quickly apparent that PIRATE is still able to rule out more of these false-positives, simply by resolving the *shape* of the light curve to greater precision for a given apparent magnitude.

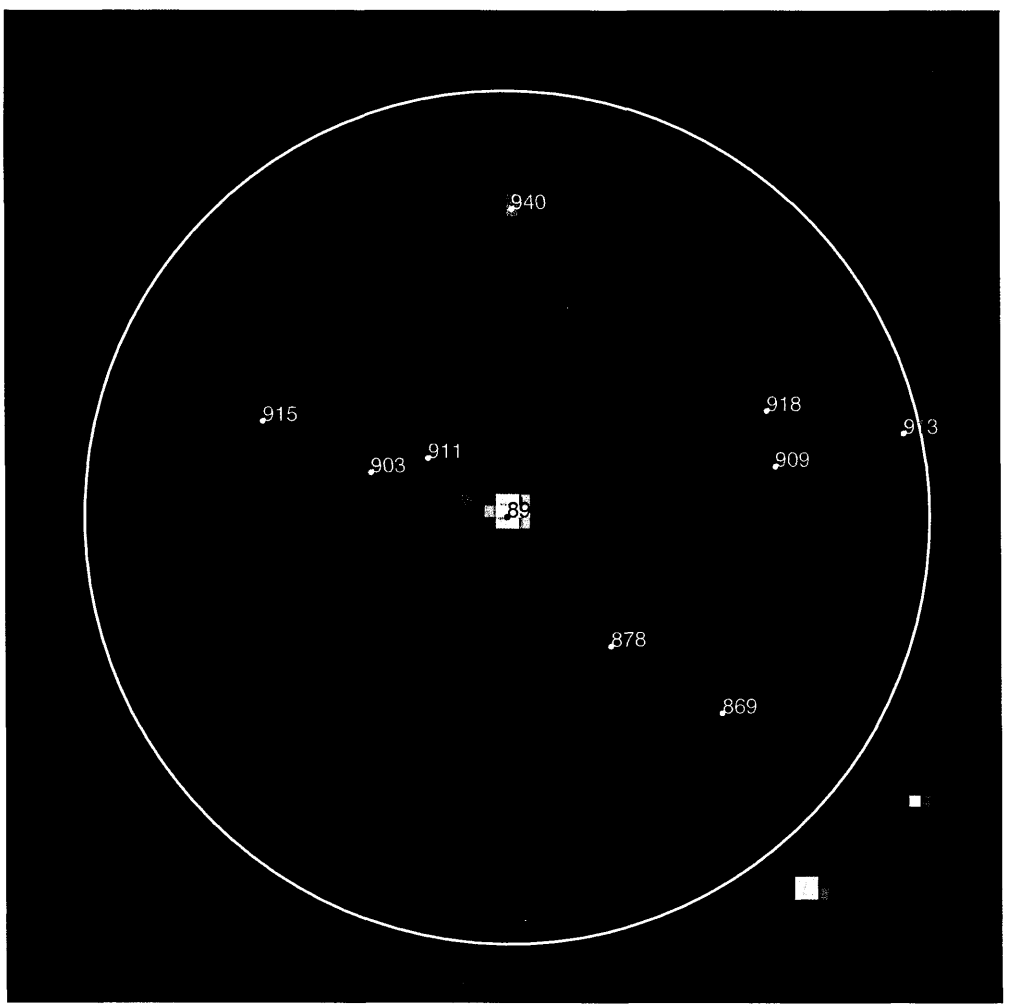

Figure 5.4: PIRATE image of the SuperWASP candidate 1SWASPJ232332.12+522539.5. The yellow circle represents the SuperWASP 4.5 pixel radius aperture, and each detected source in the PIRATE image within the aperture is labelled according to its Star ID in the pipeline reduction. The light curves of these objects can be seen in Fig. 5.5 .

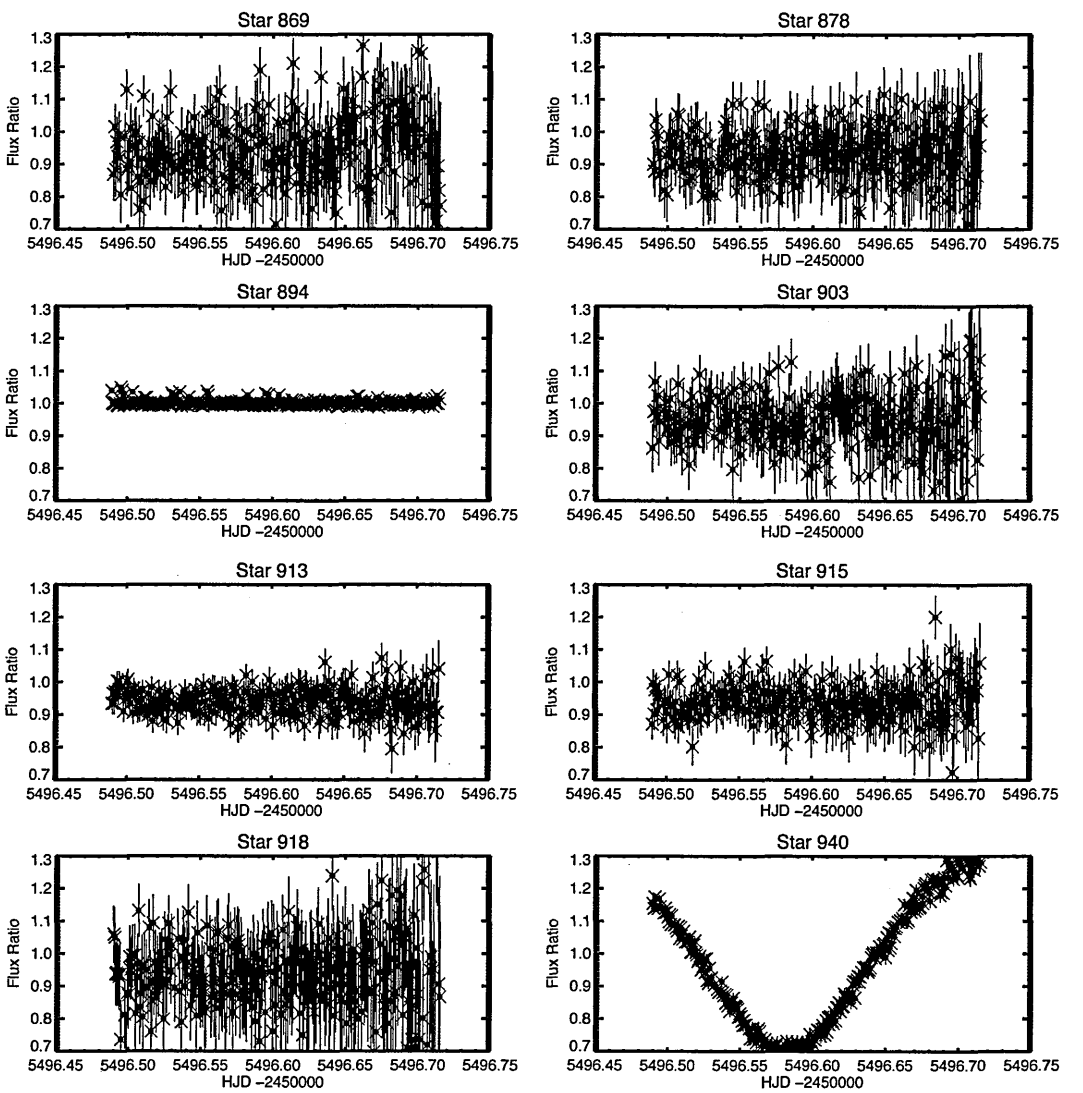

Figure 5.5: The light curves of a selection of the detected sources within 61.65" of the target 1SWASPJ232332.12+522539.5. The Star IDs for each plot are referenced in Fig. 5.4.

To demonstrate this, a selection of fake blend scenarios have been created, depicted in Fig. 5.6. The light curve of 1SWASPJ232332.42+522624.2 (the EB responsible for the blend scenario in Figs. 5.4 and 5.5), as observed with PIRATE was used as a starting point. The catalog magnitude of this object is $m_V = 13.8$. Five artificial light curves were generated at apparent magnitudes of 9, 9.5, 10, 11, and 14. A fractional uncertainty of 0.6% was assumed for the light curve of 1SWASPJ232332.42+522624.2. The noise properties for the artificial stars were then generated with this as a reference point (i.e. that a $m_V = 13.8$ star has a fractional uncertainty of 0.6%). The source of the noise was assumed to be photon noise alone, so the size of the 1-sigma photometric uncertainties for the artificial stars scale as if poissonian, but the noise itself was generated by drawing from a gaussian distribution, to give it a visual appearance more grounded in reality. To generate the combined light curve, the non-varying artificial stars are combined with the flux from the eclipsing binary, weighted by their respective flux ratio. In Fig. 5.6, we see the effect of dilution on the eclipse depth. In the case of 1SWASPJ232332.42+522624.2 being diluted by an artificial 9.5th mag star, we see a plausible planetary transit of depth ($^{\Delta F}/_{F}$) around 0.8%. The dilution masks the extent of the 'V' shape; one could plausibly fit a high impact parameter ($b > 0.5$) planetary ('U' shape) transit to this combined light curve. However, at lower levels of dilution (by fainter stars), the 'V' shape characteristic becomes more prominent and resolvable, thanks to the increased transit depth.

What is it about the shape that betrays the true eclipsing binary nature of the object? We can look at the analytical transit equations of Seager and Mallén-Ornelas (2003), specifically Eq. 5.1, which links t_F and t_T to the ratio of radii of the two bodies ($^{R_p}/_{R_*}$) and the impact parameter, b. Here, t_F and t_T are the duration of time between contacts 2 and 3, and 1 and 4, respectively:

$$\frac{\sin(t_F\pi/P)}{\sin(t_T\pi/P)} = \frac{([1-(R_p/R_*)]^2 - b^2)^{1/2}}{([1+(R_p/R_*)]^2 - b^2)^{1/2}} \tag{5.1}$$

If we assume ($^{t_T\pi}/_{P}$) $\ll 1$, we can use the small angle approximation, $\sin x \approx x$, and in turn remove the dependence on P. We can also substitute in for $(^{R_p}/_{R_*})^2 = \Delta F$:

$$\left(\frac{t_F}{t_T}\right)^2 = \frac{\left(1-\sqrt{\Delta F}\right)^2 - b^2}{\left(1+\sqrt{\Delta F}\right)^2 - b^2} \tag{5.2}$$

What does this tell us already about 'V' shaped eclipses? For any given value of

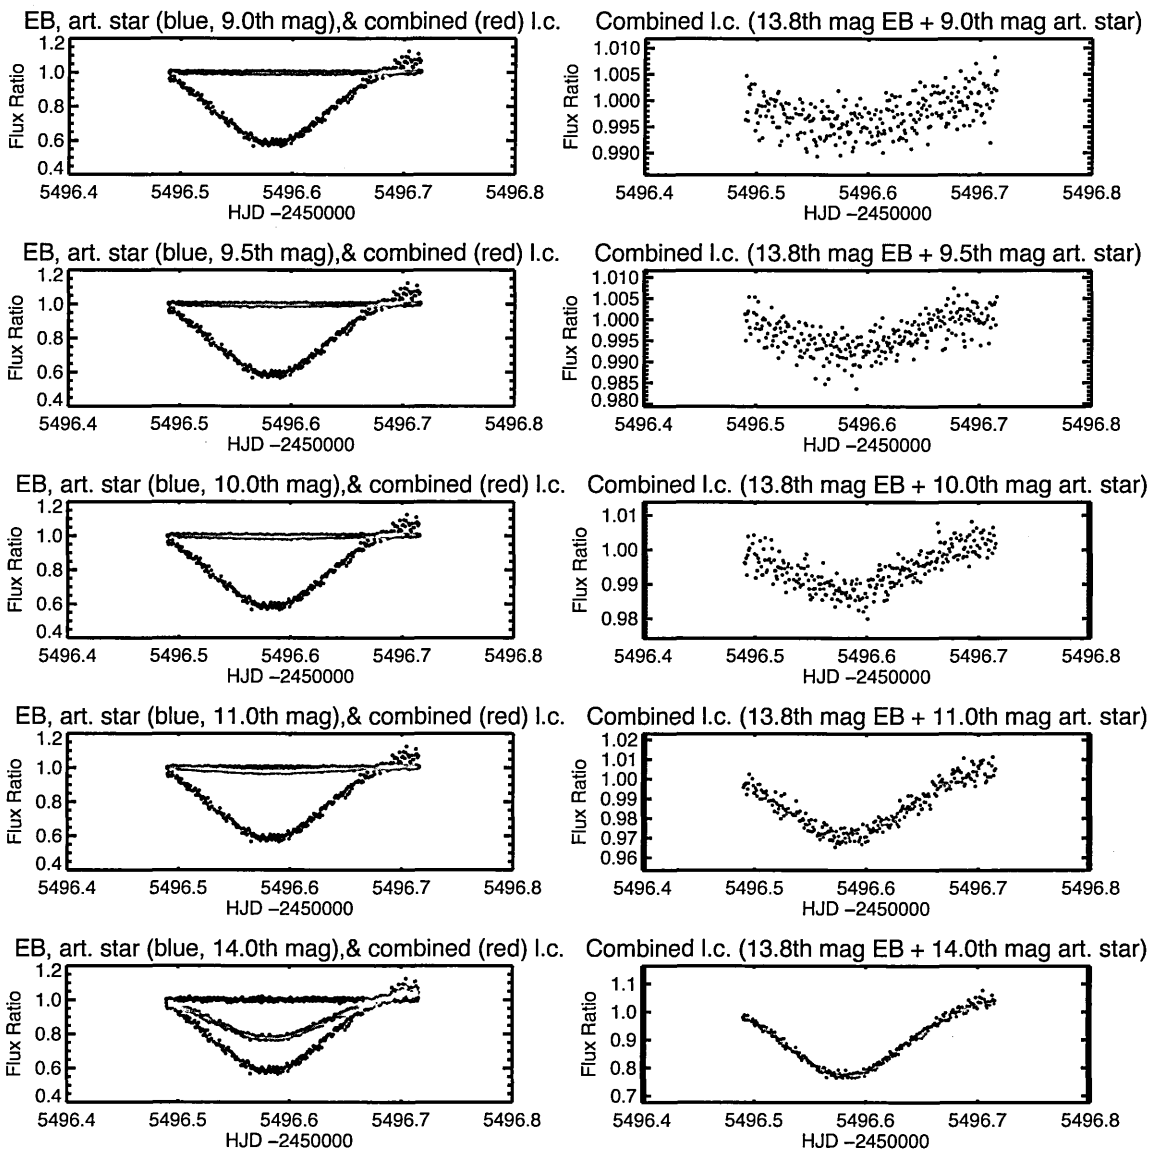

Figure 5.6: The light curve of 1SWASPJ232332.42+522624.2 (EB in Fig. 5.5) diluted by a range of (9^{th}, 9.5^{th}, 10^{th}, 11^{th}, 14^{th} mag) artificial stars. The left column shows the original input light curve (black, as recorded by PIRATE), the flux from the artificial source (blue), and the resultant combined (red) light curve. The right column shows just the combined flux (EB + artificial flux) with appropriate scaling to match the size of the apparent eclipse depth.

impact parameter, the lower the value of the ratio (t_F/t_T) (as we observe with the 'V' shape) the higher the value of ΔF, and hence the greater the size of the eclipsing body relative to the host star. Any transit light curve that exhibits a flux deficit ΔF that yields a secondary radius (in combination with an estimate of primary radius) suggestive of a planetary sized body coupled with a low value of (t_F/t_T) must necessarily have a high impact parameter. High values for the impact parameter necessarily reduce the ratio (t_T/P), as $t_T \rightarrow 0$ as $b \rightarrow 1$. If the light curve is of sufficient quality that the transit duration can be measured, together with the ratio (t_F/t_T) and ΔF, and the period is well known, it is possible to reject candidates of planetary transit depth on the basis of an excessive transit duration for the measured impact parameter. In effect, such a light curve breaks the assumption $(t_T\pi/P) \ll 1$, so Eq. 5.2 ceases to hold, and there ceases to be a unique solution for the stellar and planetary parameters.

Fig. 5.7 shows the dependence of the impact parameter upon the flux deficit for a range of values (0, 0.2, 0.4, 0.6, 0.8) of (t_F/t_T), produced by solving Equation 5.2 for b:

$$b^2 = 1 + \frac{2\sqrt{F}\left((t_F/t_T)^2 + 1\right)}{(t_F/t_T)^2 - 1} + \Delta F \tag{5.3}$$

It is readily apparent from Fig. 5.7 that for light curves exhibiting transit depths indicative of a transiting exoplanet (3% and less), and a 'V-shaped' transit ($t_F/t_T \lesssim 0.2$), only a small range of (high) impact parameter values are allowed; typically $b \gtrsim 0.8$. This already points to a low probability of the light curve being of planetary origin from orbital inclination arguments alone. The contra-planetary argument is typically bolstered further by measurements of (t_T/P).

5.1.2.2 Grazing eclipsing binary stars

If the orbital inclination of a binary system is such that a full-eclipse (where the full disc of the smaller transiting body is not encapsulated by the disc of the larger body) does not occur, but instead a partial eclipse occurs, a planetary-mimicking eclipse depth can occur. These are typically referred to as 'grazing eclipsing binary' systems. A unique solution does not exist to distinguish such systems from blended systems using photometry alone. Because of the nature of a partial eclipse, the eclipsed area of the primary star is never constant throughout the transit event, and so such eclipses are always 'V-shaped'. The same arguments apply for the elimination of

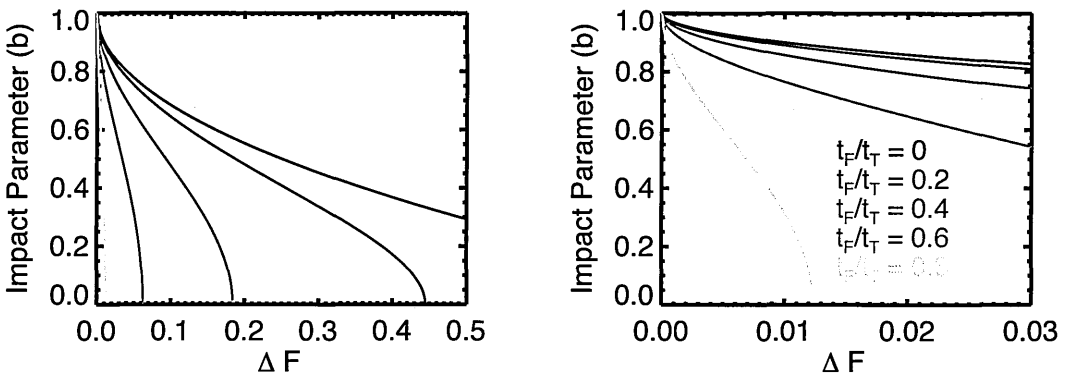

Figure 5.7: Dependence of b upon ΔF for a range of values of (t_F/t_T). The right hand plot is the of the same functions but with the x-axis scaled to values of ΔF typical for that of transiting extrasolar planets.

such targets from photometry alone, where data of a sufficient S/N ratio can reveal a transit duration in conjunction with a flux deficit, orbital period and ratio (t_F/t_T) that implies a non-planetary eclipse. If the eclipse depth is such that an object cannot be ruled out by photometry alone, and the object is not a deep-eclipsing binary diluted by the flux of a bright non-varying star, a spectroscopic follow-up would reveal a double-lined spectroscopic binary; where emission or absorption lines from stars of similar (or the same) spectral type are simultaneously shifted to the blue and red indicating one body moving away from the observer, and the other moving towards the observer.

5.1.3 Plausible planetary transits

These are, as the name suggests, follow-up observations which yield transit light curves for which plausible star-planet system parameters may be derived. These targets have no nearby blend candidates, and exhibit typically non-V-shaped transits (though some transits of planetary objects are particularly 'V-shaped', see for example TrES-3b (O'Donovan et al., 2007), the PIRATE light curve of which can be seen in Fig. 6.8). Such observations will fall into either of three categories after the initial follow-up photometry; they will either be rejected as blended binary systems or grazing binary systems via reconnaissance spectroscopy, be revealed to be the transit of early type main sequence star by a late type, low mass companion, or will eventually be revealed as true planetary mass systems from their radial velocity curve.

5.1.3.1 1SWASPJ002328.03-022914.5

This target was detected within the 2008 and 2009 seasons of SuperWASP data, with an ephemeris of:

$$T_{mid} = 2454945.893\ (\pm 0.002) + 3.50253E\ (\pm 0.00004)\ \mathrm{d}$$

and a transit depth of ~ 0.007 magnitudes. From the MCMC solution of SuperWASP data alone, a high impact parameter transit ($P(b < 0.8) = 0.15$ - where ($P(b < 0.8)$ is the probability that the impact parameter is less than 0.8) is detected with plausible stellar parameters ($M_* = 0.93M_\odot$, $R_* = 0.98R_\odot$), yielding a planetary radius of $0.8R_J$.

1SWASPJ002328.03-022914.5 was observed with PIRATE on 10/12/2010 (Mark 1.5) in the R band, with 295 exposures taken, each for 20s. A full transit was captured, with a transit depth consistent with that seen in the SuperWASP data. There are no nearby neighbours in the PIRATE images; and the light curve details a plausible planetary transit shape and depth. This object cannot be ruled out therefore by PIRATE follow-up photometry alone. The PIRATE data (binned to $\phi = 0.001$) and model fit from the combined MCMC analysis of SuperWASP & PIRATE R band data are shown in Fig. 5.8.

This object was subsequently ruled out by a spectroscopic analysis within the consortium using CORALIE (Queloz et al., 2000) mounted on the 1.2m telescope at La Silla, using the cross-correlation function (CCF) method as detailed in the original 51 Peg b paper (Mayor and Queloz, 1995) to find the radial velocity shift between a reference stellar spectrum template (in this case of a G2 star) and the observed spectrum. This analysis showed two distinct peaks in the CCF, the principal peak at $\sim 9\mathrm{kms}^{-1}$, and a broader secondary peak at $\sim -40\mathrm{kms}^{-1}$, so this target was classified as a double lined spectroscopic binary. As the light from two stars contributes to the spectrum, we can conclude that no third light contamination is present, and the plausible planetary transit actually derives from the grazing orbital inclination of a binary star system. Examples of actual planetary transits observed with PIRATE can be found in Chapter 6 for comparison.

5.1.4 Non-detections

Follow-up observations categorised as a 'Non-Detection' (ND) are observations in which no discernible signal is detected in either the target star or any nearby neigh-

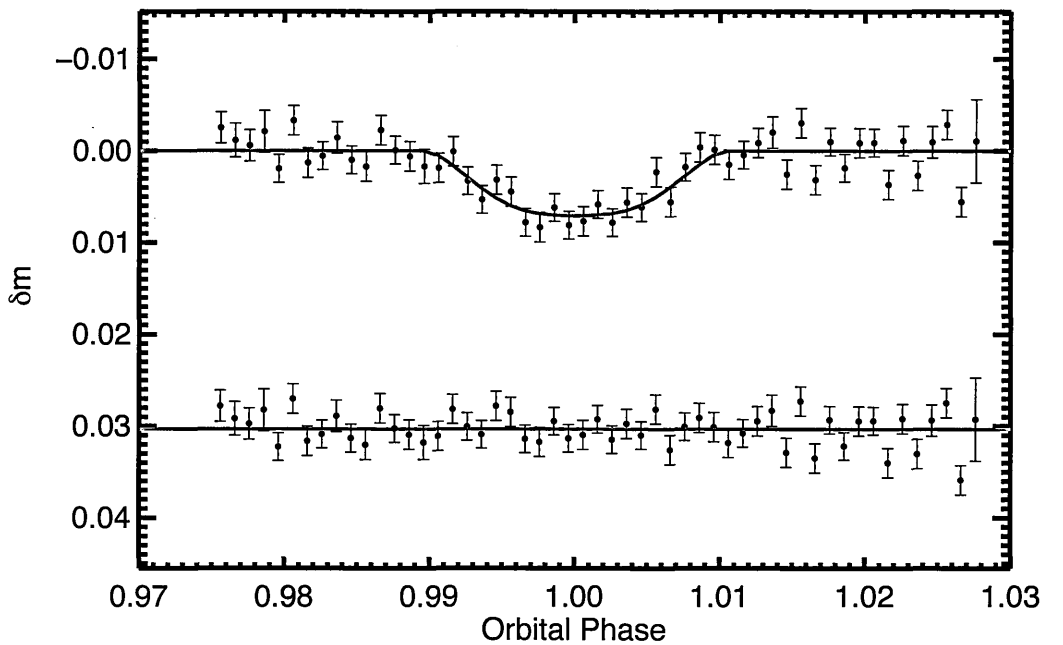

Figure 5.8: PIRATE light curve of the eclipse of 1SWASPJ002328.03-022914.5, R band. The data are binned to $\phi = 0.001$ (in this case $\sim$ 5 minutes).

bours during the observation window.

Such situations can arise for one of four reasons:

- The target does not show a 'real' periodic dip, but has been detected due to systematics in the detection dataset.
- The dip is 'real' (i.e. astrophysical in nature), but the ephemeris is poorly constrained, and hence the event possibly occurred outside the observation window.
- The dip is too shallow to be reliably detected in a single PIRATE observation. For example, a 3 mmag dip in a 12th mag star would only result in a 1σ 'detection'.

5.1.4.1 Non-astrophysical signal

The HUNTER algorithm (Collier Cameron et al., 2006) can often be fooled into detecting systematic noise with periodic modes. Most common of these are 'end-of-night' effects where targets heading for the horizon are subject to very high airmass observations at the end of an observation run, and thus suffer dramatically increased noise. The repetition of such events at similar times of the evening can install a fake

periodic signal in the target's light curve, which can subsequently be detected by the HUNTER algorithm. Such signals often occur therefore at integer-value periods, or at significant fractions of integer periods, a common period being 16 hours or 0.66 days. If the target is observed by PIRATE at a reasonable airmass ($z \leq 2$), it will not suffer any such trends, and most likely be revealed to be non-varying throughout the observation window, and a likely spurious candidate.

5.1.4.2 Ephemeris drift

In cases where the ephemeris is poorly constrained in the current epoch, either via very large uncertainities in the Period (P) and/or Epoch (T_0), or via smaller uncertainties that propagate and accumulate over long baselines, the likelihood of the transit event occurring within the observation window is significantly reduced. To demonstrate this, we look at the case of 1SWASPJ032259.**+******.* (SuperWASP ID blanked out as candidate nature still to be determined).

This object has a detected ephemeris of:

$$T_{\mathrm{mid}} = 2454346.60497\,(\pm 0.0061) + 1.2234940E\,(\pm 0.00018)\,\mathrm{d}$$

It was observed on 22/01/2012 with PIRATE Mark 2 (between HJD 2455949.302 - 2455949.501), and so it does not formally fall amongst the data discussed in this thesis, but is a perfect example of ephemeris drift, and so is included by exception. The uncertainties on the period and initial epoch are not in themselves particularly excessive ($\sim$15s & $\sim$ 9 minutes respectively, c.f. 8 minute cadence of SuperWASP), however, the last data taken by SuperWASP on this target was in 2007. Consequently, the error has accrued over the last 4-5 years to produce a significant ephemeris drift, approximately ± 0.24 d at the time of the PIRATE observation. No transit signal was detected in the target star or nearby companions, suggesting the event's timing had 'drifted' outside the PIRATE observation window.

Fig. 5.9 provides a visual depiction of the ephemeris drift by way of an O-C diagram. For a description of the figure see the caption. It is readily apparent how the ephemeris has very likely led to PIRATE detecting no transit signal during the observation window. For direct comparison, Fig. 5.10 depicts the same plot but for an observation of WASP-12b on 13/01/2011, as part of the ephemeris refinement study (see Chapter 6). The ephemeris is tightly constrained ($P = 1.09142189 \pm$

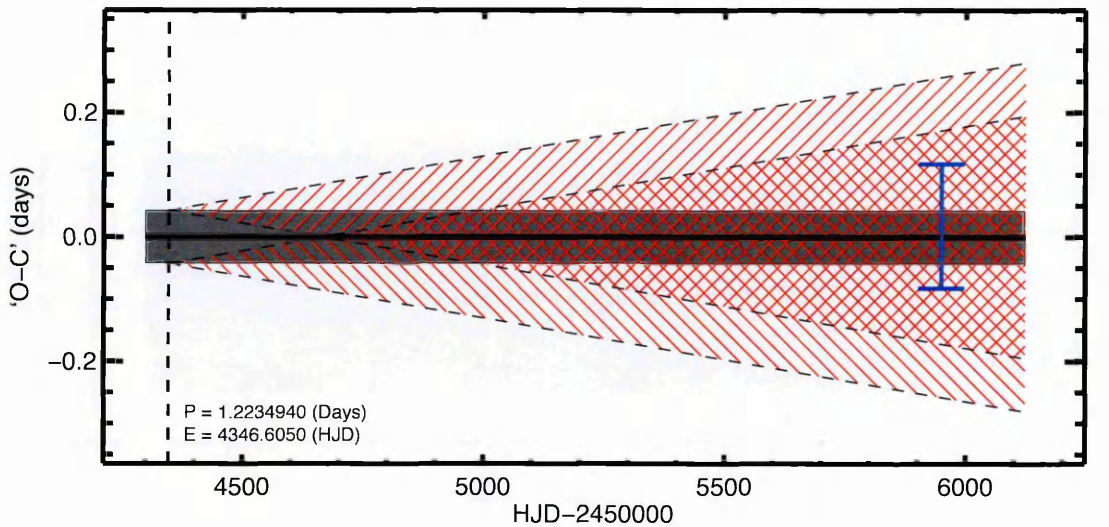

Figure 5.9: Visual depiction of the ephemeris drift of 1SWASPJ032259.**+******.* using an O-C diagram. The dark grey shaded area represents the width of the transit either side of (O-C = 0), the predicted transit midpoint. The width is assumed to remain constant. The vertical dashed line represents the initial epoch, 2454346.6050 (HJD). The shaded red regions represent the accumulated uncertainty through time for both the start and the end of the transit, hence they emanate from both the top and the bottom of the grey shaded area at T_0. The blue marker indicates the PIRATE observation window, bounded by the horizontal 'hats'.

0.00000044 d, $T_0 = 2454826.57975 \pm 0.00017$), and so the shaded red zones that represent the propagating uncertainty in transit timing are negligable. Consequently, as can be determined from Fig. 5.10, the transit mid-point and egress were indeed observed, see Fig. 5.11.

5.1.4.3 Shallow transits

For targets that exhibit small transit signals of depth 0.1-0.3% a single PIRATE observation will prove inconclusive, and most likely be classified as an ND ('Non Detection'). It is typical for any given 12^{th} magnitude star to display photometric scatter of between 0.2% and 0.3% (let us take 0.3%) for a typical operating cadence for PIRATE. Ignoring the nature of the contributing noise sources, and assuming these noise sources are uncorrelated (an often invalid assumption but one which doesn't break the validity of this illustration), a transit must exhibit a depth of 0.9% to register as a 3σ detection, a commonly accepted minimum detection criterion. With 0.3% scatter in our target light curve, a transit signal of depth 0.3% is statistically

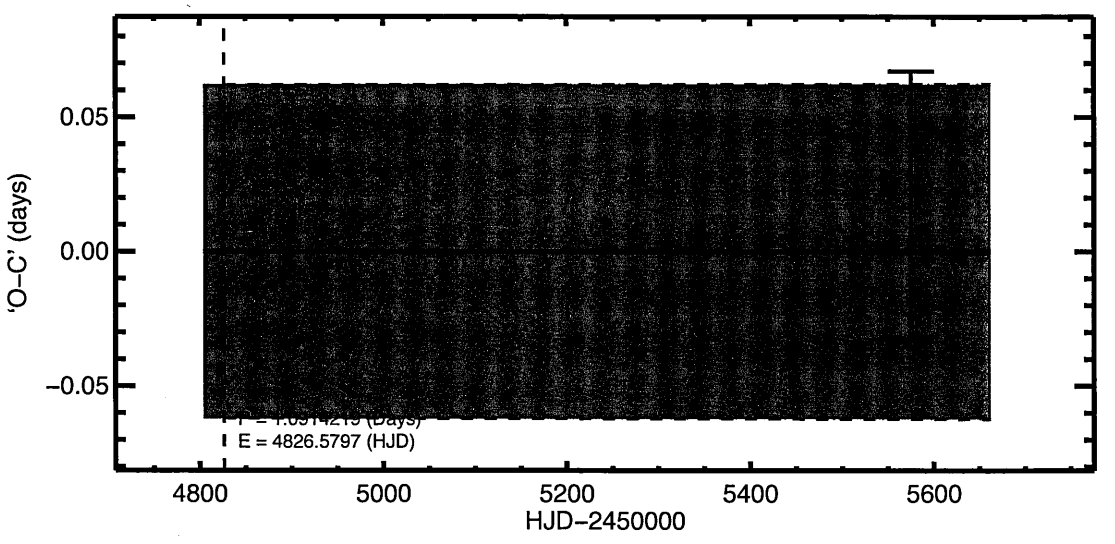

Figure 5.10: The equivalent diagram to Fig. 5.9, but for an observation of the transit of WASP-12b on 13/01/2011. The shaded red zone bound by black dashed lines, indicating the propagating uncertainty in the ephemeris from T_0 onwards, is barely visible, as the uncertainty is so small. Therefore the black dashed lines simply sit on top of the edge of the grey shaded zone.

insignificant.

If we think of a light curve as a sampling distribution of a given statistic, in this case the mean magnitude of (or flux from) a star, then the scatter in the light curve is representative of the standard error of the mean. For gaussian-distributed 'scatter' (a sound approximation to the scatter in a light curve due to the various different noise types and sources), the standard error of the mean can be reduced if we observe the target once more, and bin the light curve to the same time interval. The standard error of the mean scales as $N^{-1/2}$, where N is the number of observations.

If we observe our object undergoing a 0.3% dip with 100 points in the lightcurve, at fixed time intervals (light curve cadence) Δt, and with a photometric precision on each point in the light curve of 0.3%, we obtain a 1σ detection. If, when the object undergoes a transit once more, we observe again with the same photometric precision and time interval, and again another 100 points in the light curve, and combine this light curve with the previous one into a binned light curve, we would expect scatter in the binned light curve of $\sigma_{phot}/\sqrt{2} = 0.0021 = 0.21\%$ ($\equiv 1.43\sigma$ detection) . To obtain a 3σ detection of our imaginary 12^{th} magnitude star undergoing a 0.3% dip, we must therefore observe it on $N = (0.003/0.001)^2 = 9$ separate occasions.

Obviously, giving up the requisite telescope time to make 9 separate observations

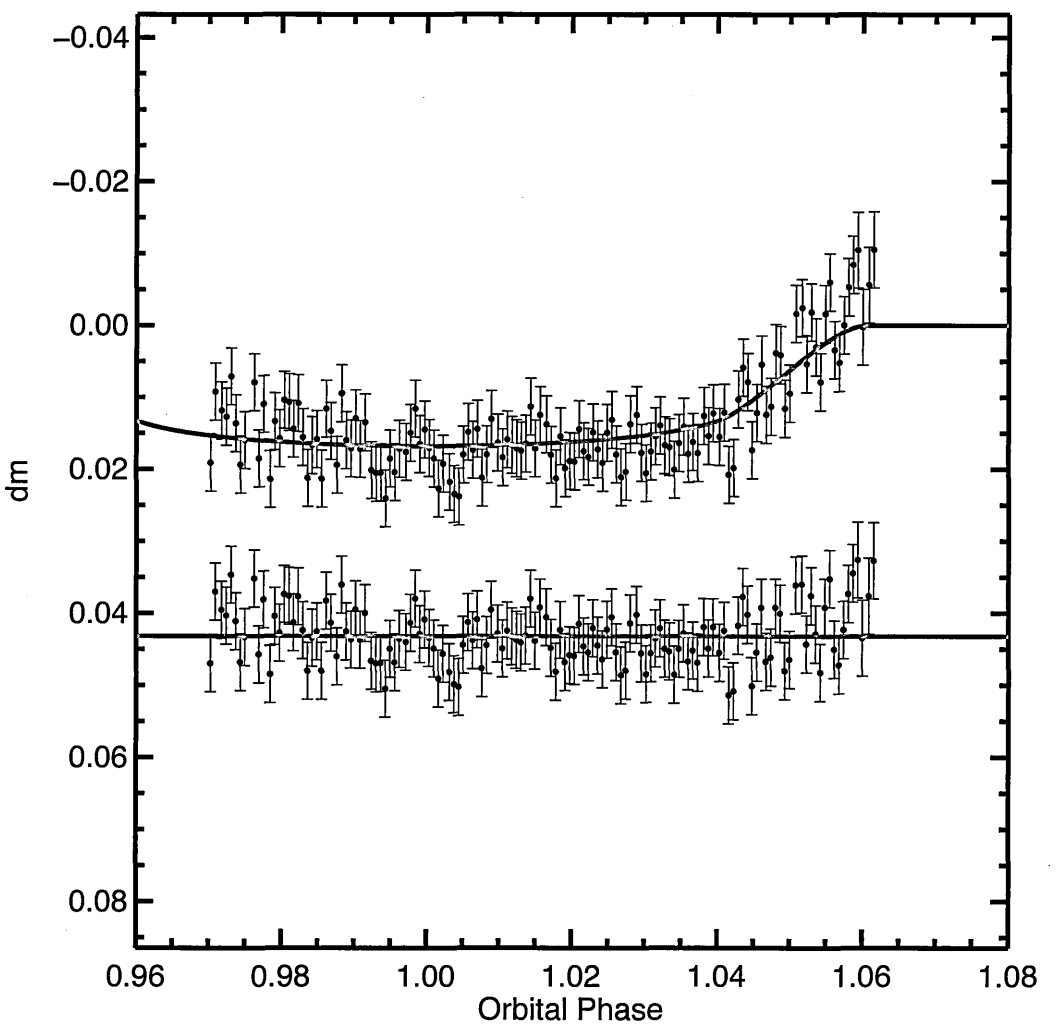

Figure 5.11: The observed transit of WASP-12b corresponding to the observing window plotted in Fig. 5.10. As predicted by the transit timing plot, the transit mid-point and egress were indeed observed. The residuals to the model fit (from the MCMC fit described in Chapter 6) are shown offset.

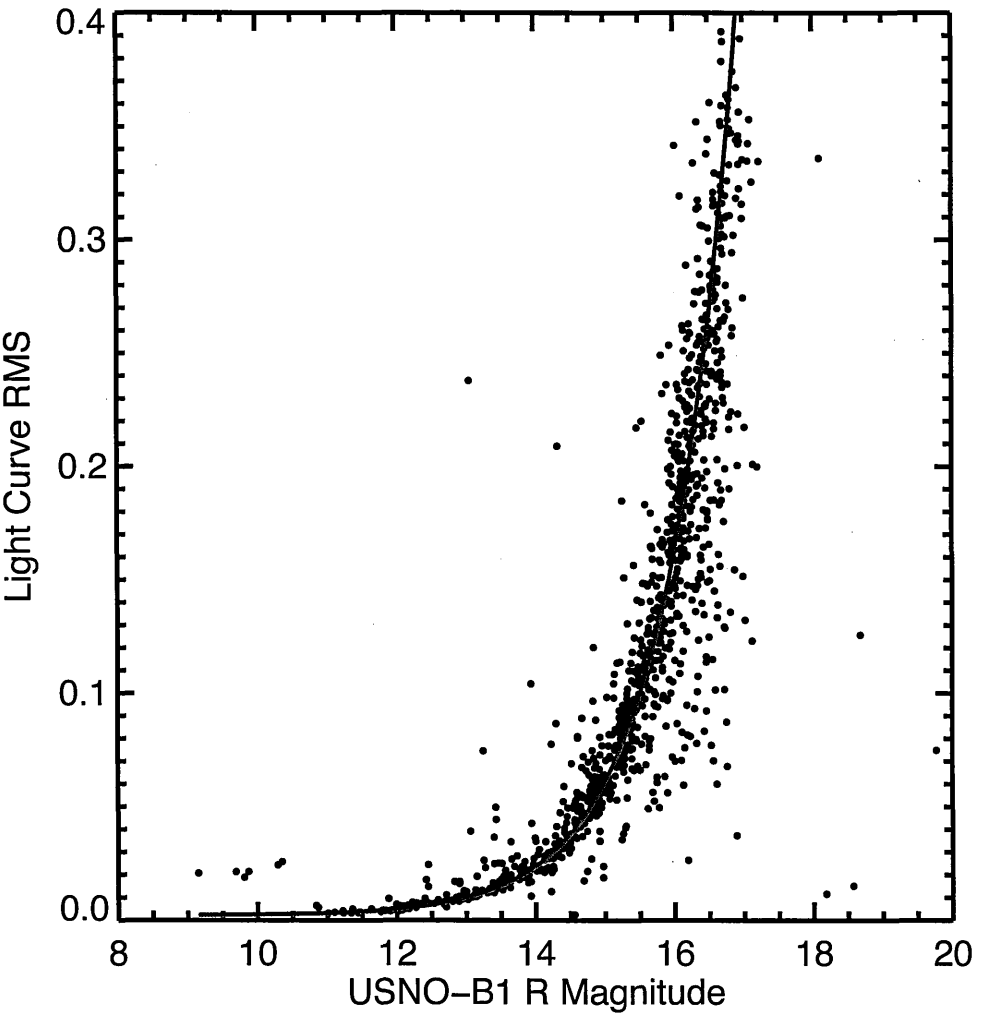

Figure 5.12: Exponential fit of the form RMS $= \alpha_1 e^{m_R} + \alpha_2$ to the plot of light curve RMS versus USNO R catalogue magnitude for the WASP-12 field, observed on 13/01/2011 with PIRATE Mark 1.5, 45s exposures in the R band.

of the target is only a possibility if the transit ephemeris is well constrained. Even then, the observing time required to produce a bare-minimum 3σ detection is most probably unreasonable, given the chance it is most likely to be an astrophysical false-positive. On the upside, such small transit depths, if real, yield precious small planets. Pragmatically, such time cannot be spent on such shallow-transit objects, and their importance as follow-up targets for PIRATE should be downgraded according to their brightness and transit depth.

This discussion of 'worthiness' of PIRATE follow-up for shallow transit candidates can be put into a more rigorous framework by generalising the example from the previous two paragraphs, and investigating the number of required repeat observations

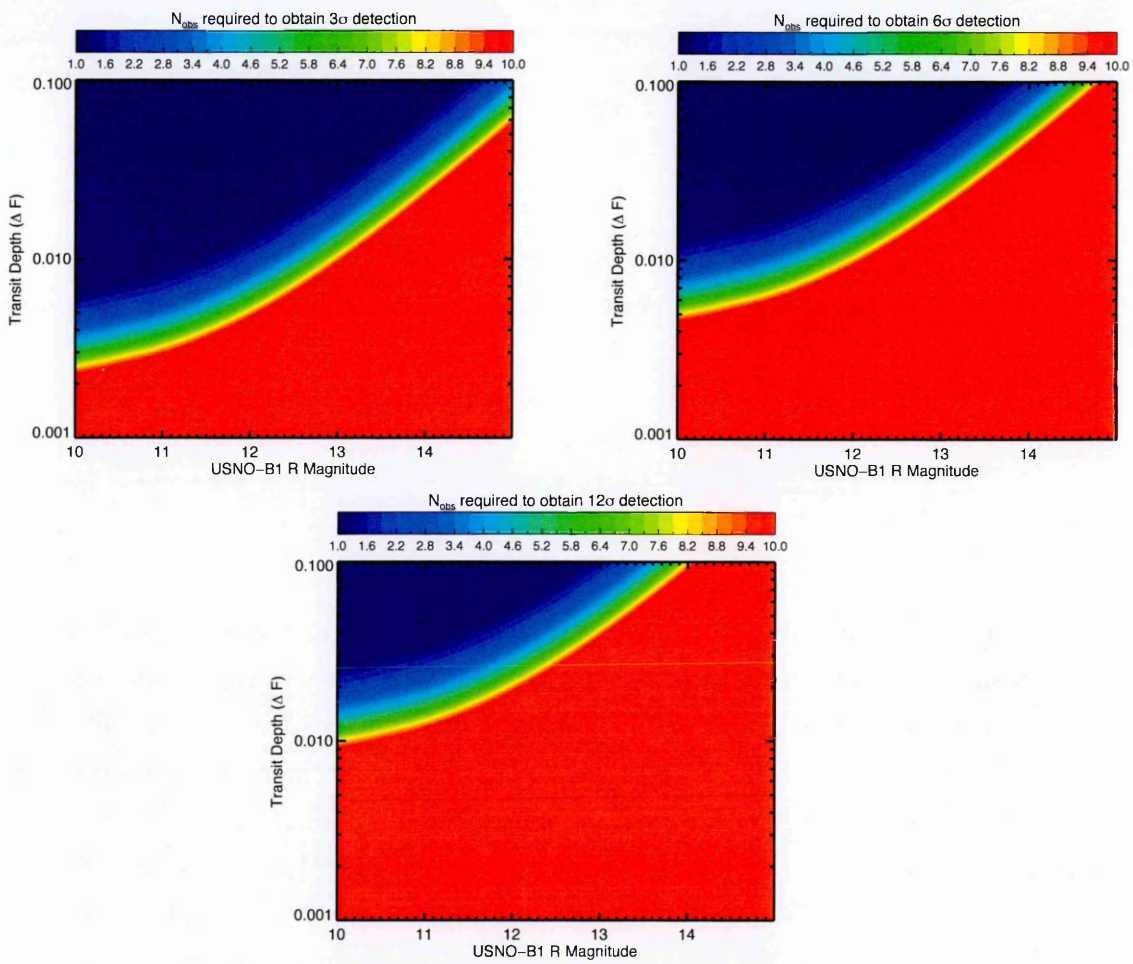

Figure 5.13: The number of repeat observations required (denoted by colour) to achieve a 3σ (top left), 6σ (top right), & 12σ (bottom) detection with PIRATE Mark 1.5, for a range of USNO-B1 R magnitudes and candidate transit depths.

for a range of stellar magnitudes and transit depths. First, let us derive the general expression for the required number of repeat observations of a transit event to achieve a target 'sigma detection threshold'. The light curve RMS of the binned light curve is:

$$\sigma_{binned} = \frac{\sigma_{phot}}{\sqrt{N_{obs}}}$$

where σ_{binned} is the RMS of the binned light curve, σ_{phot} is the photometric uncertainty estimated from the light curve RMS of a single light curve, and N_{obs} is the number of separate observations of a transit event. Let us define the 'sigma detection threshold', σ_{det}, as:

$$\sigma_{det} = \frac{\Delta F}{\sigma_{binned}}$$

where ΔF is the transit depth, or flux deficit. We can trivially combine these two expressions to show that :

$$N_{obs} = \left(\frac{\sigma_{det}\sigma_{phot}}{\Delta F}\right)^2 \tag{5.4}$$

To understand the combinations of stellar magnitude and transit depth that we should avoid and those that we should consider, possibly even for multiple repeat observations, it would be useful to have a map of N_{obs} for a range of magnitudes and transit depths. Before this can be produced, a fit must be made to the data from 11/01/2011 in the 'RMS vs. USNO-B1 R magnitude' plot (Fig. 3.4), so that for a a given input stellar magnitude, a best-estimate light curve RMS may be known. Fig. 5.12 shows an exponential relation of the form RMS $= \alpha_1 e^{m_R} + \alpha_2$ fit to the very same data. Values of $\alpha_1 = 1.8 \times 10^{-9}$ & $\alpha_2 = 0.002$ were used for the co-efficients. Using this fit, it is then possible for a given stellar magnitude to know σ_{phot}, and therefore N_{obs} for a given transit depth. Fig. 5.13, shows this map of N_{obs} for a 3σ, 6σ, and 12σ transit detection with PIRATE Mark 1.5. The plots are capped so that all values $N_{obs} < 1$ and $N_{obs} > 10$ are set to $N_{obs} = 1$ and $N_{obs} = 10$ respectively, as 'less than one observation' doesn't make much meaningful sense, and observing an object's transit more than 10 times is unlikely due to time constraints.

5.2 Full follow-up programme results

5.2.1 Overview

Whilst some SuperWASP 'A' class candidates were observed in 2009 and early 2010, the follow-up programme using PIRATE was fully initiated from October 2010 onwards. For this thesis, follow-up observations made with PIRATE Mark 2 (i.e. using the STX-16803 camera - see Table 2.1) are excluded. Though this author made many follow-up observations with PIRATE Mark 2, the new hardware was added towards the end of the data-collecting phase of this project; and no commissioning and characterisation of the instrument had been performed, and thus this data is excluded from this thesis. All 'attempted observations' are also excluded from this work, so that only successful observing runs are included. 'Attempted observations' include those for which:

- persistent cloud cover blocked successful running of the telescope
- hardware issues occured, be they loss of connection between components, or any other issue disrupting the running of the facility
- weather-triggered shutdown (often humidity) occurred, cutting an observation run short.

Typically, a run is considered 'cut short' and of little use if less than one hour's worth of observations have been collected.

Between October 2010 and October 2011, 56 SuperWASP 'A candidates' were observed with 76 observing runs, so some targets were observed more than once. Some of these repeat observations became redundant however, as they were made before the initial observing run was reduced. In the case where the initial observing run conclusively excludes a candidate as a false-positive, the repeat observations were not subsequently reduced.

The follow-up observations are classified into five distinct groups:

- 'V' - 'V-shaped' - for transits events which are 'V-shaped' and likely to be a grazing EB or blended EB.
- 'Bl' - 'Blend' - for resolved blend scenarios, with a nearby deep eclipsing binary shown to be a blend culprit.

- 'Pl' - 'Plausible' - for 'plausible' planetary-shaped transits events, for which photometry alone can not determine the true nature of the object.
- 'ND' - 'Non-Detection' - for no detected event in the observation window.
- 'U' - 'Unknown' - for transit events where classification is not possible, due to suspicions over the quality of the data, or any other factor that makes classification impossible.

Table 5.2: Follow-up classification by SuperWASP ID.

The table is ordered by increasing right ascension

SuperWASP ID	Date of Obs.	Filter used	Exp. Time (s)	n_{frames}	Class
J001101.**+******.*	17/09/2011	V	45	141	Pl
J002328.**-******.*	10/12/2010	R	20	295	Pl
J003437.**+******.*	19/09/2011	V	30	135	Pl
J003750.**+******.*	29/12/2010	B, R	60, 30	103, 103	Pl
J012918.63+415105.3	12/09/2011	V	45	261	ND
J021225.84+335846.9	14/09/2011	V	45	134	Bl
J022921.**+******.*	20/09/2011	V	45	182	U
J023140.65+504840.7	16/11/2010	R	25	199	ND
J030556.**+******.*	20/08/2011	V	30	297	Pl
J030647.88+395500.8	13/09/2011	V	60	115	V
J031533.65+203903.0	16/09/2011 21/09/2011	V	45	118	ND
J032723.70+060123.3	03/11/2010	R	45	172	V
J032959.**+******.*	21/09/2011	V	30	166	Pl
J040403.**+******.*	11/11/2010	R	25	332	U
J042142.87+513305.6	08/09/2011 15/09/2011	V, V	20, 30	141, 42	Bl
J051109.87+222428.3	16/12/2010 03/01/2011	R, R	45, 45	245, 247	Bl
J053512.**+******.*	13/09/2011	V	50	382	Pl

J061332.**+******.*	09/02/2011	V	45	103	Pl
J065617.52+192123.5	07/04/2011	V	120	228	Bl
J070238.**+******.*	19/09/2011	V	45	124	U
J072230.10+482503.2	23/11/2010	R	120	63	Bl
J075548.**+******.*	10/12/2010	R	20	417	Pl
J084457.35+411945.3	14/03/2010	I	120	179	V
J084943.30+525745.4	06/04/2011	V	60	249	ND
J090926.34+553836.8	11/01/2010 15/01/2010 28/01/2010 01/02/2010	I, I, I, I	60, 60, 60, 60	182, 241, 101, 213	V
J092349.39+503821.4	08/02/2011 07/04/2011	V, V	45, 60	127, 228	Bl
J125234.11+350344.6	04/04/2011	V	45	247	ND
J145526.85+264408.1	20/07/2011	V	90	105	Bl
J152718.**+******.*	14/06/2011	R	60	221	Pl
J171804.**+******.*	13/06/2011	R	45	55	U
J172431.36+412230.4	05/04/2011	R	45	246	ND
J172917.65+065655.0	06/04/2011 13/07/2011	V, V	60, 45	136, 151	Bl
J173037.**+******.*	01/08/2011	V	45	167	U
J174155.50+214427.1	14/08/2011	R	45	87	ND
J174313.27+093649.6	10/07/2011	R	45	80	Bl
J174434.32+103839.2	31/07/2011	R	60	103	ND
J175856.**+******.*	13/06/2011	R	45	188	U
J180509.14+255516.6	21/07/2011	V	45	170	ND

J181022.15+172132.3	03/10/2010	R	60	93	Bl
J181113.13+141441.9	25/07/2011	R	45	206	V
J182429.**+******.*	08/08/2011	R	45	200	U
J184527.98+604038.6	04/08/2011 11/08/2011	R, R	60, 60	174, 157	Bl
J193542.37+602705.4	27/07/2011	R	45	160	Bl
J195009.77+581730.4	16/07/2011	V	45	161	V
J205053.86+183333.9	16/10/2010 11/11/2010	R, R	45, 45	169, 159	Bl
J210318.**+******.*	15/08/2011	R	45	232	U
J210841.23+104830.3	03/08/2010 06/08/2010 07/08/2010	R, R, R	120, 120, 120	124, 118, 94	Bl
J211916.02+152057.7	16/12/2010	R	45	169	Bl
J212707.48+131231.0	23/11/2010	R	60	87	Bl
J213320.84+543037.7	05/10/2010 14/10/2010	R, R	30, 30	133, 216	Bl
J213602.95+161602.2	03/11/2010 06/08/2011	R, R	45, 45	191, 102	Bl
J215642.90+333410.9	30/09/2010	R	60	65	ND
J223531.**+******.*	05/10/2010 20/10/2010 24/11/2010	R, R, R	15, 60, 30	403, 123, 300	Pl
J232332.12+522539.5	26/10/2010 02/11/2010	R, R	45, 45	307, 120	Bl
J232718.73+551928.2	10/08/2011	R	45	203	ND

J233129.75+293952.2	30/09/2010 18/10/2010	R, R	45, 45	67, 162	ND
J234639.98+310921.4	27/10/2010	R	45	134	ND

Classification type	Number of candidates
Plausible planetary candidate	11
Resolvable blend	19
V-shaped transit	6
Null result - no signal detected	13
Unclassifiable	8

Table 5.3: Summary table detailing the number of candidates in each classification category

5.2.2 Notable candidates

Rather than present the light curves and observation details of all objects observed in Table 5.2, a few have been selected from the table for presentation. As archetypal candidates representative of each category have already been presented, the following notable candidates represent those that have either proved troublesome in their classification, or require a more nuanced argument for their classification.

5.2.2.1 1SWASPJ181113.13+141441.9

This target was detected in two separate fields, one from the 2004 season, the other from the 2007 season, with a period of $\sim$ 3.28 days, a transit depth of $\sim$ 0.02 mag, and a transit width of $\sim$3.5 hours. There are 4 faint objects in the NOMAD catalogue (Zacharias et al., 2005) within one 3.5 pixel SuperWASP aperture that are within 5 magnitudes of the target star, and so this object was flagged for follow-up photometry - to determine whether the bright target was undergoing eclipse, or a nearby faint neighbour. Fig. 5.14 shows the field as seen in the PIRATE images. Multiple sources are detected within the 4.5 SuperWASP pixel aperture (yellow), though all are very faint.

The ephemeris as determined by an MCMC run on the 'ORFG_TAMTFA' field, which aggregates both the 2004 and 2007 data, is:

$$T_{\rm mid} = 2453431.14584(\pm 0.00266) + 3.2826447E\ (\pm 0.0000176)\text{d HJD}$$

The ephemeris was thus sufficiently well constrained for the PIRATE follow-up photometry, for which the observing window was 2455768.358 - 2455768.566 (HJD). We can clearly see from Fig. 5.15 that the observing window of the follow-up photometry should result in either egress and some post-transit baseline being captured, or mid-transit, egress, and post-transit baseline being captured, and this is indeed what was seen, see Fig. 5.16. As can be seen from the light curve, 1SWASPJ181113.13+141441.9 is indeed the object undergoing eclipse, and not the nearby faint neighbours. The blend scenario is thus ruled out. However, as Table 5.2 reveals, this object was classified 'V' for 'V-shaped transit', and was thus rejected as a possible planetary candidate within the SuperWASP archive.

What characteristics of the light curve betray its true nature? It is apparent from the PIRATE light curve that the transit depth is at least 2%, though possibly greater. However, the duration of egress looks suspicious. Though the transit duration is not measurable from the PIRATE light curve alone, we can put a lower limit on the duration of ingress/egress; which appears to be (from a simple by-eye estimation) $\gtrsim$ 100 minutes. When combined with the prior expectation of the full transit duration being $\sim$3.5 hours / $\sim$ 210 mins (from all of the various HUNTER runs on the SuperWASP data), however, it is apparent that the duration of the transit 'floor' (t_F in Eq. 5.2) might be very small indeed; perhaps with an upper limit of $\sim$10 minutes. Using this information we might infer that $t_F/t_T \lesssim 0.05$ for this system. So from simply having prior knowledge of the MCMC-determined transit duration from the SuperWASP data, combined with a single PIRATE observation of what appears to be the majority of egress, we can, in combination with Eq. 5.2 and Fig. 5.7, deduce that this is likely to be a high impact parameter light curve, with $b \gtrsim 0.8$. If the host star is a main-sequence star then it will prove troublesome to reconcile the high impact parameter with the measured transit duration, orbital period, and scale of the system ($R_* : R_p : a$).

These are, of course, imprecise 'by-eye' estimations, but we can see these estimations and suspicions of the non-planetary nature of this system confirmed by running two MCMC simulations: with and without the main-sequence mass-radius constraint. This constraint is applied via a Bayesian prior 'penalty' to the fitting statistic for solutions that stray too far from a main-sequence stellar mass and radius.

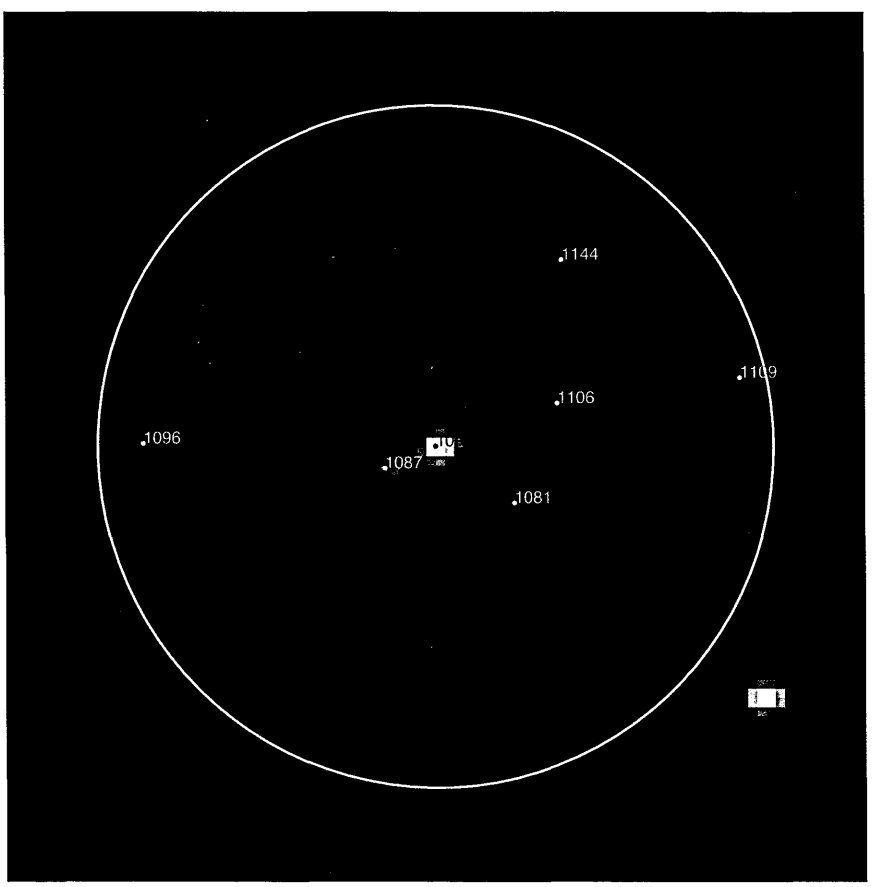

Figure 5.14: PIRATE image of the field surrounding 1SWASPJ181113.13+141441.9 (V = 12.62, ID=1094). The largest SuperWASP measuring aperture is overlayed in yellow, with a radius of 61.65". A faint companion (ID = 1087) is present, so this target was selected for follow-up photometry to determine if the target is the object undergoing eclipse.

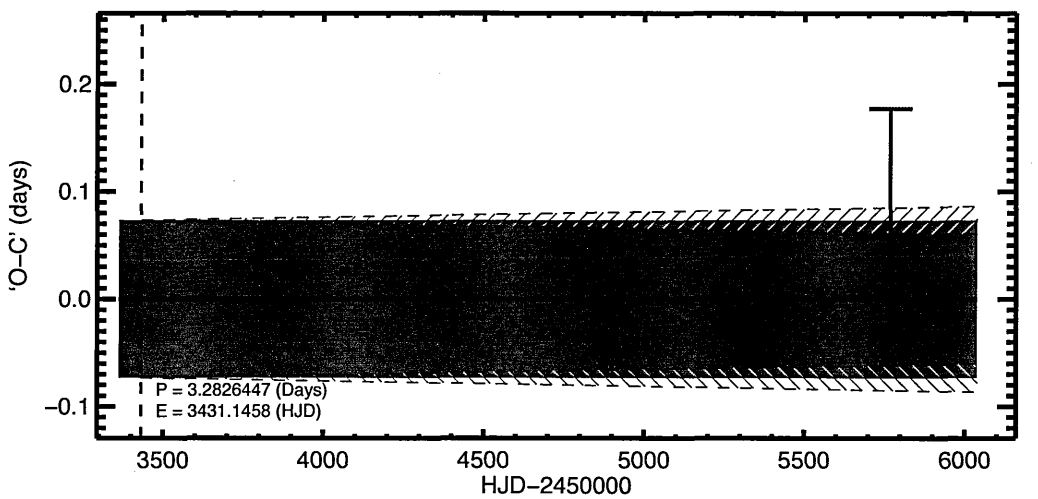

Figure 5.15: Visualisation of the potential ephemeris drift for 1SWASPJ181113.13+141441.9. The capped bar represents the PIRATE observing window, see Fig. 5.9 for further explanation.

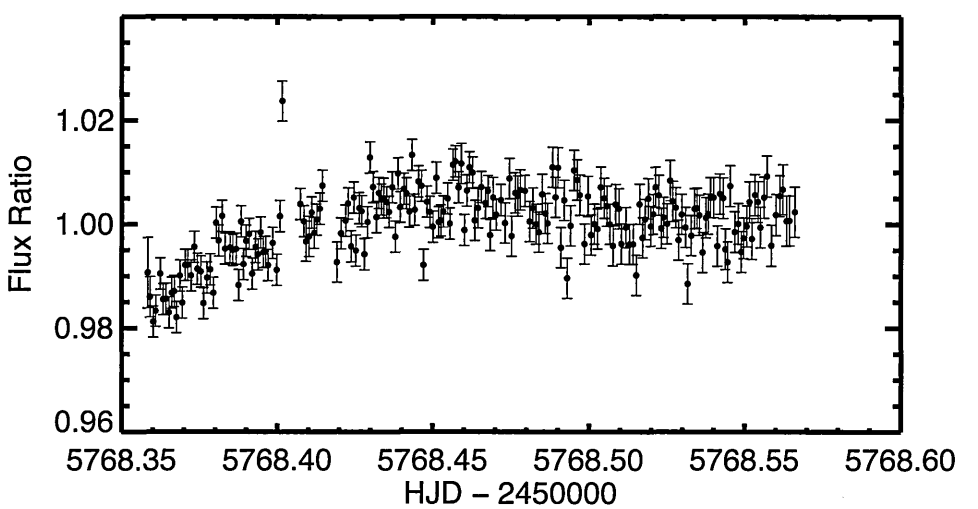

Figure 5.16: The light curve of 1SWASPJ181113.13+141441.9 from the night of 25/07/2011.

The best-fit parameters from the simulation with the mass-radius constraint turned on are presented in Table 5.4, with the corresponding plots of both the PIRATE follow-up photometry and the SuperWASP photometry in Fig. 5.17. The fit suggests an inflated (1.7 R_J) planet orbiting a sub-Solar density host star. We recover a transit width of 216 minutes, akin to that from the 'discovery' fit to SuperWASP data alone. However, visual inspection of the model fit to the PIRATE follow-up photometry suggests an under-estimation of the egress duration, and a possible underestimation of the total transit duration. The under-estimation of the transit duration is borne out somewhat when comparing the fit to the SuperWASP data, which also apparently exhibits a greater transit duration than that afforded by the model. As expected, it has proved tricky for the fitting process to reconcile an apparent high impact parameter light curve with the measured transit duration, orbital period, and scale of the system ($R_* : R_p : a$).

Switching the Bayesian prior off, however, is illuminating; see Table 5.5. The fitting procedure achieves a greater transit width (~300 mins.), transit depth (3.7% instead of 2%), and an impact parameter greater than unity. The primary and secondary radii (2.76 $R_\odot$ and 8.62 R_J respectively) coupled with the high impact parameter point to a grazing eclipse between two stellar objects. In effect, with the stellar radius free to roam from the main sequence value inferred from the $J-H$ colour, the stellar radius has been dramatically increased in order to reconcile the high impact parameter with the ratio $^{t_T}/_P$. Notable however is the very low stellar density, as $M_* = 1.45 M_\odot$. This may actually add weight to the argument that the light curve is not that of a grazing system, but that of a fully-eclipsing system diluted by third light in order to achieve the observed flux deficit (Seager and Mallén-Ornelas, 2003). Either way, the validity of the non-planetary solution is further supported by the superior fit to the binned SuperWASP data, shown in Fig. 5.18. The fitting statistics also bear this out. For the solution with the mass-radius constraint switched on, $\chi^2_{\rm min} = 5261.55$; whereas $\chi^2_{\rm min} = 5249.69$ for the solution with the mass-radius constraint switched off.

5.2.2.2 1SWASPJ090926.34+553836.8

Aside from the incompatibility of a V-shaped transit of shallow depth with that of a true exoplanetary transit, another feature that betrays the imposter's true nature (be it grazing EB or blended EB) is a secondary eclipse of shallower depth than the primary. 1SWASPJ090926.34+553836.8 is a target that demonstrates the benefit of

Parameter	Value	1 σ (upper)	1 σ (lower)
Transit Epoch (T_0) (HJD)	5531.9699	0.0030	0.0029
Orbtal Period (days)	3.282546	0.000006	0.000009
$(R_p/R_*)^2$	0.020	0.001	0.001
Transit Width (days)	0.150	0.007	0.006
b (impact parameter)	0.10	0.08	0.06
Stellar mass ($M_\odot$)	1.14	0.03	0.03
Stellar radius ($R_\odot$)	1.23	0.06	0.05
Stellar density ($\rho_\odot$)	0.61	0.08	0.08
Orbital semi-major axis (a/AU)	0.0452	0.0004	0.0004
Orbital inclination (i)(°)	89.3	0.5	0.6
Planetary radius (R_p/R_J)	1.70	0.09	0.08

Table 5.4: Best fit parameters from the MCMC analysis of SuperWASP & PIRATE data for 1SWASPJ181113.13+141441.9, with the main sequence mass-radius constraint turned on.

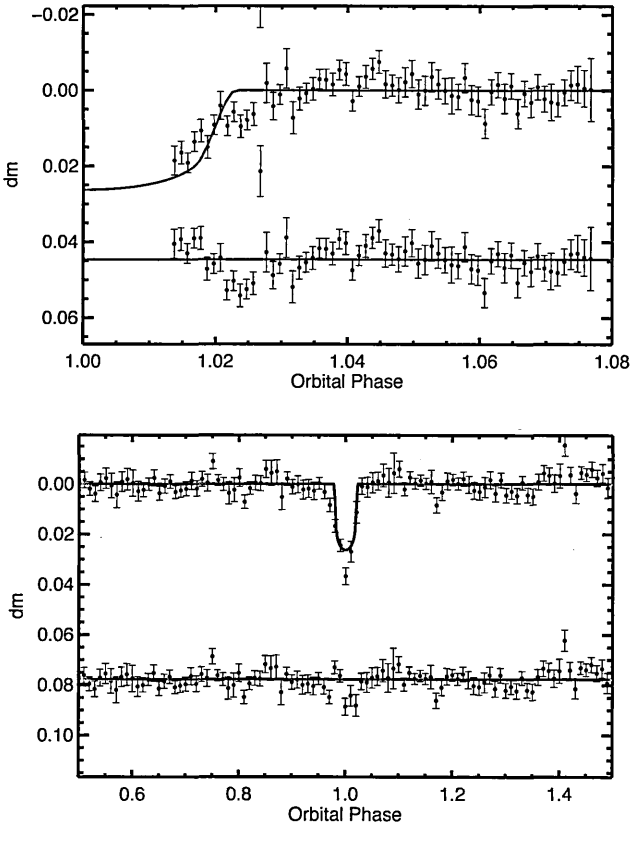

Figure 5.17: Top: PIRATE R band data binned to $\phi = 0.001$. Bottom: All SuperWASP data binned to $\phi = 0.01$. The model fit was made with the main sequence mass-radius constraint turned on. Below each light curve are the residuals of each fit.

Parameter	Value	1 σ (upper)	1 σ (lower)
Transit Epoch (T_0) (HJD)	5531.9669	0.0034	0.0037
Orbtal Period (days)	3.282530	0.000006	0.000007
$(R_p/R_*)^2$	0.037	0.005	0.006
Transit Width (days)	0.206	0.011	0.009
b (impact parameter)	1.08	0.09	0.12
Stellar mass ($M_\odot$)	1.45	0.05	0.05
Stellar radius ($R_\odot$)	2.76	0.20	0.17
Stellar density ($\rho_\odot$)	0.07	0.01	0.01
Orbital semi-major axis (a/AU)	0.0490	0.0006	0.0006
Orbital inclination (i)(°)	73.4	2.4	1.8
Planetary radius (R_p/R_J)	8.62	2.01	2.51

Table 5.5: Best fit parameters from the MCMC analysis of SuperWASP & PIRATE data for 1SWASPJ181113.13+141441.9, with the main sequence mass-radius constraint turned off.

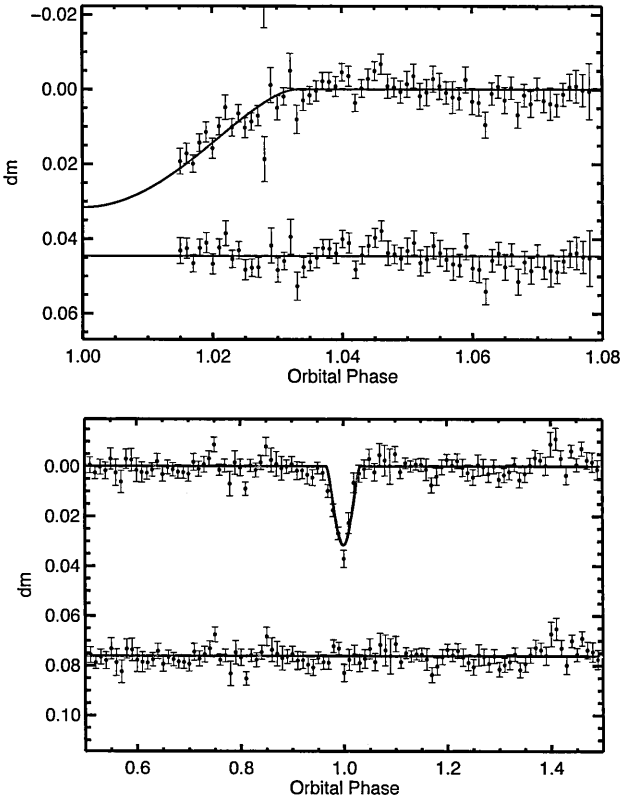

Figure 5.18: Top: PIRATE R band data binned to $\phi = 0.001$. Bottom: All SuperWASP data binned to $\phi = 0.01$. The model fit was made with the main sequence mass-radius constraint turned off. Below each light curve are the residuals of each fit.

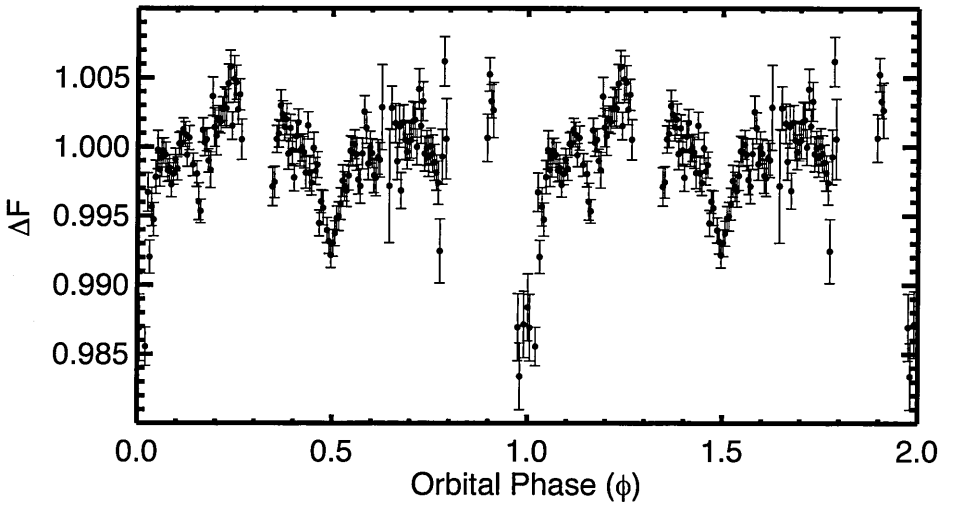

Figure 5.19: Phase-folded light curve of all the PIRATE data for the object 1SWASPJ090926.34+553836.8; binned to bin size $\phi = 0.005$. The data is repeated over two full cycles.

following up such a target at orbital phase $\phi = 0.5$, as opposed to observing solely at $\phi = 0$ (or 1). This target was detected with a principal period of 0.628501d and epoch 2455071.3629, with a depth in the SuperWASP data of ~0.006 magnitudes. Suspicions are aroused however by the MCMC solution indicating a $0.55R_{\rm J}$ body orbiting a $0.75R_{\odot}$, $1.2M_{\odot}$ host star. The fact that ${}^{R_*}/_{R_{\odot}} \ll {}^{M_*}/_{M_{\odot}}$ already points to a probable false positive detection, due to the implied stellar density ($2.84\rho_{\odot}$) not being compatible with that of a main sequence star.

This target was observed on four separate occasions (see Table 5.2), once at $\phi = 0$, and 3 times at $\phi = 0.5$. A separate observation made by Grant Miller at St. Andrews with the James Gregory Telescope (data not included here) of the primary transit confirms the depth seen in the PIRATE data (~1.3%) at $\phi = 0$. A follow-up observation was made with PIRATE at $\phi = 0.5$, which showed signs of a secondary eclipse of depth $\sim$ 0.6%. As the depth of the eclipse was shallow, two further observations were made in order to confirm the depth of the secondary eclipse. All of the PIRATE observations are shown phase-folded and binned to $\phi = 0.005$ in Fig. 5.19. Two different eclipse depths are apparent, leading to the classification of this object as an EB with either third-light contamination or a grazing orbital inclination.

Chapter 6

Observations of known planetary systems

6.1 WASP-12b

Since its discovery in 2009, WASP-12b (Hebb et al., 2009) has received a lot of follow-up attention. At the time of its discovery it was the hottest known transiting exoplanet due to the high level of irradiation from its host star throughout its close-in orbit. The planet is an inflated (1.79 R_J) hot Jupiter in a 1.09 day orbit around an F9V ($T_{\rm eff} = 6300^{+200}_{-100}$K) host star of super-solar metallicity ([Fe/H]=$0.3^{+0.05}_{-0.15}$). The planet has a mass of 1.4±0.1 M_J, and resides in a short period orbit, requiring that it orbits at a distance of just 3.1 solar radii from its star, meaning the planet is subject to intense tidal forces. Li et al. (2010) predicted significant mass-loss through L1, and a spiral-in-time of ∼10 million years, assuming a tidal quality factor for the star of $< 10^6$. They also posited tidal heating as a potential driver for the puffed-up nature of the planet, replacing energy lost through radiated heat flux on the planet's night-side. To further probe the bloated nature of this very Hot Jupiter, near-UV transmission spectroscopy was obtained by Haswell (Open University) using the Cosmic Origins Spectrograph (COS) on the Hubble Space Telescope. As a part of this study (Fossati et al., 2010), PIRATE was employed to update the transit ephemeris as determined by Hebb et al. (2009) in the visible part of the spectrum. Initial results from the COS data suggested an early ingress attributable to absorption by resonance lines in a metal-rich and extended exosphere.

6.1.1 Confirmation of optical ephemeris

WASP-12 was observed on 6 separate occasions in Autumn 2009 by PIRATE Mark 1, detailed in Table 6.1. No single observation captured a full transit, though a partial transit was captured on 4 occasions. The data were collected via the usual method, with relevant calibration frames taken each night (flats, bias, dark) and processed by the pipeline in the usual manner. All of the observations were made in the Johnson R filter, with the camera cooled to -20°C. As these data were processed with the PIRATE pipeline described in Chapter 6, each light curve is made by comparing the flux of the target (WASP-12) with that of a uniquely defined ensemble flux, picked by the pipeline to minimise the RMS of the final target light curve. The initial purpose of these observations was to confirm the discovery ephemeris of Hebb et al. (2009), which was performed once full phase coverage of the transit was achieved after the 6^{th} successive observation, by simply phase folding the PIRATE data and comparing it to the discovery paper model, as supplied by Hebb (private communication). It was readily apparent that the discovery (optical) ephemeris was still valid (within the uncertainties), and therefore the opportunity was presented to update the ephemeris. A seventh observation was made in January 2011 with PIRATE Mark 1.5 as part of its commissioning, and this was included for the benefit of ephemeris refinement. All of the light curves can be seen in Fig. 6.1, overplotting the transit model determined in subsection 6.1.2.

The orbital period of the system is arguably the most effectively improved orbital parameter from small-telescope follow-up observations of known transiting planets. Follow-up observations that significantly extend the baseline of observations (and therefore increase the number of cycles from T_0) can dramatically reduce the uncertainty in the period and initial epoch, even if the photometry itself is taken in non-photometric conditions. The impact of a poorly defined transit mid-point is negated by extending the observation over many cycles. This is demonstrated neatly by Table 6.2, which details the determined orbital period (via an MCMC run in each case) for a variety of sources. Note that by combining 4 partial transits of WASP-12b (made with PIRATE in 2009) with the pre-existing SuperWASP data, the period is improved by more than an order of magnitude.

6.1.2 Ephemeris refinement with additional light curves

A second HST visit to WASP-12b in NUV was performed in March 2010, and a further extended analysis of the deeper transit in the NUV and its implications for

Date	HJD -2450000 (start)	HJD - 2450000 (end)	Integration time (s)	Comments
15/10/2009	5121.53078	5121.68142	120	Out-Of-Transit (OOT) only
18/10/2009	5123.49117	5123.67301	120	OOT only
11/11/2009	5147.47507	5147.62549	120	Egress & OOT
12/11/2009	5148.47187	5148.51421	60	OOT & ingress
22/11/2009	5158.37221	5158.69467	60	Egress & OOT
23/11/2009	5159.44126	5159.59863	60	Mid-transit, egress & OOT
13/01/2011	5575.26270	5575.36230	45	Mid,transit & egress. (PIRATE Mark 1.5)

Table 6.1: Observing log for WASP-12 observations with PIRATE. The 6 observations in 2009 were made as part of a program to check and update the optical ephemeris of the system for use in conjunction with HST observations of the system in the near-UV (Haswell et al., 2012 - in press). The 2011 observation was initially taken as part of the commissioning program for PIRATE Mark 1.5

Source	Period	Uncertainty
SuperWASP alone (19 transits)	1.09132 days	$\pm 5 \times 10^{-5}$ days
PIRATE alone (4 partial transits in 2009)	1.09143 days	$\pm 7 \times 10^{-5}$ days
SuperWASP + PIRATE	1.091422 days	$\pm 2 \times 10^{-6}$ days
Published (Tenagra II (0.81m), Liverpool 2m, SOPHIE RV (1.93m) (Hebb et al., 2009)	1.091423 days	$\pm 3 \times 10^{-6}$ days

Table 6.2: Comparison of the period determined from a fit using the SuperWASP MCMC routine on a selection of different input sources.

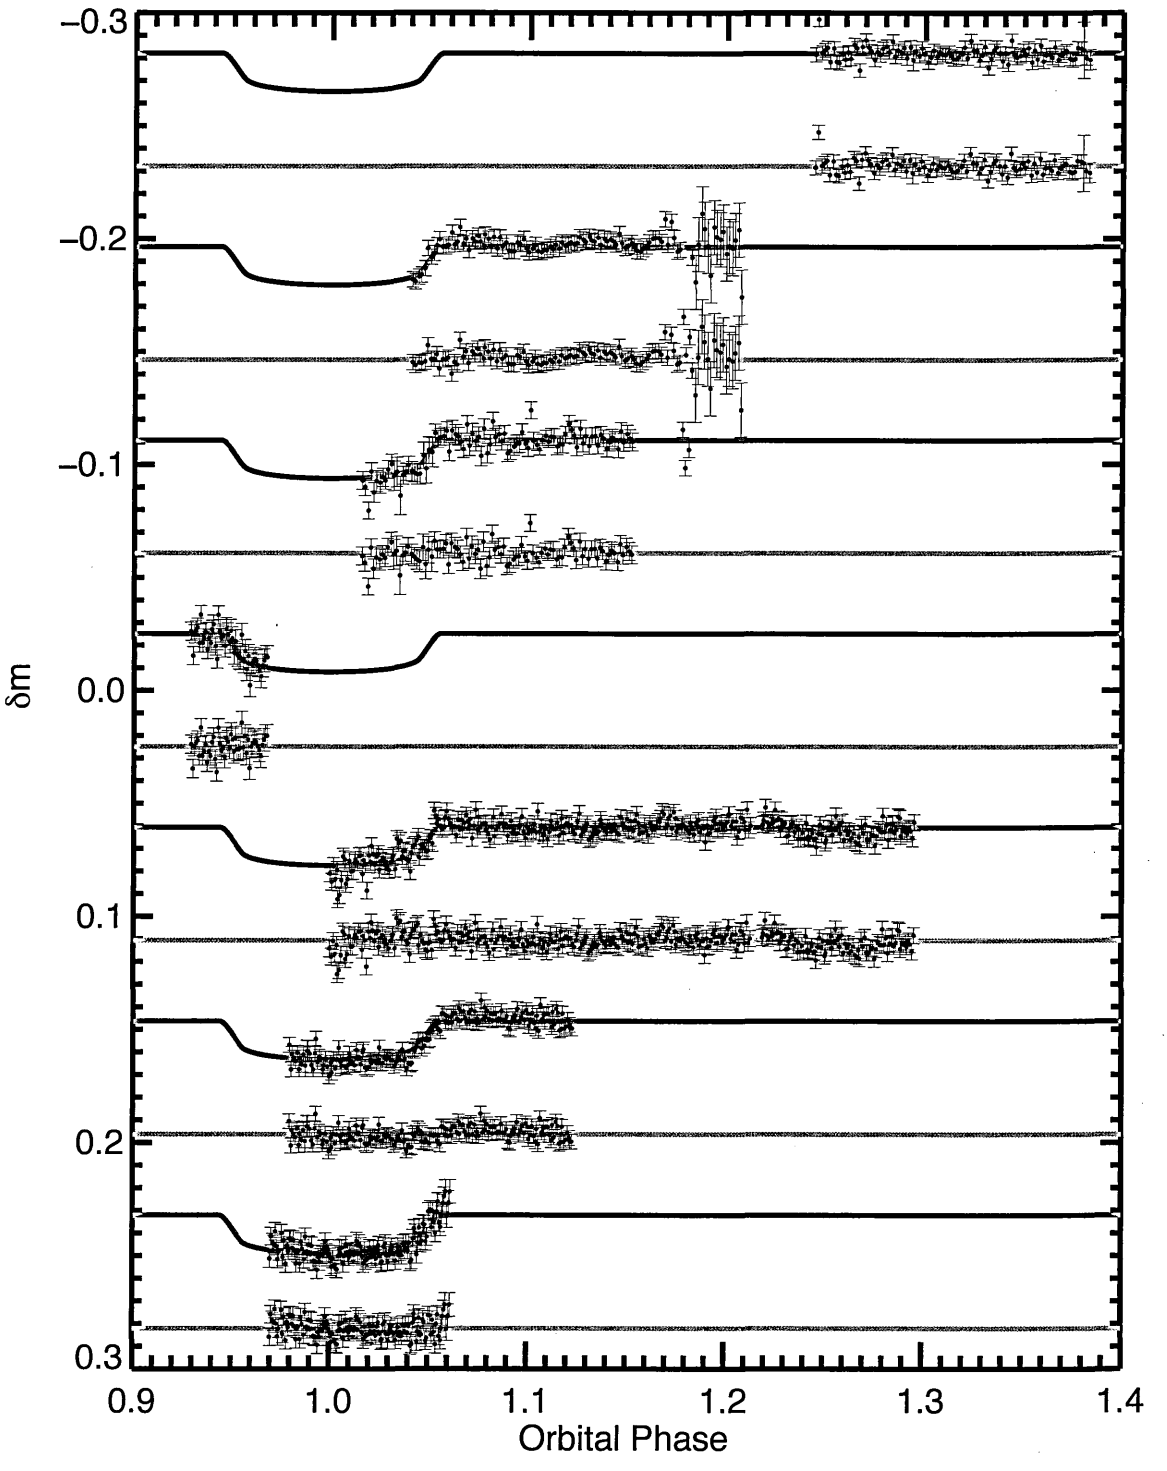

Figure 6.1: Light curves of the individual observations of WASP-12 made with PI-RATE Mark 1 & Mark 1.5 as detailed in Table 6.1. Red lines denote the model fit with the MCMC routine of Enoch et al. (2010) & Collier Cameron et al. (2007), with their corresponding residuals directly below (grey line). The observation dates run from top (earliest) to bottom (latest).

the star-planet interaction is in press (Haswell et al. 2012). As part of this work, the refinement of the optical ephemeris was further extended, including transit observations from other facilities, such as the Faulkes Telescope North, the Liverpool Telescope, the James Gregory Telescope (University of St. Andrews). Comparison of these observations with the transit observations reported in Maciejewski et al. (2011) and Chan et al. (2011) was also undertaken, especially in light of the claim by Maciejewski et al. (2011) of possible TTVs (and thus possibly more planets in the WASP-12 system). Thus a linear ephemeris fit was first determined, and then the deviation of the individual observations from the linear fit examined in the form of an O-C diagram.

6.1.2.1 Faulkes Telescope North observation

Faulkes Telescope North (FTN) data of the transit of WASP-12b were taken on 17/12/2010 in the Bessell-B filter with the Spectral camera. The Spectral camera is a 4k×4k detector (used in binning 2 mode), which, in conjunction with the 2m FTN telescope yields a $0.304''\text{pixel}^{-1}$ plate scale and a $10' \times 10'$ field of view. 22s exposures were used for the observing run, giving a cadence of $\sim$ 47s. WASP-12 was observed between HJD UTC = 2455547.91399 and 2455548.08510. The data were preprocessed using the ARI pipeline and aperture photometry was then performed in IRAF. Thanks go to Rachel Street for obtaining and producing the light curve.

6.1.2.2 James Gregory Telescope observation

The James Gregory Telescope (JGT) is the University of St. Andrews' on-site 0.95m Schmidt-Cassegrain telescope. It has a 1k×1k detector, a $17'$ field of view, and a resultant plate scale of $1''\text{pixel}^{-1}$. WASP-12 was observed on 30/11/2008 in the Cousins-R band with 90s exposures taken over a period of 3 hours from HJD UTC = 2454801.38830 to 2454801.51706.

6.1.2.3 Liverpool Telescope observation

This transit observation made with the Liverpool Telescope (LT) is the same Z band light curve described in the WASP-12b discovery paper (Hebb et al., 2009). The Liverpool Telescope is a 2m robotic telescope on the island of La Palma in the Canary Islands. See discovery paper for further details.

T_{mid}('O') (HJD-2450000)	$\sigma\,(T_{\mathrm{mid}})$ (s)	'O-C' (s)	Source
4515.52455	11.32	5.65793	LT (Z)
4801.47593	35.83	-98.07082	JGT (R)
5123.44587	60.35	-60.35127	PIRATE (R)
5147.45781	77.33	-3.77195	PIRATE (R)
5148.55083	113.16	133.90439	PIRATE (R)
5158.37107	49.04	-86.75496	PIRATE (R)
5159.46575	52.81	194.25566	PIRATE (R)
5548.00980	20.75	3.77195	FTN (B)
5575.29570	54.69	33.94759	PIRATE (R)

Table 6.3: Transit mid-times for the 9 separate observations of the transit of WASP-12b

6.1.2.4 Establishing a linear ephemeris and creating the O-C diagram

The PIRATE, FTN, JGT, LT follow-up observations, in conjunction with all available SuperWASP data were input to the SuperWASP consortium transit-fitting code (see Chapter 4). A new ephemeris of:

$$T_{\mathrm{mid}}(\mathrm{HJD}) = 2454852.7739^{+0.00014}_{-0.00014} + N \times 1.09142206^{+3.3\times10^{-7}}_{-3.1\times10^{-7}}$$

was determined, further improving upon the discovery paper ephemeris by an order of magnitude. In order to look for a TTV signal, the individual mid-transit times of each observation must be compared with the linear ephemeris to look for significant deviations from the expected mid-transit times as predicted by the linear model. The deviations of the observed ('O') mid-transit times from the calculated ('C') mid-transit times can then be plotted as a function of the calculated mid-transit points; the well-known 'O-C' diagram, which can be seen in Fig. 6.2.

The SuperWASP data and any observations which did not capture any part of the transit were excluded from the TTV analysis. The nine remaining light curves were phase-folded and renormalised against a model light curve appropriate for their photometric band using χ^2 minimisation. A phase offset from the linear model was then determined for each light curve, again using χ^2 minimisation. The phase offsets were then converted back to seconds using the period established from the MCMC fit. 1σ uncertainties for the O-C values were derived from the χ^2 minimisation. The observed transit midpoints and their deviations from the linear ephemeris are tabulated in Table 6.3.

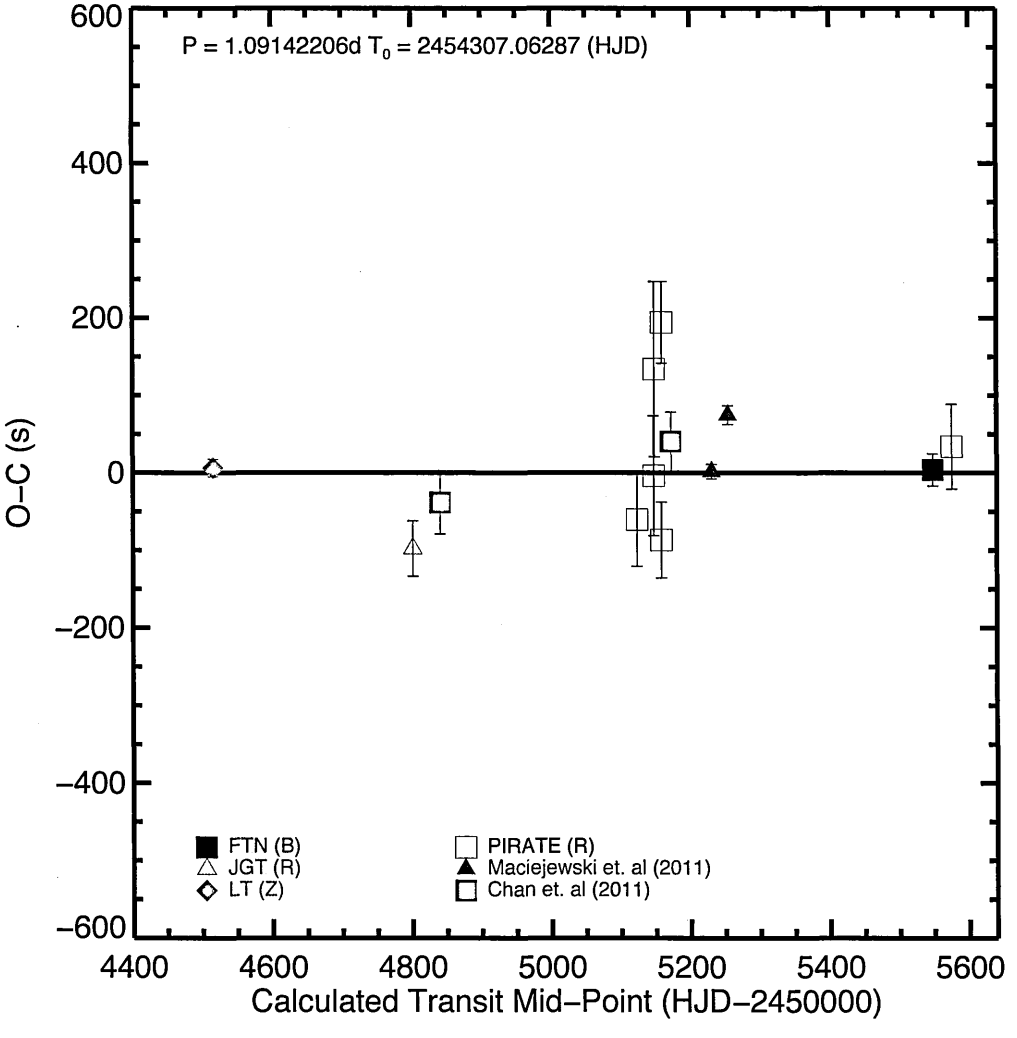

Figure 6.2: O-C diagram for 13 individual transit observations of WASP-12b, incorporating data from the Faulkes Telescope North, the Liverpool Telescope, the James Gregory Telescope (University of St. Andrews), the observations reported in Maciejewski et al. (2011) and Chan et al. (2011), and the previously reported PIRATE observations.

The linear ephemeris is well constrained at both ends by precise light curves obtained with the Liverpool Telescope and later with the Faulkes Telescope North. As previously mentioned, Maciejewski et al. (2011) observed two transits and suggested their measurements might represent the detection of a TTV signal. Whilst this still remains possible, the uncertainties on the Maciejewski et al. (2011) timings are optimistically small. Given that the Chan et al. (2011) measurements agree within 1σ with our linear ephemeris, we suspect the Maciejewski et al. (2011) observations have underestimated uncertainties. TTVs cannot be ruled out however, and so further full high-quality transit observations of this system are encouraged.

It is also noted that this method does not consider the contribution of red noise to the timing uncertainty, as the χ^2 minimisation assumes only gaussian noise. While the high precision light curves (LT and FTN) do not exhibit noticable red noise, some will inevitably be present; and red noise is certainly present in the PIRATE follow-up observations. Red noise in the light curves will further increase the uncertainties of the individual timing measurements; and this is not accounted for in this analysis.

6.1.2.5 Timing accuracy considerations and methods of keeping time

Looking for TTVs of order $\sim$ tens of seconds requires high precision photometry, but also high precision time keeping. In order to claim a transit arrival time is late or early by say 20s, one must be sure of the ability of the instrument to produce accurate time stamps in the FITS headers, especially if the duration of a single exposure alone is greater than the amplitude of the TTV signal. Combining data from multiple sources is hazardous for this reason. In the absence of detailed knowledge of the instrument, we must assume all other intruments are accurately correcting temporal drift in their instruments and synchronising their clocks. Beyond the accuracy of the system clock, this also extends to any timing deficits between the reported time of opening the shutter and the actual opening of the shutter; and any other hardware or software issues that may interfere with timing.

Beyond the accuracy of the FITS header time stamps, care must also be taken to ensure that all time stamps (and hence transit mid-points) are reported in the same timing standards, and for the same astrophysical reference frame. What follows is a brief summary of some timing standards in common use, and the extent to which carelessness with timing standards and reference frames can lead to large discrepancies in reported times. This summary and its implications for the TTV analysis in the previous subsection is based on Eastman et al. (2010).

First it is important to note the difference between one day as defined by the rotation of the Earth on its axis, a so-called 'mean solar day', and a day as defined by an atomic clock. Universal Time (UT) is defined by the duration of a mean solar day. However, the Earth's rotation rate is (in general) gradually slowing down due to tidal braking from the moon, and is therefore an inaccurate measure of the passing of time. International Atomic Time (TAI) is defined by the transition rate between two hyperfine levels of the ground state of Caesium 133; it is consistent and accurate, and serves as the fundamental basis of many other time standards. Coordinated Universal Time (UTC) is commonly used in astronomy, and runs at the same rate as TAI, except that it is not allowed to differ from UT by more than 0.9s. Every six months the International Earth Rotation and Reference Systems Service (IERS) meet to determine whether the addition or subtraction of a second to UTC is required in order to keep it in step with UT. The most recent leap second was added on 30th June 2012, and brings the total difference in seconds between UTC and TAI to 35. UTC maintains the precision of TAI in maintaining the duration of an SI second, but by keeping in step with UT it is, in civil terms, non-disruptive. If TAI was employed as the civil time standard of choice, the constant drift of UT away from TAI would eventually result in mid-day occurring in the middle of the night for any given time-zone. UTC is the time stamp recorded by PIRATE in its FITS headers, and is currently the most commonly used time standard in astronomy.

Terrestial Time (TT) is a simple offset from TAI of 32.184s, used to maintain continuity between TT and its (now-defunct) historical predecessor Ephemeris Time (ET). I reference TT as it relates to the final time standard in use, Barycentric Dynamical Time (TDB), which corrects for the Einstein delay to the geocenter; the delay due to time dilation and gravitational redshift from the motions of the sun and other bodies in the solar system. This is a periodic offset with amplitude ± 3.4 ms, and period 1 year. Knowledge of this time system is relevant as observations are often made in TDB and not UTC. For example, the Maciejewski et al. (2011) observations were initially made in TDB. Clearly, inclusion of data points from multiple instruments in a TTV analysis requires close scrutiny of the time systems used, as TDB $\simeq$ TT=UTC + 35s + 32.184s.

Further to this, differences in the choice of reference frame must be accounted for. The PIRATE pipeline takes the UTC time stamp in conjunction with the right ascension and declination of the target to determine the Heliocentric Julian Date (HJD) of the observation, which adheres to the time keeping principles of UTC but tracks time in fractional days, not the typical sexagesimal format of UTC. This corrects for the

difference in light travel time between the target and the Earth due to the position of the Earth in its orbit around the Sun. In making the conversion, the reference frame is switched from the Earth (Geocentric Julian Date) to the Sun (Heliocentric Julian Date). The choice of implementing HJD in the PIRATE pipeline was made for consistency reasons, as SuperWASP also uses HJD as its reference frame of choice. However, the Sun undergoes a reflex motion about the common barycentre of the Solar System, and so light travel times to the Sun are also a function of the position of the Sun in its orbit around the common barycentre. A more precise reference frame then is used for Barycentric Julian Date, which uses the barycentre of the Solar System as its reference frame. Again, the Maciejewski et al. (2011) data points were produced in BJD, and thus required conversion to HJD before inclusion in the same O-C diagram. Over the timescale of all observations included in Fig. 6.2, BJD and HJD differ by $\sim \pm 2$s.

6.2 Other systems

6.2.1 WASP-10b

WASP-10b is a massive ($3M_J$) Hot Jupiter planet with an orbital period of 3.09d and a large transit depth, thanks to an apparently unusually dense K5 host star (Christian et al., 2009). The abnormal density of the host star make this an unusual system, so perhaps it comes as little surprise that the parameters have been through significant revision, with the discovery paper (Christian et al., 2009) apparently over-estimating the planetary radius by almost 30%, according to Johnson et al. (2009). The target therefore proves to be an interesting case for testing the modelling procedures employed in this thesis, specifically the effect of including follow-up photometry alongside the discovery SuperWASP photometry. The high-precision light curve (RMS $\sim$0.5mmag) from Johnson et al. (2009) is also included in the analysis, as the data is available from the NASA Exoplanet Archive[1].

6.2.1.1 Observations

Data on this object were taken on 11/12/2010, with PIRATE Mark 1.5. The observing run consisted of 168 frames each of 45s exposure time (resultant cadence of $\sim$1 minute)

[1] http://exoplanetarchive.ipac.caltech.edu/

taken in the Johnson R band, with the camera cooled to an operating temperature of -20 °C. Some light cirrus clouds were present, though they had little ill-effect on the quality of the photometry, as the RMS residuals from the model fits (described later in the subsection) are around 8 mmag, which is to be expected for a V=12.7 mag star (see Fig. 5.12). The light curve of WASP-10b for this night's data was produced in the usual way using the PIRATE data reduction pipeline as described in Chapter 3.

6.2.1.2 Modelling the photometry

The SuperWASP consortium code (Collier Cameron et al., 2007; Pollacco et al., 2008; Enoch et al., 2010) as described in Chapter 4 was employed to simultaneously fit to the SuperWASP observation data and the single, partial PIRATE observation. Two MCMC simulations were run, each consisting of a maximum burn-in time of 2000 steps (both took around 500 steps to burn-in), and a 12000 step chain post burn-in to evaluate the posterior probability distribution around the determined global solution. One of the simulations had the main-sequence mass radius prior applied (see Chapter 4), the other did not. The latest iteration of the SuperWASP consortium MCMC code removes stellar mass as a jump parameter, instead using the mass-radius calibration of Torres et al. (2010). Here, stellar mass is a no longer a jump parameter, replaced instead with T_{eff} and $[Fe/H]$, used in combination with $\log\rho_*$ (derived from the light curve fit) to determine the stellar mass. More detail can be found in Enoch et al. (2010), but consequently T_{eff} and $[Fe/H]$ must be initialised as proposal parameters. For all the MCMC simulations run in this section, T_{eff} was initialised as $T_{eff} = 4675^{+100}_{-100}$K , and $[Fe/H] = 0.03^{+0.2}_{-0.2}$, taken from the spectroscopic analysis in Christian et al. (2009). Alongside these two MCMC simulations on SuperWASP data plus PIRATE data, simulations were also run on SuperWASP data alone, and SuperWASP, PIRATE & the Johnson et al. (2009) Sloan z' band light curve, each with the main sequence prior on and off. Tables 6.4 and 6.5 contain the parameter estimates for all combinations of data source, with the main-sequence prior switched on and off for each data source combination. Table 6.4 contains the system parameters as reported in Christian et al. (2009) and Johnson et al. (2009), alongside the determined parameters on the MCMC runs on SuperWASP data alone. Table 6.5 contains all the combinations of SuperWASP and follow-up observations. The model light curves for all MCMC runs can be seen in Fig. 6.5.

Of principle interest is the ability to, with the main-sequence prior on, repeat the discovery parameters from Christian et al. (2009) (SuperWASP & PIRATE data);

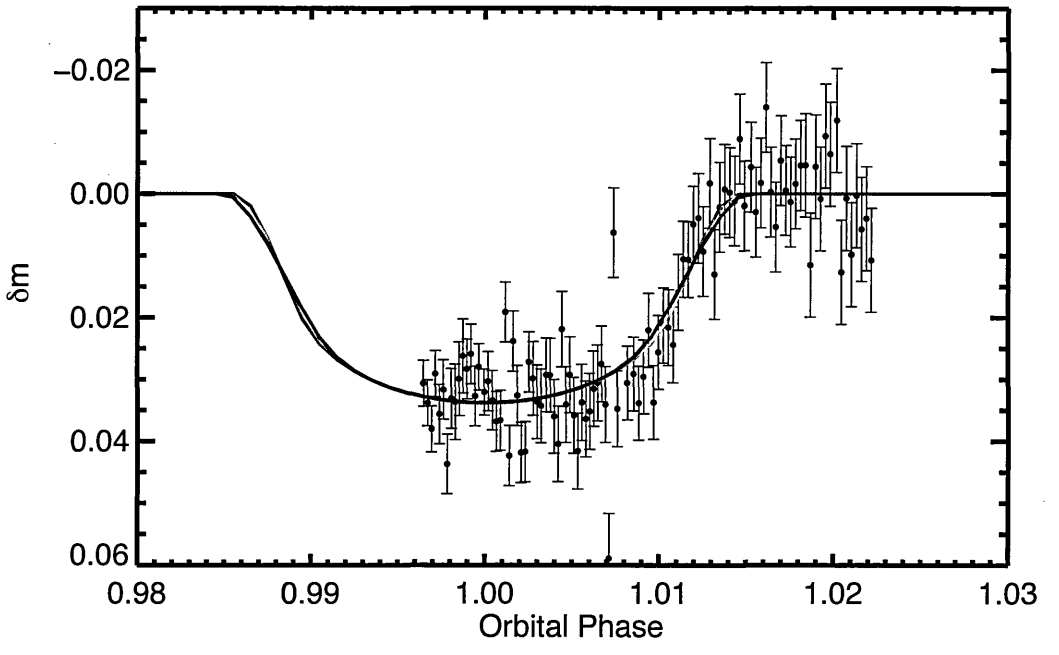

Figure 6.3: PIRATE Mark 1.5 observations in the R band of the transit of WASP-10b, taken on 11/12/2010. The two overlaid model light curves are from MCMC fits to the SuperWASP and PIRATE data alone. The blue model shows the fit with the main sequence mass-radius prior turned on, the red with it turned off.

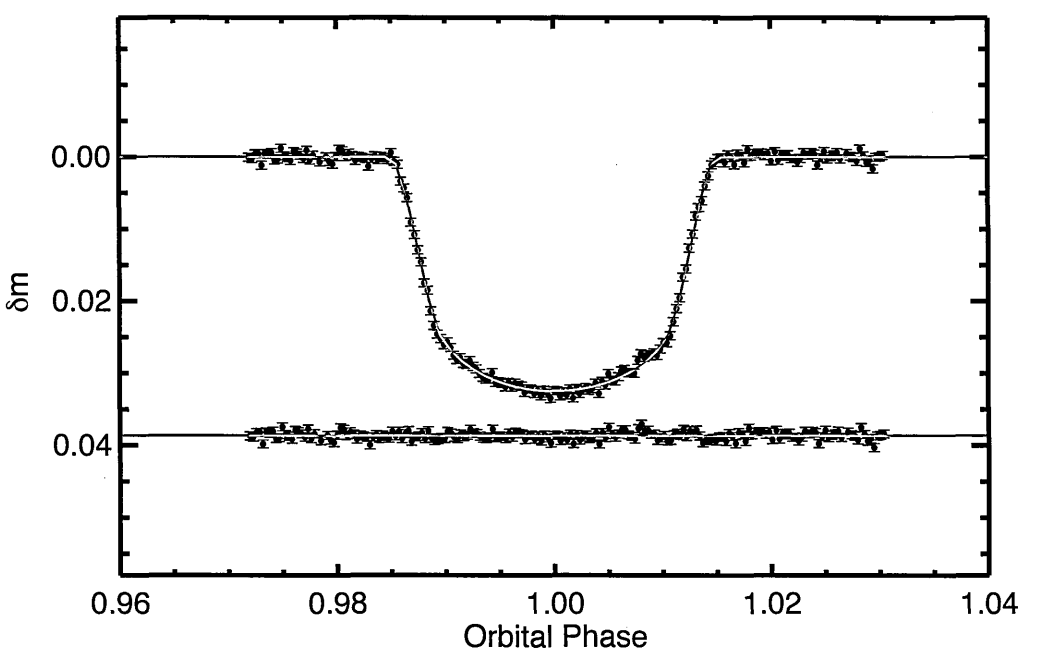

Figure 6.4: The high precision observations from Johnson et al. (2009), taken in the Sloan z' band. The model shown is that determined by MCMC fitting to the combined SuperWASP, PIRATE, and Johnson et al. (2009) data, with the main sequence mass-radius prior switched off.

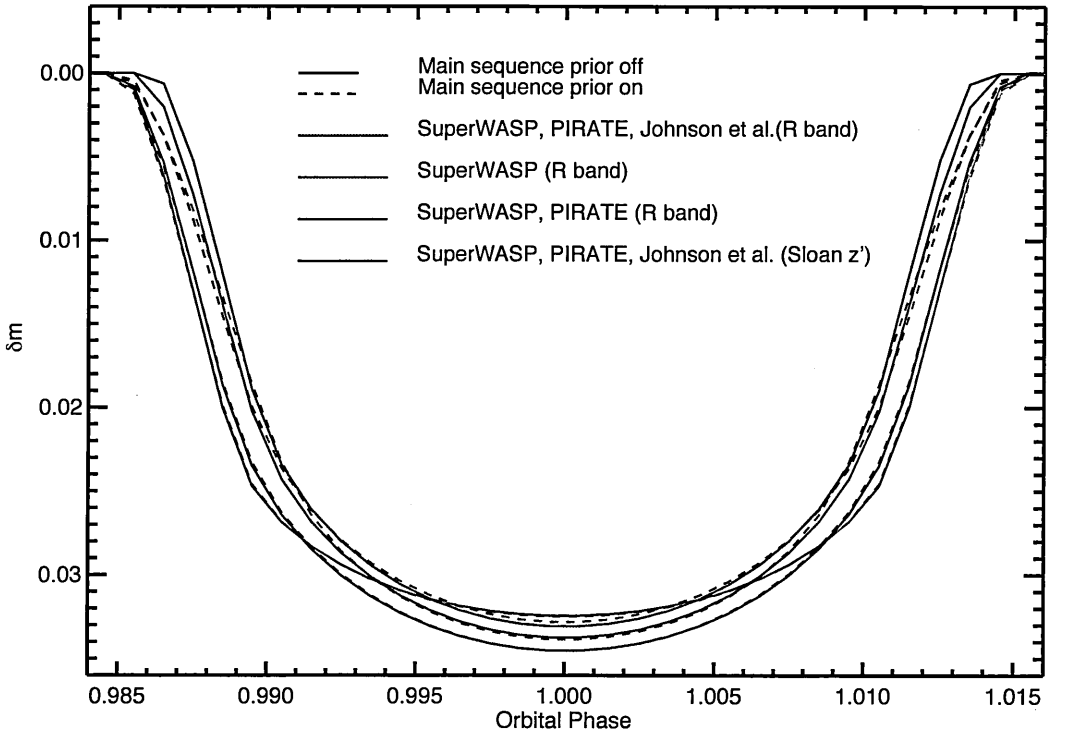

Figure 6.5: The model fits for all data combinations. The two red light curves fit the data for the Sloan z' band, whilst the rest are all R band model fits.

whilst also being able to recreate the parameters from the follow up paper by Johnson et al. (2009) with the main sequence prior off for this combination. The system is also of interest due to the unusual nature of the host star, which is low mass and active (Smith et al., 2009). Such a host star should, according to Torres et al. (2010) and Enoch et al. (2010), fail to have its mass and radius correctly determined by the T_{eff}, $[Fe/H]$ and $\log\rho_*$ calibration. Stars with $T_{eff} < 5000$K do not feature in the calibrating sample of binary systems employed in Torres et al. (2010), and thus the calibration may yield incorrect stellar radii.

As previously mentioned, with just SuperWASP & PIRATE data, it is possible to recreate both the discovery parameter estimations and a set of parameter estimations that largely agree with the revising follow-up paper. This seems to point to two things. Firstly, the main-sequence prior may be inappropriate for this particular system, as the host star is apparently far from the main sequence. The prior then acts to penalise solutions that lie further from the main sequence, pulling the solution back towards the main sequence, resulting in a lower measured stellar density (in the discovery paper, which employs a main sequence prior) of $1.51^{+0.25}_{-0.20}\rho_\odot$. This in turn leads to the reported larger planetary radius of $1.28R_J$. Secondly, the PIRATE data could prevent the solution from moving to larger stellar radii, but it is clearly not of sufficient SNR to 'combat' the main sequence prior. The same cannot be said of the exquisite Johnson et al. (2009) light curve taken with the Orthogonal Parallel Transfer Imaging Camera mounted on the UH 2.2m telescope, which achieves a light curve RMS of 0.5mmag and negligable correlated noise. As can be seen in Table 6.5, even with the main sequence prior switched on, the solution cannot be forced back to the main sequence due to the precision of the Sloan z' band light curve.

In conclusion, with relatively low SNR data, MCMC fitting can produce results that are heavily dependent on the priors adopted. When using modelling processes that employ Bayesian priors, it is imperative to examine the sensitivity of the results to the adopted priors. The best possible SNR photometry should be sought in order to achieve the most reliable parameter estimations.

Parameter	Christian et al. (2009)	Johnson et al. (2009)	SuperWASP data alone (M-S ON)	SuperWASP data alone (M-S OFF)
Transit Epoch (T_0) (HJD-2450000)	$4357.85808^{+0.00041}_{-0.00036}$	$4664.030913^{+0.000082}_{-0.000082}$	$3927.9661^{+0.0011}_{-0.0012}$	$3927.9655^{+0.0012}_{-0.0011}$
Orbital Period (days)	$3.0927636^{+0.0000094}_{-0.0000021}$	$3.0927616^{+0.0000112}_{-0.0000112}$	$3.0927227^{+0.0000062}_{-0.0000063}$	$3.0927198^{+0.0000061}_{-0.0000063}$
$(R_p/R_*)^2$	$0.029^{+0.001}_{-0.001}$	$0.02525^{+0.00024}_{-0.00028}$	$0.0253^{+0.0014}_{-0.0015}$	$0.0245^{+0.0013}_{-0.0011}$
Transit Width (days)	$0.098181^{+0.0019}_{-0.0015}$	$0.0928^{+0.0003}_{-0.0003}$	$0.0922^{+0.0023}_{-0.0022}$	$0.0856^{+0.0030}_{-0.0032}$
b (impact parameter)	$0.568^{+0.054}_{-0.084}$	$0.299^{+0.029}_{-0.042}$	$0.501^{+0.66}_{-0.13}$	$0.306^{+0.14}_{-0.10}$
Stellar mass ($M_\odot$)	$0.703^{+0.068}_{-0.080}$	$0.75^{+0.04}_{-0.03}$	$0.77^{+0.02}_{-0.02}$	$0.80^{+0.03}_{-0.03}$
Stellar radius ($R_\odot$)	$0.775^{+0.043}_{-0.040}$	$0.698^{+0.012}_{-0.012}$	$0.730^{+0.031}_{-0.030}$	$0.648^{+0.035}_{-0.030}$
Stellar density ($\rho_\odot$)	$1.51^{+0.25}_{-0.20}$	$2.19^{+0.06}_{-0.06}$	$1.98^{+0.23}_{-0.21}$	$2.93^{+0.45}_{-0.44}$
Orbital semi-major axis (a/AU)	$0.0369^{+0.0012}_{-0.0014}$	$0.0378^{+0.0003}_{-0.0004}$	$0.0381^{+0.0004}_{-0.0004}$	$0.0385^{+0.0004}_{-0.0004}$
Orbital inclination (i)(°)	$86.9^{+0.6}_{-0.5}$	$88.49^{+0.22}_{-0.17}$	$87.44^{+0.72}_{-0.08}$	$88.65^{+0.65}_{-0.75}$
Planetary radius (R_p/R_J)	$1.28^{+0.08}_{-0.09}$	$1.080^{+0.020}_{-0.020}$	$1.133^{+0.064}_{-0.076}$	$0.982^{+0.073}_{-0.050}$

Table 6.4: System parameters for WASP-10b as derived by the consortium MCMC code for a variety of different combinations of data. The first two columns contain the discovery parameters and subsequently revised parameters (Christian et al. (2009) and Johnson et al. (2009) respectively). The last two columns contain the derived parameters for SuperWASP data alone, with the main sequence mass-radius constraint both on and off.

Parameter	SuperWASP & PIRATE (R band) (M-S Prior ON)	SuperWASP & PIRATE (R band) (M-S Prior OFF)	SuperWASP, PIRATE (R band) & Johnson et al. (Sloan z') (M-S prior ON)	SuperWASP, PIRATE (R band) & Johnson et al. (Sloan z') (M-S prior OFF)
Transit Epoch (T_0) (HJD-2450000)	$4911.45665^{+0.00084}_{-0.00085}$	$4911.45627^{+0.00080}_{-0.00075}$	$4670.21648^{+0.00006}_{-0.00006}$	$4670.21647^{+0.00006}_{-0.00005}$
Orbital Period (days)	$3.09273601^{+0.0000036}_{-0.0000029}$	$3.09273415^{+0.0000030}_{-0.0000026}$	$3.09273953^{+0.0000023}_{-0.0000019}$	$3.09274001^{+0.0000022}_{-0.0000022}$
$(R_p/R_*)^2$	$0.0283^{+0.0018}_{-0.0015}$	$0.02513^{+0.00106}_{-0.00096}$	$0.02528^{+0.00034}_{-0.00019}$	$0.02507^{+0.00011}_{-0.00011}$
Transit Width (days)	$0.0923^{+0.0021}_{-0.0022}$	$0.0884^{+0.0021}_{-0.0023}$	$0.0929^{+0.0005}_{-0.0004}$	$0.0924^{+0.0002}_{-0.0002}$
b (impact parameter)	$0.618^{+0.055}_{-0,047}$	$0.376^{+0.075}_{-0.087}$	$0.312^{+0.059}_{-0.037}$	$0.262^{+0.016}_{-0.027}$
Stellar mass ($M_\odot$)	$0.786^{+0.025}_{-0.024}$	$0.797^{+0.026}_{-0.025}$	$0.770^{+0.024}_{-0.023}$	$0.786^{+0.022}_{-0.022}$
Stellar radius ($R_\odot$)	$0.781^{+0.032}_{-0.029}$	$0.679^{+0.025}_{-0.024}$	$0.692^{+0.021}_{-0.011}$	$0.685^{+0.007}_{-0.007}$
Stellar density ($\rho_\odot$)	$1.65^{+0.18}_{-0.17}$	$2.54^{+0.29}_{-0.26}$	$2.31^{+0.08}_{-0.14}$	$2.43^{+0.05}_{-0.03}$
Orbital semi-major axis (a/AU)	$0.0383^{+0.0004}_{-0.0004}$	$0.0385^{+0.0004}_{-0.0004}$	$0.0381^{+0.0004}_{-0.0004}$	$0.0383^{+0.0004}_{-0.0004}$
Orbital inclination (i)(°)	$86.65^{+0.35}_{-0.41}$	$88.24^{+0.44}_{-0.42}$	$88.49^{+0.19}_{-0.33}$	$88.75^{+0.14}_{-0.08}$
Planetary radius (R_p/R_J)	$1.28^{+0.08}_{-0.07}$	$1.05^{+0.05}_{-0.05}$	$1.071^{+0.039}_{-0.021}$	$1.056^{+0.013}_{-0.013}$

Table 6.5: MCMC parameter determinations for SuperWASP & follow-up photometry combined. The first two columns contain SuperWASP & PIRATE data (for both main-sequence prior on and off); the second two columns contain SuperWASP, PIRATE, and the Johnson et al. (2009) data.

6.2.2 HAT-P-20b

HAT-P-20b (Bakos et al., 2011) is a dense $7.25M_J$, $0.87R_J$ planet orbiting a $V = 11.339$ K3 dwarf star with a period $P = 2.875317 \pm 0.000004$ d and epoch $T_0 =$

2455080.92661 ± 0.00021 (BJD_{UTC}).

This target was observed with PIRATE Mark 1.5 as part of the commissioning phase on the night beginning 2nd November 2010. Initial test observations at high airmass (> 2) of the object produced a peak pixel value for the target of 27000 ADU for a 30s exposure. Given that the target was rising, the decision was made to test the process of deliberately defocusing the images. This spreads the PSF out, preventing saturation from occuring even for longer exposures of 60s (see Fig. 6.6). The duty cycle (ratio of time spent with the shutter open to the total duration of observing run) is also improved by defocusing. However, it is not without its problems, as described later. FocusMax was used to find a best focus position of 17224μ for the microfocuser. This was then bumped up a further 1mm to 18224μ for the observing run. A sequence of 212 science frames of 60s were taken in the R band, with the camera cooled to an operating temperature of -20 °C. The effective cadence was therefore $\sim$ 75s. The target was observed whilst rising, from an airmass of 2.115, decreasing throughout the entire run to 1.038.

This output from the PIRATE pipeline for this observation run can be seen at the top of Fig. 6.7, binned to $\phi = 0.001$. Overlaid is the model fit from a combined SuperWASP data & PIRATE data MCMC fit using the standard consortium code. As can be seen, the fit model apparently underestimates the amplitude of the flux deficit, and thus underestimates (R_P/R_*). The bottom of the same figure shows the SuperWASP data for the object also binned to $\phi = 0.001$. The model apparently *overestimates* the flux deficit for the SuperWASP data. The answer for the discrepancy between both data sets lies in their respective spatial resolution, and the presence of nearby companions to the target star. Fig. 6.6 shows a DSS image of HAT-P-20b in the top left, with the SuperWASP 2.5, 3.5 and 4.5 pixel radius measuring apertures overlaid. The image in the top right of this Figure is image number 140 in the PIRATE observation sequence. The two images are not equally scaled. The effect of the defocusing is readily apparent in the PIRATE image, with all the PSFs showing the signature 'doughnut' profile of a defocused image. The bottom of Fig. 6.6 shows a surface plot of a 20 pixel × 20 pixel square surrounding the HAT-P-20 PSF; the square centred on image pixel [514, 511].

It is evident from the DSS image that the target star suffers from dilution by nearby fainter companions, which all contribute to the total flux within the aperture and consequently reduce the observed transit depth. PIRATE also suffers from dilution by the very close companion whose PSF overlaps that of the target star. The close proximity of the companion in the PIRATE image coupled with the use of

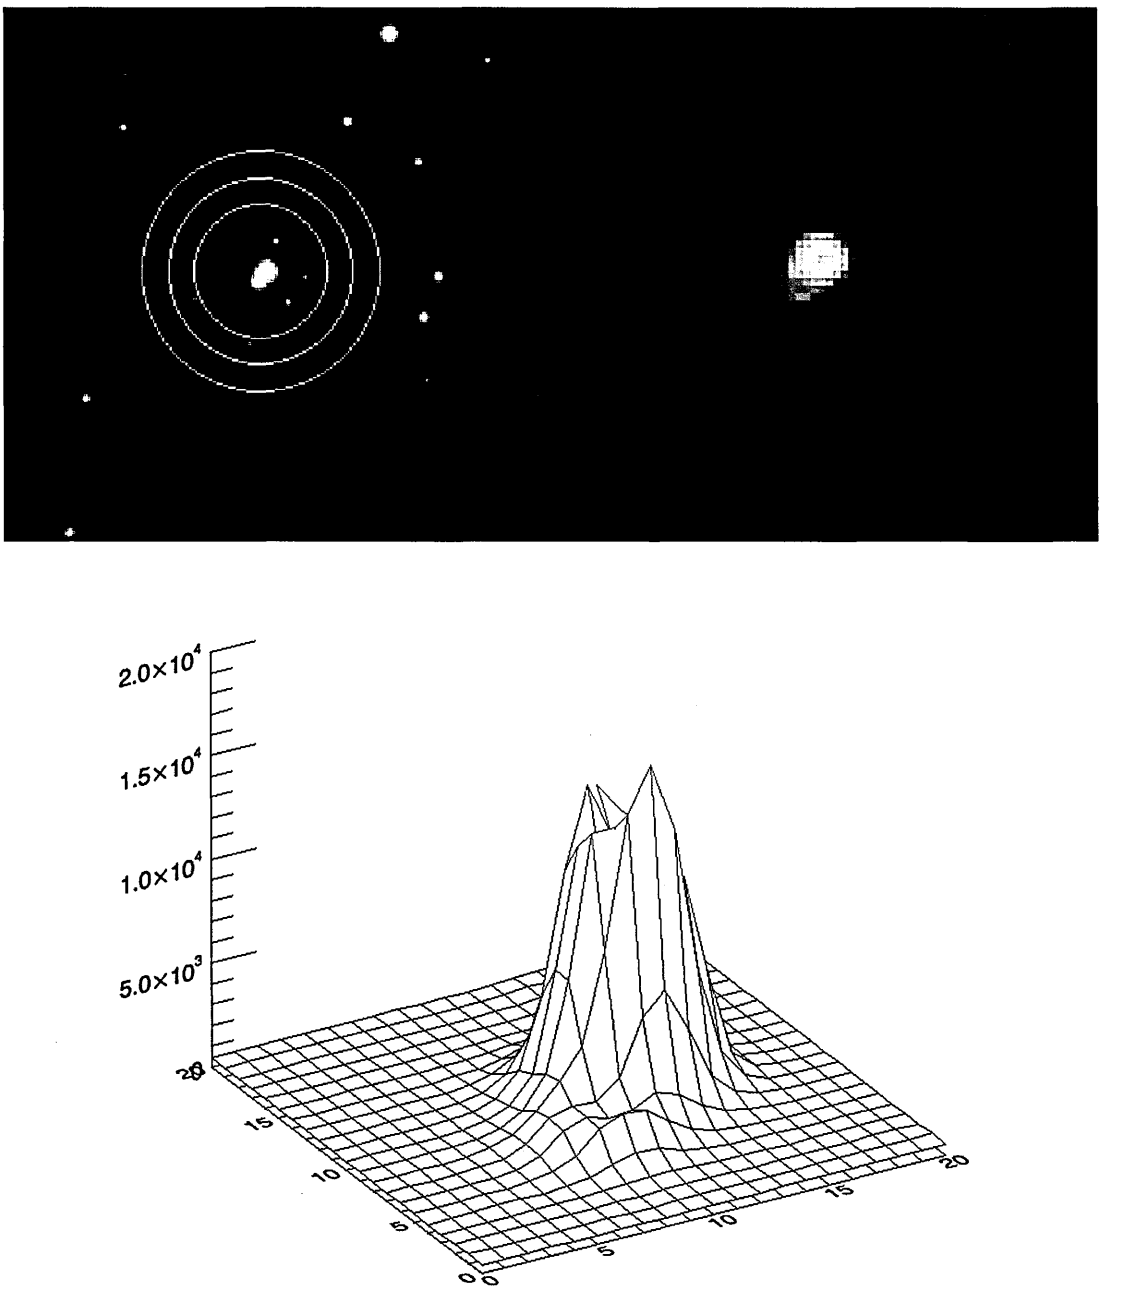

Figure 6.6: *Top Left:* HAT-P-20b in a DSS image with the 2.5, 3.5 and 4.5 pixel SuperWASP apertures overlaid. *Top Right:* defocused PIRATE image of the same target. The two images are not comparably scaled, this figure is just for illustrative purposes. *Bottom:* Wire-mesh surface image of the PIRATE HAT-P-20b PSF. The faint companion is visible in the foreground. The peak pixel value in this contour plot is 17083 ADU. The surface plot is centred on image pixel coordinates [514, 511].

defocusing result in the dilution of the PIRATE transit light curve as well, though to a less extent. A measuring aperture of radius 11 pixels was used on the PIRATE photometry, meaning the flux of the overlapping companion was included in the measuring aperture, but not all of the surrounding companions that are included in the SuperWASP apertures. This explains the difference in flux deficit depths between the PIRATE and SuperWASP datasets, and the apparent understimation in ΔF and $(^{R_p}/_{R_*})$. The derived model parameters are not included here due to their invalidity in correctly fitting the transit depth of the target; so they are expect to be (and are) different from the discovery paper parameters.

6.2.3 TrES-2b, TrES-3b, XO-1b

Three other known transiting systems were observed with PIRATE Mark 1 and 1.5, but no modelling work was done on these, so the light curves are simply presented along with the details of the observations.

6.2.3.1 TrES-2b

TrES-2b (O'Donovan et al., 2006) is a 1.28 M_J mass planet orbiting a $V = 11.41$, G0 V, 1.08 $M_\odot$ host star every 2.47063 days. The transit duration is short, only ~1.6 hours, which, in conjunction with the 2.47 d period, implies a high impact parameter transit; confirmed by the shape of the transit which has a very short duration between contacts two and three. O'Donovan et al. (2006) measure a value for the impact parameter of $b = 0.84 \pm 0.02$.O'Donovan et al. (2006) observe a transit depth of $\sim$ 1.4 %, and calculate a radius for the planet of 1.24 R_J.

This target was observed on 7th September, 2010 with PIRATE Mark 1.5, and was the first transiting planet observation to be made with the new CDK17 OTA (see Chapter 2). 148 V band images each of 30s exposure time were taken as a test sequence to assess performance whilst auto-guiding and pointing were still being calibrated and fine-tuned during the post-installation period. The target was observed initially at airmass 1.019, culminating at 1.014, and finishing at airmass 1.45. The resulting output from the PIRATE pipeline can be seen at the top of Fig. 6.8. As can be seen, the light curve apparently suffers from significant correlated noise which is attributed to the sub-optimal auto-guiding in use. However, the expected features of the light curve are apparent, including the 'V-shape' of a high impact parameter, and the $\sim$ 1.4% transit depth.

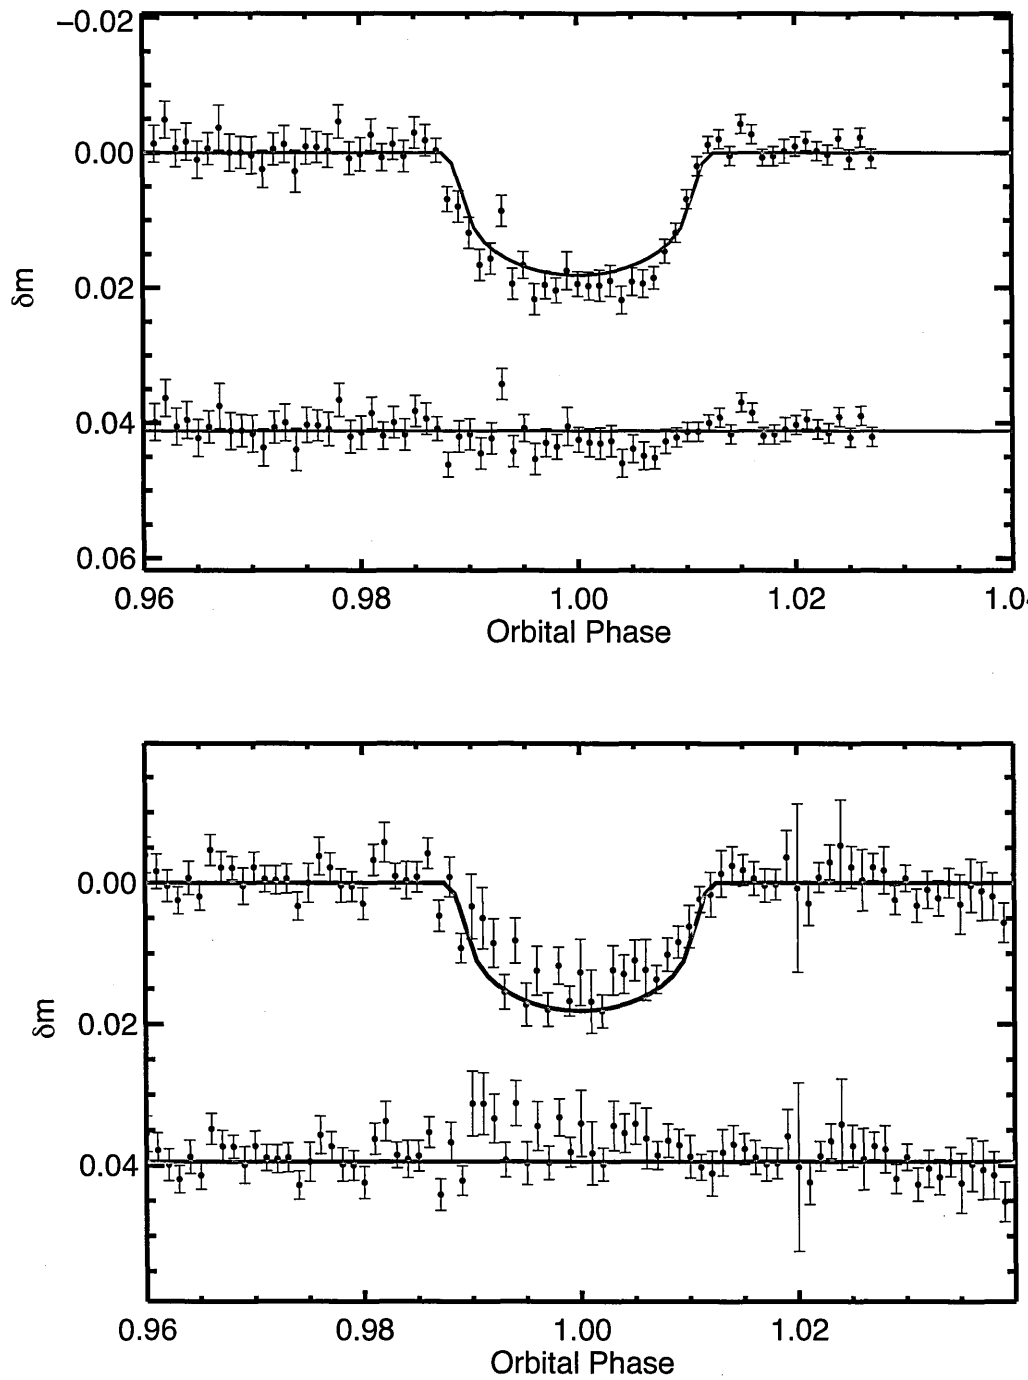

Figure 6.7: PIRATE follow-up observation (top) and SuperWASP data (bottom) of HAT-P-20b, both binned to $\phi = 0.001$

6.2.3.2 TrES-3b

TrES-3b (O'Donovan et al., 2007) is a 1.91 M_J planet (Sozzetti et al., 2009) orbiting a $V = 12.40$ G dwarf host star with a period of 1.30619 days. This system also has a short transit duration of ~ 1.4 hours, and thus a high impact parameter of 0.84 ± 0.01 (Sozzetti et al., 2009). The planet is of a similar size to TrES-2b, with a radius of 1.336 R_J, but the system has a smaller host star with $R_* = 0.83 R_\odot$, yielding a greater transit depth of $\sim 2.7\%$.

This target was observed on 8th April 2011 with PIRATE Mark 1.5, when a sequence of 156 frames each of 60s exposure time were taken in the R band. The target was observed whilst rising, between airmasses 1.55 and 1.01. The light curve generated by the PIRATE pipeline can be seen in the middle plot of Fig. 6.8, where the distinctive high impact parameter features and 2.7% transit depth of the system are recovered well.

6.2.3.3 XO-1b

XO-1b (McCullough et al., 2006) is a $M_p = 0.90$ M_J, $R_p = 1.30 R_J$ planet orbiting a Solar analogue G1 V, $V = 11.3$, $M_* = 1.0 M_\odot$, $R_* = 1.0 R_\odot$ host star every 3.94153 days

This target was the first transiting extrasolar planet observation made with the PIRATE facility, on 17^{th} March 2009. It was made with PIRATE Mark 1, and the data were first reduced manually, without the use of the PIRATE pipeline. Due to a lack of decent comparison stars in the field of view, a single star comparison yielded the light curve least affected by systematics. This observation was the first to clearly highlight the issues faced with the 'Light Curve Discontinuities' (LCDs) discussed in Chapter 2. As can be seen, soon after fourth contact the mount pier flips, resulting a drop in out of transit flux of 1.5-2%, roughly equivalent to that seen in the transit itself. This observation led to further study of the LCD problem, discussed in Chapter 2.

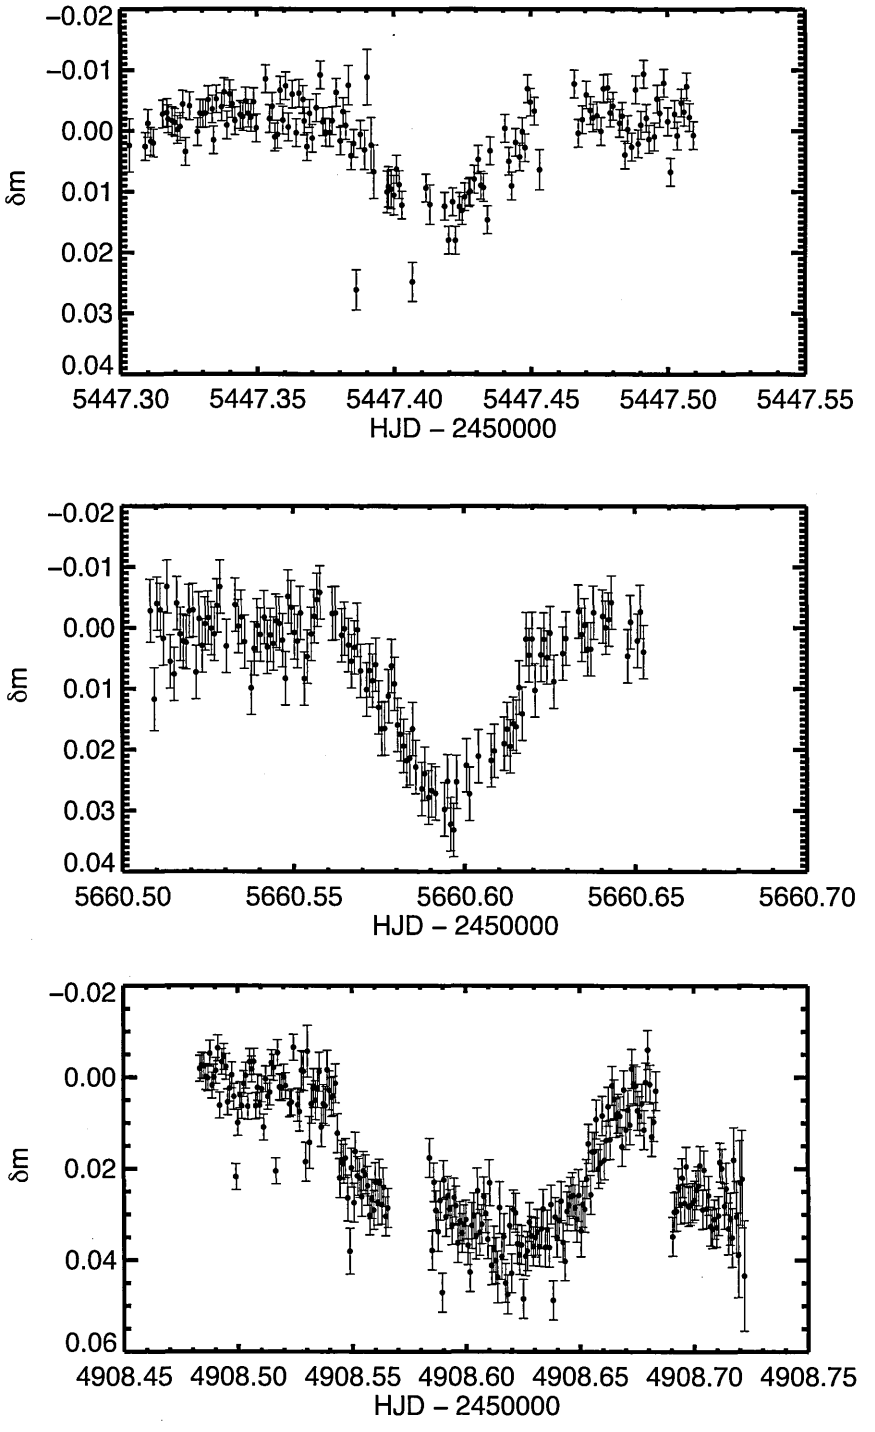

Figure 6.8: PIRATE observations of TrES-2b (07/09/2010), TrES-3b (08/04/2011), and XO-1b (17/03/2009).

Chapter 7

Discussion & Conclusions

7.1 Success of the PIRATE facility

The PIRATE facility has undergone great strides of progress since it started life in March 2008. In that time it has been transformed from a fledgling, nascent bundle of possibility to a fully accomplished, semi-autonomous remote facility, now routinely delivering useful scientific results. It has been used to make observations of possible planets, discoveries of novae in other galaxies, and even a fortuitous observation of a supernova that helped constrain the radius (and therefore nature) of the exploding progenitor system (Bloom et al., 2012). The aim of this PhD project was always to develop the research capacity of the facility; to turn it into a efficacious instrument for performing astronomy research in reach of a small facility such as PIRATE, and I feel that this has been suitably accomplished. A principle goal of this development was to allow for a huge increase in the volume of observations possible with the facility; extending the range of conditions under which observations are possible, and increasing the safety of the facility and the ease with which observations are made. The success of SuperWASP follow-up observations made in the Summer of 2011, and the way in which they were obtained is strong evidence that the development program has been a success. These observations were obtained with the facility running in a semi-autonomous mode; where user-submitted plans were used to initialise a night's data collection, and automate the start-up procedure - but the facility took over the monitoring and execution of these plans, finally shutting itself down before sunrise and leaving the user with the simple task of collecting the data the following day. This brought a great increase in the scientific yield and productivity of the instrument, and

this knowledge and experience has been carried through to the latest iteration of the facility, PIRATE Mark 2.

Alongside the development of the research capacity of the facility, much success has also been had in installing PIRATE as a leader in the field of remote astronomy education. This began with the first presentation of S382 'Astrophysics' in April 2010, and has expanded to include a new, shorter module for for the second level general practical science course SXP288. In all instances of its educational role, PIRATE is used remotely by a group of students coming together online, communicating over VOIP (Voice Over IP), and working 'live' on the telescope through the night to obtain real data. In its educational guise, PIRATE's autonomous functionality is switched off in order to teach as much as possible about the operational aspects of a telescope facility to the students. For the recent presentation of SXP288, the efficiency of learning how to use the facility was boosted via the introduction of a PIRATE simulator, which I set up in the run-up to the module. The simulator runs on a virtual machine and allows users to log in to a virtual representation of PIRATE, use the same operation software, and even take and download simulated images, allowing the students to learn how to use the facility without actually operating the hardware. Clearly, the students are gaining a lot of vital experience in using the facility. The data taken in some of the S382 projects has been of high quality, leading some of the students to publish their findings, see Rodda et al. (2012), Faillace et al. (2012).

The facility is now firmly in its next operational phase with a new camera, Mark 2, and is continuing to make a strong contribution to research projects and to the field of remote higher eduation. A recent addition is the introduction of ACP Scheduler[1] which will, in the future, be used to gain the ability of interrupting scheduled follow-up observations in response to received VOEvent messages that indicate current outbursts of optical transients. Such a facility will add another string to PIRATE's bow, and stand it in good stead for the future.

7.2 Improving the PIRATE facility & follow-up program

This thesis focuses entirely on PIRATE Mark 1 and Mark 1.5, mostly as the transition to PIRATE Mark 2 in September 2011 brought about an entirely new and

[1] http://scheduler.dc3.com

much improved camera in the form of the SBIG STX-16803. Due to the extent of the commissioning work that would be required for the PIRATE Mark 2 set-up, the transition from Mark 1.5 to 2 presented a natural 'line in the sand' breakpoint, and it was quickly obvious that this thesis should reference the new set-up, but not contain any data from the new instrument due to time constraints. The new camera represents a substantial upgrade over the STL-1001E, bringing with it a larger field of view, and a significantly improved plate scale (see Chapter 2). This brings with it added benefits for the follow-up program. Firstly, the larger field of view introduces more comparison stars available for use with the target star, which the ensemble compilation routine will benefit from. Secondly, the smaller plate scale will ensure that point-spread function fitting techniques (not employed in this thesis as attempts to use them for PIRATE Mark 1.5 were unsuccessful) should function correctly. With PIRATE Mark 1.5's plate scale of $1.69''\text{pixel}^{-1}$, if seeing conditions were good, the point spread function was often undersampled (the PSF consisted of too few pixels) which interfered with the success of PSF fitting procedures. With PIRATE Mark 2's plate scale of $0.634''\text{pixel}^{-1}$, this is no longer a problem. PSF fitting is particularly useful in the case of blending from close companions which have overlapped PSF functions, which aperture photometry cannot distinctly resolve. Having the ability to successfully PSF fit the data will provide PIRATE with greater opportunity to remove trickier blend scenarios from the SuperWASP candidate lists.

Whilst the new camera brings with it great improvements, throughout my time working on PIRATE Mark 1 and Mark 1.5 I have also identified a number of areas in which the facility and the follow-up program could be improved in the future, which are broken down here in no particular order of significance or magnitude:

7.2.1 Areas for future work

7.2.1.1 Further automation

With the recent introduction of Scheduler to the facility, a step further has been made towards full automation of the facility. Scheduler provides a crucial piece of functionality previously missing from running the system in the absence of scheduler: the ability to have PIRATE automatically restart observation plans after a weather (or other) interrupt. Previously, a weather interrupt would see the weather script compiled and run by ACP, interrupting and stopping the image acquisition script, as scripts cannot be run concurrently by ACP. The weather shut down would be

enforced, but no monitoring of weather conditions for improvement would occur; this would be entirely the responsibility of the user. However, with Scheduler, persistent monitoring of conditions does occur, and if conditions improve, the image acquisition script is restarted. This is a significant boost to the follow-up program as automatic monitoring of conditions and restarting of the imaging schedule for a given night makes the system more efficient, losing less time to unnecessary down time where the user has failed to spot that conditions have improved sufficiently.

However, even with Scheduler in place, it is still not possible to run the follow-up program in a fully robotic and autonomous fashion. Scheduler is a 'dispatch scheduler'. It holds a database of targets and the details of the observations required for the target, and then chooses the best target to be observing at a given time. With Scheduler running, the system can run in a fully robotic mode: so long as a list of targets is input to the database, the full operation (start-up, shutdown, target acquisition and image acquisition, weather monitoring and safety shutdowns) is handled on a daily basis by Scheduler, so the user has even fewer attendance requirements. However, a crucial limitation of this mode of operation is the lack of appreciation of time-dependence for a given observation of a target. If, say, a 300 exposure image sequence in the R band of a SuperWASP candidate is required, Scheduler would observe the target when it was most convenient according the target's rise and set times, distance from the moon etc. This would of course miss the point of the observation, which is the need to observe it according to the predicted ephemeris. Scheduler has no way (as of yet) to include this extra temporal constraint for periodic variables. Discussion was entered into with ACP and Scheduler's author, Robert Denny, about how this could be achieved, but no satisfactory solutions currently exist, bar the significant modifcation of the underlying codebase.

7.2.1.2 Forcing a plate solve for every image acquired

This feature was previously not implemented for PIRATE Mark 1 and Mark 1.5, but has since been implemented for PIRATE Mark 2. This setting forces a plate solve attempt for every final image taken by ACP. With this feature turned off, ACP would make a limited number of attempts to plate solve a particular image before giving up on plate solving the image, and turning off plate solving for all subsequent images in the sequence. This would allow for any errors in pointing to go unchecked, and also removes some of the functionality of the image, by stripping its FITS header of useful plate solve parameters. On the downside, forcing plate solving in poor conditions

increases the duration of a successful plate solution, impacting cadence negatively. However, this is usually of little concern if conditions are already unsuitable, as the quality of the photometry will be poor.

Forcing all images to be plate solved ensures all images will have the parameters of a plate solution stored in the FITS headers, which allows this information to be extracted and put to good use by the pipeline. Further more, if a plate solution is not attainable, this usually indicates that conditions have deteriorated and barely any stars are visible in the image; so time spent trying to plate solve is not good observing time lost. The frames without a plate solution could then be automatically rejected by the pipeline.

7.2.1.3 Improvements to the follow-up program

Though the SuperWASP candidate follow-up observation program proved successful, with a total of 57 candidates observed and categorised (with even more taken with PIRATE Mark 2 in September 2011, though not included in this thesis for cohesion's sake), there are areas in which the follow-up program may be improved.

Firstly, 13 out of the 57 candidates were observed with no detectable signal. This could have been mitigated by setting strict bounds on the acceptable level of ephemeris uncertainty. Incorporating the method of visualising the ephemeris drift (see 5.9) as a simple tool involved in the observation planning sequence would also be of benefit. This would reduce the amount of telescope time spent hunting for transit signals for which there is little chance of observing. Further to this, restricting candidates by expected transit depth is also advised, in conjunction with Fig. 5.13; ensuring that only targets with a chance of retrieving a 6σ detection or better (for example) in a single observation are observed. Alternatively, shallower transits may be repeatedly observed, but only if the ephemeris is suitably well constrained.

Another area for possible expansion is the investigation of multi-colour follow-up photometry. The filter wheel offers the ability to run multi-colour light curve sequences, interleaving filters throughout the sequence. This allows for observations in which a colour index, for example $(B-V)$, can be observed as a function of time, under the assumption that any two consecutive observations (each in a different filter) are essentially coincidental, and can thus be convolved temporally as a single measurement. This could prove to be a valuable follow-up tool as planetary transits are largely 'colourless', i.e. the colour of the measured starlight does not change with time (this is not strictly true, they actually form a minimal 'double-horned' profile

- see Tingley, 2004). Conversely, binary star transits do undergo a change in colour throughout the transit event, assuming the stars in the binary system have a different effective temperature. Such a technique may therefore be used to distinguish the transits of low mass dwarf stars from those of planets. More work would be required to assess the suitability of such a technique to the PIRATE hardware; specifically with regard to the size of the color signal $\Delta(B-V)$ in typical low mass eclipsing binary systems, and its detectability with PIRATE.

7.2.1.4 Improvements to the pipeline

The following is a list of possible improvements to the PIRATE pipeline, not implemented due to time constraints:

- FWHM from FITS headers. ACP employs the Pinpoint software to plate solve each image taken. In doing so, it should write the mean FWHM of the stars PSF-fitted in the process into the FITS header. However, this feature appears not to work for unknown reasons. If this worked, the section of the pipeline that requires the user to determine and then input the FWHM of the master frame could be automated by simply pulling the figure from the FITS header. In the case that it is not possible to fix the issue of incorrectly recorded FWHM values in the FITS headers, another possible solution exists to automate this part of the pipeline procedure. An initial DAOFIND search could be employed with an arbitrary (but plausible) FWHM input value (to be employed in star-finding procedure). Even if this guessed value is incorrect, the process would still return the locations of bright sources. This information could then be fed back into IDL and the GAUSS2DFIT routine could then be found to measure the actual FWHM value once the source positions were known.

- BJD instead of HJD. HJD was adopted for the pipeline to ensure congruity with the SuperWASP dataset, which also employs HJD. However, BJD is now seen as a superior time system (due to using the barycentre of the solar system, rather than the centre of the Sun) and more transit ephemerides are being publised in $\mathrm{BJD_{TDB}}$.

- Incorporation of the full CCD equation uncertainties. The full CCD equation (see Eq. 3.1) is preferable to the cut-down version employed in uncertainty calculations by IRAF; which holds for bright sources but is less accurate for sources that are just above background level.

- Automatic frame rejection by condition of plate solution. This works in conjunction with forcing all images to be plate solved by ACP. In the situation where a frame has not been successfully plate-solved, this can be determined with a simple check in IDL and the frame can be rejected from the image sequence automatically.

- Use of more than one aperture size. Currently the pipeline adopts a single aperture size for performing photometry; a multiple of the FWHM. In order to optimise the photometry further, a range of aperture sizes should be attempted for all frames, and the aperture size that minimises the RMS of the resultant light curves should be employed.

- Greater tolerance for cosmic ray hits. As it currently stands, only stars which have a photometric measurement in all frames left after the frame rejection cuts (manual, and the cuts made by the photometric failure threshold parameter, eq. 3.5) pass through to the final stages of the pipeline. This means that any star that is otherwise fine throughout a run that is unlucky enough to suffer a cosmic ray strike within the bounds of the measuring aperture (thus sending the struck pixel to the saturation point) will be needlessly struck from the list of stars. Tolerance could be built into to allow one or two bad frames to pass, keeping the star in the final list of stars.

7.2.2 A new method for determining the optimum comparison ensemble

An alternative, more advanced method for optimally compiling comparison ensembles was briefly investigated by the author but not fully tested, and so it is detailed in this chapter more as a possible route for future work, rather than a report on a successfully completed investigation.

The basic principle is similar to the optimal ensemble compilation routine described in chapter 3, in that comparison stars are selected for inclusion in the ensemble according to a statistical figure of merit, though the determination of the figure of merit is more complicated in the proposed new version.

Instead, each differential light curve between any pair of stars is split into numerous segments of equal size. Within each segment, a linear fit is made to the data within the segment, and the fit parameters are stored. Alongside this, the following

statistics are recorded for each segment: the RMS deviation of the data within the segment, the residual RMS (post linear fit) and the median value of the data in the segment. Together with these statistics for the individual segments in a light curve, overall light curve statistics are recorded, such as the RMS of the segments' median values, the RMS of the gradients of all of the linear fits for the light curve segments, and the median and RMS of the light curve segment residuals.

Using these segments, one can overcome a limitation of the current formulation of the ensemble compilation routine, namely its susceptibility to suppressing variability for large amplitude variables in certain situations. In the current iteration, faint (and therefore noisy) comparison stars may be selected in preference to comparison stars with a good SNR as their RMS may prove less than that of a strong SNR comparison star that reveals the true amplitude of the target's variability. With the new version, the segment residuals could be used to simultaneously find good comparison stars for large amplitude variables, and can then be used in conjunction with the overall light curve statistics to even flag variable stars within PIRATE's field.

Some initial prototyping of this concept proved promising, with the routine successfully flagging variable stars artificially injected into a pre-existing dataset, whilst also providing sensible stars for inclusion in the comparison ensembles for the artificial large-amplitude variables. Unfortunately, the code was unwieldy and took a long time to run, and the project was de-prioritised and then terminated, but this still remains a promising avenue for further work.

7.3 Conclusion

PIRATE is a remotely-operable telescope facility built primarily from readily available 'off-the-shelf' components, used in both education and research. In its education role, it has successfully supported small groups of simultaneous student users in their efforts to observe and classify variable stars in the SuperWASP archive. It has been shown that relatively inexpensive equipment can play a useful role in modern astrophysical research, including high-precision time-series photometry required for transiting Jupiter-size exoplanets. The PIRATE system employs a GEM which provides exceptional stability and all-sky pointing accuracy for its cost. The pier-flip introduces LCDs which are difficult to correct by calibration procedures alone; the best strategy is to treat pre- and post-flip light curves as separate observing runs.

The use of a remotely-operable telescope in the OU module S382, and subsequenty the second level module S288 proved the feasibility of deploying such complex hardware in real-time to create an inspirational teaching tool for distance education. The great enthusiasm of the students involved, and the quality of the acquired data and of the scientific reports generated at the end of the project demonstrate the success of the PIRATE teaching project. Use by level 1 (year 1) OU students is currently under consideration.

A comprehensive follow-up program was established in which many observations were made, and the tools developed to make the routine follow-up of possible planets found in the SuperWASP dataset a simple and swift process. This program continues to the day. Other research work includes the regular monitoring of M31 for a team at the Max Planck Institute for Extraterrestrial Physics, who use the data in the hunt for novae, and the unexpected delight of assisting in pinning down the nature of the progenitor of supernova SN2011fe (Bloom et al., 2012).

I am very proud to have been instrumental in helping to make the PIRATE the accomplished facility that it is today, and hope that it continues to bounce from success to further success in the future.

Bibliography

G. Á. Bakos, J. Hartman, G. Torres, D. W. Latham, G. Kovács, R. W. Noyes, D. A. Fischer, J. A. Johnson, G. W. Marcy, A. W. Howard, D. Kipping, G. A. Esquerdo, A. Shporer, B. Béky, L. A. Buchhave, G. Perumpilly, M. Everett, D. D. Sasselov, R. P. Stefanik, J. Lázár, I. Papp, and P. Sári. Hat-p-20b-hat-p-23b: Four massive transiting extrasolar planets. *ApJ*, 742:116, December 2011. doi: 10.1088/0004-637X/742/2/116.

C. Beaugé, S. Ferraz-Mello, and T. A. Michtchenko. *Planetary Masses and Orbital Parameters from Radial Velocity Measurements, in Extrasolar Planets: Formation, Detection and Dynamics.* Wiley-VCH Verlag GmbH & Co., 2007.

J. S. Bloom, D. Kasen, K. J. Shen, P. E. Nugent, N. R. Butler, M. L. Graham, D. A. Howell, U. Kolb, S. Holmes, C. A. Haswell, V. Burwitze, J. Rodriguez, and M. Sullivan. A compact degenerate primary-star progenitor of sn 2011fe. *ApJL*, 744:L17, January 2012. doi: 10.1088/2041-8205/744/2/L17.

J. Bochinski, R. Busuttil, S. Holmes, C. A. Haswell, and Kolb U. Exploring the treasure trove: Pirate as a semi-robotic exoplanets winnower. Poster presented at RAS-NAM 2012, Manchester, March 2012.

I. A. Bond, A. Udalski, M. Jaroszyński, N. J. Rattenbury, B. Paczyński, I. Soszyński, L. Wyrzykowski, M. K. Szymański, M. Kubiak, O. Szewczyk, K. Żebruń, G. Pietrzyński, F. Abe, D. P. Bennett, S. Eguchi, Y. Furuta, J. B. Hearnshaw, K. Kamiya, P. M. Kilmartin, Y. Kurata, K. Masuda, Y. Matsubara, Y. Muraki, S. Noda, K. Okajima, T. Sako, T. Sekiguchi, D. J. Sullivan, T. Sumi, P. J. Tristram, T. Yanagisawa, P. C. M. Yock, and OGLE Collaboration. Ogle 2003-blg-235/moa 2003-blg-53: A planetary microlensing event. *ApJL*, 606:L155–L158, May 2004. doi: 10.1086/420928.

R. P. Boyle. A method to correct ccd flatfields for precise stellar photometry. *Baltic Astronomy*, 16:459–466, 2007.

T. Brown. Expected detection and false alarm rates for transiting jovian planets. *ApJ*, 593:125–128, 2003.

T.M. Brown, D. Charbonneau, R. L. Gilliland, R. W. Noyes, and A. Burrows. Hubble Space Telescope Time-Series Photometry of the Transiting Planet of HD 209458. *ApJ*, 552:699–709, May 2001. doi: 10.1086/320580.

C. J. Burke, B. S. Gaudi, D. L. DePoy, and R. W. Pogge. Survey for transiting extrasolar planets in stellar systems. iii. a limit survey for transiting extrasolar planets in stellar systems. iii. a limit on the fraction of stars with planets in the open cluster ngc 1245. *AJ*, 132:210, 2006.

T. Chan, M. Ingemyr, J. N. Winn, M. J. Holman, R. Sanchis-Ojeda, G. Esquerdo, and M. Everett. The transit light-curve project. xiv. confirmation of anomalous radii for the exoplanets tres-4b, hat-p-3b, and wasp-12b. *ApJ*, 141:179, June 2011. doi: 10.1088/0004-6256/141/6/179.

D. Charbonneau, T. M. Brown, D. W. Latham, and M. Mayor. Detection of planetary transits across a sun-like star. *ApJ*, 529:L45–L48, 2000.

D. Charbonneau, T. M. Brown, R. W. Noyes, and R. L. Gilliland. Detection of an extrasolar planet atmosphere. *ApJ*, 568:377–384, 2002.

D. Charbonneau, T. M. Brown, E. W. Dunham, D. W. Latham, D. L. Looper, and G. Mandushev. Astrophysical false positives encountered in wide-field transit searches. In *THE SEARCH FOR OTHER WORLDS: Fourteenth Astrophysics Conference*, volume 713, pages 151–160, 2004.

D. J. Christian, D. L. Pollacco, I. Skillen, R. A. Street, F. P. Keenan, W. I. Clarkson, A. Collier Cameron, S. R. Kane, T. A. Lister, R. G. West, B. Enoch, A. Evans, A. Fitzsimmons, C. A. Haswell, C. Hellier, S. T. Hodgkin, K. Horne, J. Irwin, A. J. Norton, J. Osborne, R. Ryans, P. J. Wheatley, and D. M Wilson. The superwasp wide-field exoplanetary transit survey: candidates from fields 23 h $<$ ra $<$ 03 h. *MNRAS*, 372:1117–1128, November 2006. doi: 10.1111/j.1365-2966.2006.10913.x.

D. J. Christian, N. P. Gibson, E. K. Simpson, R. A. Street, I. Skillen, D. Pollacco, A. Collier Cameron, Y. C. Joshi, F. P. Keenan, H. C. Stempels, C. A. Haswell,

K. Horne, D. R. Anderson, S. Bentley, F. Bouchy, W. I. Clarkson, B. Enoch, L. Hebb, G. Hébrard, C. Hellier, J. Irwin, S. R. Kane, T. A. Lister, B. Loeillet, P. Maxted, M. Mayor, I. McDonald, C. Moutou, A. J. Norton, N. Parley, F. Pont, D. Queloz, R. Ryans, B. Smalley, A. M. S. Smith, I. Todd, S. Udry, R. G. West, P. J. Wheatley, and D. M. Wilson. Wasp-10b: a $3m_J$, gas-giant planet transiting a late-type k star. *MNRAS*, 392:1585–1590, February 2009. doi: 10.1111/j.1365-2966.2008.14164.x.

F. R. Chromey and D. A. Hasselbacher. The flat sky: Calibration and background uniformity in wide-field astronomical images. *PASP*, 108:944–949, 1996.

A. Claret. A new non-linear limb-darkening law for lte stellar atmosphere models. *A&A*, 363:1081–1190, November 2000.

A. Collier Cameron, D. Pollacco, R. A. Street, T. A. Lister, R. G. West, D. M. Wilson, F. Pont, D. J. Christian, W. I. Clarkson, B. Enoch, A. Evans, A. Fitzsimmons, C. A. Haswell, C. Hellier, S. T. Hodgkin, K. Horne, J. Irwin, S. R. Kane, F. P. Keenan, A. J. Norton, N. R. Parley, J. Osborne, R. Ryans, I. Skillen, and P. J. Wheatley. A fast hybrid algorithm for exoplanetary transit searches. *MNRAS*, 373:799–810, December 2006. doi: 10.1111/j.1365-2966.2006.11074.x.

A. Collier Cameron, D. M. Wilson, R. G. West, L. Hebb, X.-B. Wang, S. Aigrain, F. Bouchy, D. J. Christian, W. I. Clarkson, B. Enoch, M. Esposito, E. Guenther, C. A. Haswell, G. Hébrard, C. Hellier, K. Horne, J. Irwin, S. R. Kane, B. Loeillet, T. A. Lister, P. Maxted, M. Mayor, C. Moutou, N. Parley, D. Pollacco, F. Pont, D. Queloz, R. Ryans, I. Skillen, R. A. Street, S. Udry, and P. J. Wheatley. Efficient identification of exoplanetary transit candidates from superwasp light curves. *MNRAS*, 380:1230–1244, September 2007. doi: 10.1111/j.1365-2966.2007.12195.x.

R. Denny. A web-remote/robotic/scheduled astronomical data acquisition system. In *Telescopes from Afar*, Hawaii, February 2011.

J. Eastman, R. Siverd, and B. S. Gaudi. Achieving better than 1 minute accuracy in the heliocentric and barycentric julian dates. *PASP*, 122:935–946, August 2010. doi: 10.1086/655938.

A. Einstein. Lens-like action of a star by the deviation of light in the gravitational field. *Science*, 84:506–507, December 1936. doi: 10.1126/science.84.2188.506.

B. Enoch, A. Collier Cameron, N. R. Parley, and L. Hebb. An improved method for estimating the masses of stars with transiting planets. *A&A*, 516:A33, 2010. doi: 10.1051/0004-6361/201014326.

G. Faillace, C. Owen, D. Pulley, and D. Smith. Investigating the properties of the near contact binary system tw crb. *ArXiv e-prints*, April 2012.

L. Fossati, C. A. Haswell, C. S. Froning, L. Hebb, S. Holmes, U. Kolb, C. Helling, A. Carter, P. Wheatley, A. Collier Cameron, B. Loeillet, D. Pollacco, R. Street, H. C. Stempels, E. Simpson, S. Udry, Y. C. Joshi, R. G. West, I. Skillen, and D. Wilson. Metals in the exosphere of the highly irradiated planet wasp-12b. *ApJL*, 714:L222–L227, May 2010. doi: 10.1088/2041-8205/714/2/L222.

R. L. Gilliland, T. M. Brown, H. Kjeldsen, J. K. McCarthy, M. L. Peri, J. A. Belmonte, I. Vidal, L. E. Cram, J. Palmer, S. Frandsen, M. Parthasarathy, L. Petro, H. Schneider, P. B. Stetson, and W. W. Weiss. A search for solar-like oscillations in the stars of m67 with ccd ensemble photometry on a network of 4 m telescopes. *AJ*, 106:2441–2476, December 1993. doi: 10.1086/116814.

A. D. Grauer, A. W. Neely, and C. H. S. Lacy. Building an automated telescope with high photometric accuracy. *PASP*, 120:992–1000, September 2008. doi: 10.1086/591808.

P. Greenfield and R. L. White. A new cl for iraf based on python. In & D. Crabtree N. Manset, C. Veillet, editor, *Astronomical Data Analysis Software and Systems IX*, volume 216 of *Astronomical Society of the Pacific Conference Series*, page 59, 2000.

C. A. Haswell. *Transiting Exoplanets.* Cambridge University Press, 2010.

L. Hebb, A. Collier-Cameron, B. Loeillet, D. Pollacco, G. Hébrard, R. A. Street, F. Bouchy, H. C. Stempels, C. Moutou, E. Simpson, S. Udry, Y. C. Joshi, R. G. West, I. Skillen, D. M. Wilson, I. McDonald, N. P. Gibson, S. Aigrain, D. R. Anderson, C. R. Benn, D. J. Christian, B. Enoch, C. A. Haswell, C. Hellier, K. Horne, J. Irwin, T. A. Lister, P. Maxted, M. Mayor, A. J. Norton, N. Parley, F. Pont, D. Queloz, B. Smalley, and P. J. Wheatley. Wasp-12b: The hottest transiting extrasolar planet yet discovered. *ApJ*, 693:1920–1928, March 2009. doi: 10.1088/0004-637X/693/2/1920.

C. Hellier, D. R. Anderson, A. Collier Cameron, M. Gillon, L. Hebb, P. F. L. Maxted, D. Queloz, B. Smalley, A. H. M. J. Triaud, R. G. West, D. M. Wilson, S. J. Bentley, B. Enoch, K. Horne, J. Irwin, T. A. Lister, M. Mayor, N. Parley, F. Pepe, D. L. Pollacco, D. Segransan, S. Udry, and P. J. Wheatley. An orbital period of 0.94days for the hot-jupiter planet wasp-18b. *Nature*, 460:1098–1100, August 2009. doi: 10.1038/nature08245.

C. Hellier, D. R. Anderson, A. Collier Cameron, M. Gillon, M. Lendl, P. F. L. Maxted, D. Queloz, B. Smalley, A. H. M. J. Triaud, R. G. West, D. J. A. Brown, B. Enoch, T. A. Lister, F. Pepe, D. Pollacco, D. Ségransan, and S. Udry. Wasp-29b: A saturn-sized transiting exoplanet. *ApJL*, 723:L60–L63, November 2010. doi: 10.1088/2041-8205/723/1/L60.

J. L. Hershey. Astrometric analysis of the field of ac +65 6955 from plates taken with the sproul 24-inch refractor. *AJ*, 78:421 – 425, 1973.

E. Høg, C. Fabricius, V. V. Makarov, S. Urban, T. Corbin, G. Wycoff, U. Bastian, P. Schwekendiek, and A. Wicenec. The tycho-2 catalogue of the 2.5 million brightest stars. *A&A*, 355:L27–L30, March 2000.

S. Holmes, U. Kolb, C. A. Haswell, V. Burwitz, R. J. Lucas, J. Rodriguez, S. M. Rolfe, J. Rostron, and J. Barker. Pirate: A remotely operable telescope facility for research and education. *PASP*, 123:1177–1187, October 2011. doi: 10.1086/662148.

K. Horne. Planetary transit searches: Hot jupiters galore. In W. R. F. Dent, editor, *Techniques for the Detection of Planets and Life beyond the Solar System*, page 5, 2001.

S. B. Howell. Two-dimensional aperture photometry: Signal-to-noise ratio of point-source observations and optimal data-extraction techniques. *PASP*, 101:616, 1989.

S. B. Howell. *Handbook of CCD Astronomy.* Cambridge University Press, 2nd edition, 2006.

E. Jehin, M. Gillon, D. Queloz, P. Magain, J. Manfroid, V. Chantry, M. Lendl, D. Hutsemékers, and S. Udry. Trappist: Transiting planets and planetesimals small telescope. *The Messenger*, 145:2–6, September 2011.

J. A. Johnson, J. N. Winn, N. E. Cabrera, and J. A. Carter. A smaller radius for the transiting exoplanet wasp-10b. *ApJL*, 692:L100–L104, February 2009. doi: 10.1088/0004-637X/692/2/L100.

S. R. Kane, K. Horne, T. Lister, A. Collier Cameron, R. A. Street, D. L. Pollacco, D. James, and Y. Tsapras. Recent results from the wide angle search for planets (wasp) prototype. In & C. Terquem J. Beaulieu, A. Lecavelier Des Etangs, editor, *Extrasolar Planets: Today and Tomorrow*, volume 321 of *Astronomical Society of the Pacific Conference Series*, page 115, December 2004.

U. Kolb, R. J. Lucas, V. Burwitz, S. Holmes, C. A. Haswell, J. Rodriguez, J Rolfe, S. M.and Rostron, and J. Barker. Pirate – the picetl robotic astronomical telescope explorer, 2010. URL `http://pirate.open.ac.uk/publications.html`.

V. Kornilov, M. Sarazin, A. Tokovinin, T. Travouillon, and O. Voziakova. Comparison of the scintillation noise above different observatories measured with mass instruments. *A&A*, 546:A41, October 2012. doi: 10.1051/0004-6361/201219954.

G. Kovács, S. Zucker, and T. Mazeh. A box-fitting algorithm in the search for periodic transits. *A&A*, 391:369–377, August 2002. doi: 10.1051/0004-6361:20020802.

D. Lafrenière, R. Jayawardhana, and M. H. van Kerkwijk. Direct Imaging and Spectroscopy of a Planetary-Mass Candidate Companion to a Young Solar Analog. *ApJL*, 689:L153–L156, December 2008. doi: 10.1086/595870.

W. B. Landsman. The idl astronomy user's library. In & J. Barnes R. J. Hanisch, R. J. V. Brissenden, editor, *Astronomical Data Analysis Software and Systems II*, volume 52 of *Astronomical Society of the Pacific Conference Series*, page 246, January 1993.

B. M. Lasker, C. R. Sturch, C. Lopez, A. D. Mallamas, S. F. McLaughlin, J. L. Russell, W. Z. Wisniewski, B. A. Gillespie, H. Jenkner, E. D. Siciliano, D. Kenny, J. H. Baumert, A. M. Goldberg, G. W. Henry, E. Kemper, and M. J. Siegel. The guide star photometric catalog. *ApJS*, 68:1–90, September 1988. doi: 10.1086/191282.

S. L. Li, N. Miller, D. N. C. Lin, and J. J. Fortney. Wasp-12b as a prolate, inflated and disrupting planet from tidal dissipation. *Nature*, 463:1054–1056, February 2010. doi: 10.1038/nature08715.

R. Lucas and U. Kolb. Software architecture for an unattended remotely controlled telescope. *ArXiv e-prints*, December 2010.

G. Maciejewski, R. Errmann, S. Raetz, M. Seeliger, I. Spaleniak, and R. Neuhäuser. High-precision photometry of wasp-12 b transits. *A&A*, 528:A65, April 2011. doi: 10.1051/0004-6361/201016268.

K. Mandel and E. Agol. Analytic Light Curves for Planetary Transit Searches. *ApJL*, 580:L171–L175, December 2002. doi: 10.1086/345520.

J. Manfroid. Stellar calibration of ccd flat fielding. *A&AS*, 113:587, November 1995.

M. Mayor and D. Queloz. A jupiter-mass companion to a solar-type star. *Nature*, 378:355–359, 1995.

P. R. McCullough, J. E. Stys, J. A. Valenti, C. M. Johns-Krull, K. A. Janes, J. N. Heasley, B. A. Bye, C. Dodd, S. W. Fleming, A. Pinnick, R. Bissinger, B. L. Gary, P. J. Howell, and T. Vanmunster. A transiting planet of a sun-like star. *ApJ*, 648: 1228–1238, September 2006. doi: 10.1086/505651.

W. Merline and S. B. Howell. A realistic model for point-sources imaged on array detectors: The model and initial results. *Experimental Astronomy*, 6:163–210, January 1995.

D. G. Monet, S. E. Levine, B. Canzian, H. D. Ables, A. R. Bird, C. C. Dahn, H. H. Guetter, H. C. Harris, A. A. Henden, S. K. Leggett, H. F. Levison, C. B. Luginbuhl, J. Martini, A. K. B. Monet, J. A. Munn, J. R. Pier, A. R. Rhodes, B. Riepe, S. Sell, R. C. Stone, F. J. Vrba, R.L. Walker, G. Westerhout, R. J. Brucato, I. N. Reid, W. Schoening, M. Hartley, M. A. Read, and S. B. Tritton. The usno-b catalog. *AJ*, 125:984–993, February 2003. doi: 10.1086/345888.

A. J. Norton, P. J. Wheatley, R. G. West, C. A. Haswell, R. A. Street, A. Collier Cameron, D. J. Christian, W. I. Clarkson, B. Enoch, M. Gallaway, C. Hellier, K. Horne, J. Irwin, S. R. Kane, T. A. Lister, J. P. Nicholas, N. Parley, D. Pollacco, R. Ryans, I. Skillen, and D. M. Wilson. New periodic variable stars coincident with ROSAT sources discovered using SuperWASP. *A&A*, 467:785–905, May 2007. doi: 10.1051/0004-6361:20077084.

F. T. O'Donovan, D. Charbonneau, G. Mandushev, E. W. Dunham, D. W. Latham, G. Torres, A. Sozzetti, T. M. Brown, J. T. Trauger, J. A. Belmonte, M. Rabus, J. M. Almenara, R. Alonso, H. J. Deeg, G. A. Esquerdo, E. E. Falco, L. A. Hillenbrand, A. Roussanova, R. P. Stefanik, and J. N. Winn. Tres-2: The first transiting planet in the kepler field. *ApJL*, 651:L61–L64, November 2006. doi: 10.1086/509123.

F. T. O'Donovan, D. Charbonneau, G. Á. Bakos, G. Mandushev, E. W. Dunham, T. M. Brown, D. W. Latham, G. Torres, A. Sozzetti, G. Kovács, M. E. Everett, N. Baliber, M. G. Hidas, G. A. Esquerdo, M. Rabus, H. J. Deeg, J. A. Belmonte,

L. A. Hillenbrand, and R. P. Stefanik. Tres-3: A nearby, massive, transiting hot jupiter in a 31 hour orbit. *ApJL*, 663:L37–L40, July 2007. doi: 10.1086/519793.

D. Pollacco, I. Skillen, A. Collier Cameron, B. Loeillet, H. C. Stempels, F. Bouchy, N. P. Gibson, L. Hebb, G. Hébrard, Y. C. Joshi, I. McDonald, B. Smalley, A. M. S. Smith, R. A. Street, S. Udry, R. G. West, D. M. Wilson, P. J. Wheatley, S. Aigrain, K. Alsubai, C. R. Benn, V. A. Bruce, D. J. Christian, W. I. Clarkson, B. Enoch, A. Evans, A. Fitzsimmons, C. A. Haswell, C. Hellier, S. Hickey, S. T. Hodgkin, K. Horne, M. Hrudková, J. Irwin, S. R. Kane, F. P. Keenan, T. A. Lister, P. Maxted, M. Mayor, C. Moutou, A. J. Norton, J. P. Osborne, N. Parley, F. Pont, D. Queloz, R. Ryans, and E. Simpson. WASP-3b: a strongly irradiated transiting gas-giant planet. *MNRAS*, 385:1576–1584, April 2008. doi: 10.1111/j.1365-2966.2008.12939.x.

D. L. Pollacco, I. Skillen, A. Collier Cameron, D. J. Christian, C. Hellier, J. Irwin, T. A. Lister, R. A. Street, R. G. West, D. Anderson, W. I. Clarkson, H. Deeg, B. Enoch, A. Evans, A. Fitzsimmons, C. A. Haswell, S. Hodgkin, K. Horne, S. R. Kane, F. P. Keenan, P. F. L. Maxted, A. J. Norton, J. Osborne, N. R. Parley, R. S. I. Ryans, B. Smalley, P. J. Wheatley, and D. M. Wilson. The WASP Project and the SuperWASP Cameras. *PASP*, 118:1407–1418, October 2006. doi: 10.1086/508556.

F. Pont, S. Zucker, and D. Queloz. The effect of red noise on planetary transit detection. *MNRAS*, 373:231–242, November 2006. doi: 10.1111/j.1365-2966.2006.11012.x.

W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical recipes in FORTRAN. The art of scientific computing.* 1992.

D. Queloz, M. Mayor, L. Weber, A. Blécha, M. Burnet, B. Confino, D. Naef, F. Pepe, N. Santos, and S. Udry. The coralie survey for southern extra-solar planets. i. a planet orbiting the star gliese 86. *A&A*, 354:99–102, February 2000.

N. Regnault, A. Conley, J. Guy, M. Sullivan, J.-C. Cuillandre, P. Astier, C. Balland, S. Basa, R. G. Carlberg, D. Fouchez, D. Hardin, I. M. Hook, D. A. Howell, R. Pain, K. Perrett, and C. J. Pritchet. Photometric calibration of the supernova legacy survey fields. *A&A*, 506:999–1042, November 2009. doi: 10.1051/0004-6361/200912446.

T. Rodda, A. Bruce, S. Cruickshank, and M. Salisbury. An investigation of gsc 02038-00293, a suspected rs cvn star, using ccd photometry. *ArXiv e-prints*, June 2012.

P. D. Sackett. Searching for unseen planets via occultation and microlensing. In *Planets Outside the Solar System: Theory and Observations*, page 189, 1999.

S. Seager and G. Mallén-Ornelas. A unique solution of planet and star parameters from an extrasolar planet transit light curve. *ApJ*, 585:1038–1055, March 2003. doi: 10.1086/346105.

F. J. Selman. Photometric flats: an essential ingredient for photometry with wide-field imagers. In P. J. Quinn & A. Bridger, editor, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 5493 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 453–459, September 2004. doi: 10.1117/12.551357.

A. M. S. Smith. The impact of correlated noise on superwasp detection rates for transiting extrasolar planets. *MNRAS*, 373:1151–1158, 2006.

A. M. S. Smith, L. Hebb, A. Collier Cameron, D. R. Anderson, T. A. Lister, C. Hellier, D. Pollacco, D. Queloz, I. Skillen, and R. G. West. A superwasp search for additional transiting planets in 24 known systems. *MNRAS*, 398:1827–1834, October 2009. doi: 10.1111/j.1365-2966.2009.15262.x.

A. Sozzetti, G. Torres, D. Charbonneau, J. N. Winn, S. G. Korzennik, M. J. Holman, D. W. Latham, J. B. Laird, J. Fernandez, F. T. O'Donovan, G. Mandushev, E. Dunham, M. E. Everett, G. A. Esquerdo, M. Rabus, J. A. Belmonte, H. J. Deeg, T. N. Brown, M. G. Hidas, and N. Baliber. A new spectroscopic and photometric analysis of the transiting planet systems tres-3 and tres-4. *ApJ*, 691:1145–1158, February 2009. doi: 10.1088/0004-637X/691/2/1145.

P. B. Stetson. Daophot - a computer program for crowded-field stellar photometry. *ASP*, 99:191, 1987.

O. Struve. Proposal for a project of high-precision stellar radial velocity work. *The Observatory*, 72:199–200, October 1952.

O. Tamuz, T. Mazeh, and S. Zucker. Correcting systematic effects in a large set of photometric light curves. *MNRAS*, 356:1466–1470, 2005.

M. Tegmark, M. A. Strauss, M. R. Blanton, K. Abazajian, S. Dodelson, H. Sandvik, X. Wang, D. H. Weinberg, I. Zehavi, N. A. Bahcall, F. Hoyle, D. Schlegel, R. Scoccimarro, M. S. Vogeley, A. Berlind, T. Budavari, A. Connolly, D. J. Eisenstein, D. Finkbeiner, J. A. Frieman, J. E. Gunn, L. Hui, B. Jain, D. Johnston, S. Kent, H. Lin, R. Nakajima, R. C. Nichol, J. P. Ostriker, A. Pope, R. Scranton, U. Seljak, R. K. Sheth, A. Stebbins, A. S. Szalay, I. Szapudi, Y. Xu, J. Annis, J. Brinkmann, S. Burles, F. J. Castander, I. Csabai, J. Loveday, M. Doi, M. Fukugita, B. Gillespie, G. Hennessy, D. W. Hogg, Ž. Ivezić, G. R. Knapp, D. Q. Lamb, B. C. Lee, R. H. Lupton, T. A. McKay, P. Kunszt, J. A. Munn, L. O'Connell, J. Peoples, J. R. Pier, M. Richmond, C. Rockosi, D. P. Schneider, C. Stoughton, D. L. Tucker, D. E. vanden Berk, B. Yanny, and D. G. York. Cosmological parameters from sdss and wmap. *Phys. Rev. D*, 69(10):103501, May 2004. doi: 10.1103/PhysRevD.69.103501.

B. Tingley. Using color photometry to separate transiting exoplanets from false positives. *A&A*, 425:1125–1131, October 2004. doi: 10.1051/0004-6361:20035792.

G. Torres, J. Andersen, and A. Giménez. Accurate masses and radii of normal stars: modern results and applications. *A&A Rev.*, 18:67–126, February 2010. doi: 10.1007/s00159-009-0025-1.

A. Udalski, O. Szewczyk, K. Zebrun, G. Pietrzynski, M. Szymanski, M. Kubiak, I. Soszynski, and L. Wyrzykowski. The optical gravitational lensing experiment. planetary and low-luminosity object transits in the carina fields of the galactic disk. *Acta Astron.*, 52:317–359, December 2002.

P. van de Kamp. Parallax, proper motion, acceleration, and orbital motion of barnard's star. *AJ*, 74:238, 1969.

& Brady S. Weber, L. Fast auto-focus method and software for ccd-based telescopes. 2001.

J. N. Winn, M. J. Holman, G. W. Henry, G. Torres, D. Fischer, J. A. Johnson, G. W. Marcy, A. Shporer, and T. Mazeh. The transit light curve project. x. a christmas transit of hd 17156b. *ApJ*, 693:794–803, March 2009. doi: 10.1088/0004-637X/693/1/794.

A. Wolszczan and D. A. Frail. A planetary system around the millisecond pulsar psr1257 + 12. *Nature*, 355:145–147, 1992.

J. T. Wright, O. Fakhouri, G. W. Marcy, E. Han, Y. Feng, J. A. Johnson, A. W. Howard, D. A. Fischer, J. A. Valenti, J. Anderson, and N. Piskunov. The exoplanet orbit database. *PASP*, 123:412–422, April 2011. doi: 10.1086/659427.

N. Zacharias, D. G. Monet, S. E. Levine, S. E. Urban, R. Gaume, and G. L. Wycoff. Nomad catalog (zacharias+ 2005). *VizieR Online Data Catalog*, 1297:0, November 2005.

R. E. Zissell. Shutter mapping correction for short ccd exposures. *JAAVSO*, 28: 149–156, June 2000.

Appendix A

Proprietary code

A.1 PIRATE IDL routines

Table A.1 contains brief descriptions of the main IDL routines of the PIRATE data reduction pipeline. All routines in the table were written by the author as part of this thesis. The pipeline makes use of standard IRAF photometry routines, as well as routines from the NASA IDL Library (Landsman, 1993).

A.2 ACP Weather Script

Algorithm A.1 is the modified weather shutdown script used by ACP, written in Javascript:

Routine	Description
doirafphot.pro	This routine acts a scripting engine for basic IRAF aperture photometry tasks performed in DAOPHOT. The routine is largely automated but asks the user for input on two occasions. Responsible for handling the preparation of calibration frames, the calibration of all science frames in the folder, suggesting a master frame to the user from which to generate a master catalog of stars for analysis, generating coordinate files for all frames in the time series, and finally running the 'PHOT' task in DAOPHOT for all stars in the master catalog across all frames. Returns an output photometry file for each science frame in the folder.
lcmaker.pro	The master routine for light curve generation. Calls each of readallmag, timeandtranspose, ensemble3, fluxdivisionsauto, investigateblends in turn. Must be pointed at the output of doirafphot.
readallmag.pro	Collates each of the individual (per-frame) photometry output files and creates a data cube for use by the pipeline.
timeandtranspose.pro	Converts UT timestamps to GJD (Geocentric Julian Day) and transposes the data cube so that it is star-centric and not frame-centric.
ensemble3.pro	Operates on the data cube, generating the optimal ensemble light curve for each star in the master list, using the method detailed in Chapter 3.
fluxdivisionsauto.pro	For the top N brightest stars in the master list takes the flux ratio of each star against all the other stars - for inspection purposes.
investigateblends.pro	Allows the user to check for nearby blend candidates by outputting the lightcurves of all objects within a fixed radius of the target object. See §3.2.1.3 for more detail.
mastercoo.pro	Is used by doirafphot to generate the master coordinate file
readincol.pro	Reads in columnar data into an IDL array
findtarginimage.pro	Uses the recorded target coordinates and the WCS data stored in the fits header of a file to visually identify the location of the target star in the image. Used by doirafphot.

Table A.1: List of the main routines in the PIRATE IDL data reduction pipeline.

Algorithm A.1 The weather shutdown script (written in javascript) modified from the default version (which just parked the system, (i.e. close down and do a full a shutdown)). The modification was made so the script simply closes the dome (if it was open), and waits with all other components on and connected, to allow for quicker restart upon resumption of fair weather.

```
1  /* This is a weather shut-down script for PIRATE.  */
2
3  function main()
4          {
5           if (Dome.Shutterstatus==1)
6                  {
7                  Console.Printline("**Weather safety script initiated**");
8                  Console.PrintLine("**Baader Planetarium AllSky Dome is already closed
                       **");
9                  }
10          else if (Dome.Shutterstatus == 0)
11                 {
12                 Console.Printline("**Weather safety script initiated**");
13                 Console.Printline("**Baader Planetarium AllSky Dome is open.**");
14                 Console.Printline("**Weather is unsafe, closing dome**");
15
16                 Dome.CloseShutter();
17                 Util.WaitForMilliseconds(30000);
18
19                 if (Dome.Shutterstatus == 1) {
20                 Console.Printline("**Baader Planetarium AllSky Dome now closed.");
21                                                }
22                 }
23         }
```

Appendix B

List of Acronyms

ACP - Observatory control software (it is not clear what ACP actually stands for)

ADU - Analog-to-Digital Unit

BEB - Blended Eclipsing Binary

BJD - Barycentric Julian Date

BLS - Box Least-Square

CCD - Charge Coupled Device

EB - Eclipsing Binary

FITS - Flexible Image Transport System

FOV - Field of View

FWHM - Full-Width Half-Maximum

GEM - German Equatorial Mount

HFD - Half Flux Diameter

HJD - Heliocentric Julian Date

IDL - Interactive Data Language

LCD - Light Curve Discontinuity

MCMC - Markov-Chain Monte Carlo

ND - Non-Detection

OAM - Observatori Astronomic de Mallorca

OTA - Optical Tube Assembly

PEC - Periodic Error Correction

PID - Position Intercept Difference

PIRATE - Physics Innovations Robotic Astronomical Telescope Explorer

PSF - Point Spread Function

RH - Relative Humidity

RMS - Root Mean Square (Deviation)

RV - Radial Velocity

SNR - Signal-to-Noise Ratio

TAI - International Atomic Time

TEP - Transing Extrasolar Planet

TDB - Barycentric Dynamical Time

TT - Terrestal Time

TTV - Transit Timing Variation

WASP - Wide Angle Search for Planets

WCS - World Coordinate System

UT - Universal Time

UTC - Coordinated Universal Time