

Planetary transits and stellar variability

Suzanne Aigrain

Institute of Astronomy & Corpus Christi College, Cambridge

A dissertation submitted to the University of Cambridge
for the degree of Doctor of Philosophy

January 26, 2005

The picture on the previous page is a composite of images taken during the transit of Venus on 8 June 2004 with the Solar X-ray imager (SXI) on the GOES-12 spacecraft. It was produced by the US National Oceanographic and Atmospheric Administration (NOAA) and downloaded from www.solarviews.com. It points to one of the main difficulties in detecting transits of a terrestrial planet: the intrinsic variability of late-type stars, a major part of which is due to magnetically induced photospheric features at the base of the X-ray bright coronal loops visible here.

Summary

Most of the 130 or so exo-planets (planets orbiting a star other than the Sun) known to date, were detected via the radial velocity (RV) method, which relies on the spectroscopic detection of changes in the parent star's radial velocity as it orbits the star-planet system's centre of mass. Another promising method relies on the detection, in stellar photometric time series, of the periodic dips caused by planets as they cross the disk of their parent star: planetary transits.

Many ground-based transit search projects have been operating for several years and are expected to come to fruition soon – a handful of planets detected via their transits have been confirmed already. From the ground, both transit and RV methods are limited to giant planets. Several space-based transit search missions are thus planned to probe the terrestrial and habitable planet regimes. The preparation of data analysis tools for these missions, in particular COROT and *Eddington*, has been the focus of my PhD, with potential application to ground-based data as a secondary objective.

I first developed and tested an algorithm for the automated detection of transits in white noise, a challenge due to the rare, brief and shallow nature of the transits. One of the most important noise sources for future space-based missions is the intrinsic low-amplitude variability of the parent star on timescales of tens of minutes to weeks. I constructed an empirical model of this 'stellar micro-variability' to simulate realistic light curves for a variety of stars, and developed filters to remove micro-variability. Monte Carlo simulations were used to test the performance of these tools alone and in combination, and to identify which types of stars make the most promising targets for *Eddington* & COROT.

The algorithms' performance was tested against that of others by participating in the COROT transit detection blind exercise, in which a number of groups from across Europe applied their algorithms to a set of simulated light curves of content known only to a game master. A transit search was also performed in 5 nights of data obtained in 2003 by the UNSW transit search team using the 0.5m APT telescope in Siding Springs Observatory in the field of open cluster NGC 6633, and a handful of transit candidates with depths below 50 mmag were identified.

Declaration

I hereby declare that this thesis entitled *Planetary transits and stellar variability* is not substantially the same as any that I have submitted for a degree, diploma or other qualification at this or any other university. This thesis is less than 60 000 words in length. Parts of this thesis were published in the following articles and conference proceedings:

- *Practical planet prospecting*
Aigrain, S., Irwin, M., 2004
Monthly Notices of the Royal Astronomical Society, vol. 350, pp. 331–345.
- *Characterising stellar micro-variability for planetary transit searches*
Aigrain, S., Favata, F., Gilmore, G., 2004
Astronomy & Astrophysics, vol. 414, pp. 1139–1152.
- *Impact of stellar micro-variability on Eddington’s planet-finding capability*
Aigrain, S., Favata, F., Gilmore, G., 2004
in ‘Stellar structure and habitable planet finding’, 2nd Eddington Workshop, ed. F. Favata & S. Aigrain, ESA SP-538, pp. 215–224.
- *Detecting planetary transits in the presence of stellar variability. Optimal filtering and the use of colour information*
Carpano, S., Aigrain, S., Favata, F., 2003
Astronomy & Astrophysics, vol. 401, pp. 743–753
- *The Frequency Content of the VIRGO/SoHO Light Curves: Implications for Planetary Transit Detection from Space*
Aigrain, S., Gilmore, G., Favata, F., Carpano, S., 2003
in ‘Scientific Frontiers in Research on Extrasolar Planets’, ed. D. Deming & S. Seager, ASP Conference Series, vol. 294, pp. 441–444.
- *Bayesian detection of planetary transits. A modified version of the Gregory-Loredo method for Bayesian periodic signal detection*
Aigrain, S., Favata, F., Gilmore, G., 2002
Astronomy & Astrophysics, vol. 395, pp. 625–636.
- *A Bayesian algorithm for the detection of planetary transits*
Aigrain, S., Gilmore, G., Favata, F., 2001
in ‘Techniques for the detection of planets and life beyond the solar system’, 4th Annual ROE Workshop, ed. W. R. F. Dent, p. 8.

This thesis is essentially my own work. However, some of the work presented therein

was done in collaboration with others, as detailed below:

- Chapter 2
The Bayesian algorithm presented in Section 2.1 is essentially my own work but benefited from regular discussions with and advice from F. Favata throughout its development and testing. The box-shaped transit finder was developed in collaboration with M. Irwin. Its implementation and initial testing were carried out in parallel by both of us, though the final Monte Carlo tests are my own work.
- Chapter 3
The work presented in this Chapter is essentially my own, but benefited from regular discussions with F. Favata and G. Gilmore, and comments from A. Lanza.
- Chapter 4
The starting point for the work presented in this chapter was the optimal filter described in Carpano et al. (2003), which was developed by S. Carpano with contributions from F. Favata and myself. The filters presented in this Chapter were developed in collaboration with M. Irwin, their implementation and testing being carried out in parallel by both of us.
- Chapter 5
The work presented in this chapter is essentially my own but benefited from discussions with F. Favata, G. Gilmore and M. Irwin, as well as with M. Auvergne for the COROT simulations.
- Chapter 6
This Chapter describes the COROT blind exercise, which was carried out by a collaboration led by C. Moutou with the following participants besides myself: M. Auvergne, P. Barge, D. Blouin, R. Cautain, A. R. Erikson, V. Guis, P. Guterman, M. Irwin, A. F. Lanza, F. Pont, H. Rauer, H. Voss and S. Zucker. Sections which describe work that was not my own were included when deemed necessary to place my work in context. The contents of Section 6.1 are the result of discussions between all the participants. The contents of Section 6.2 are primarily the work of others except for Subsection 6.2.3 which describes two sets of stellar light curves, one simulated by myself and one by A. Lanza. Section 6.3 contains work carried out by myself in collaboration with M. Irwin, except for Subsection 6.3.2, which summarises the methods used by the other participating teams for comparison. The contents of Sections 6.4 and 6.4.3.1 are essentially my own although they benefited from discussions with M. Irwin.

Acknowledgements

During the three years this thesis was in the making, three people in particular have guided and encouraged me. Fabio, who first got me into this transit business – albeit by bribing me with some fantastic sailing –, who has been a continuous source of advice and discussions since then, who gave me my first editing experience, and who kept a watchful eye over my work from afar with his customary attention to detail. Gerry, who agreed to supervise my project and gave me complete freedom to pursue it, with the occasional confidence building seal of approval, who helped me see my work in the wider context, and who was the provider of the life-saving loA espresso machine. Mike, who applied his time, his mind, his tremendous experience and his patience to help me grab problems by the roots and shake them till only what mattered was left, and for those evenings spent chatting in the pub, wine-tasting or watching Shakespeare. To all three, thanks also for proof-reading this thesis and providing many suggestions for its improvement.

I would also like to thank Keith Horne and Wyn Evans for agreeing to be my examiners.

I am indebted to Thierry Appourchaux and the VIRGO team at ESTEC for access to VIRGO data, to Derek Buzasi for access to WIRE data, to Rachel Street and the St Andrews exo-planet team for access to data from their INT transit search, to the UNSW transit search team for access to APT data, and to Floor van Leeuwen for extracting Hipparcos light curves from the database for me. This research has made use of the NASA ADS service and the SIMBAD database operated by CDS (Strasbourg), and was supported by studentships from PPARC & the Isaac Newton Trust.

Thanks to John Webb and Marton Hidas for inviting me to Australia where I had a fantastic time, to Leslie for being such a great observing companion. I had many useful discussions with the members of the loA Stellar Populations and Star Formation Groups, with Andrew Collier-Cameron, Keith Horne, Jon Jenkins, Hans Kjeldsen, Nuccio Lanza, Mark McCaughrean, Sergio Messina, Giusi Micela and Sami Solanki. A special thanks to Magali Deleuil, Claire Moutou, Pierre Barge, and to Annie Baglin and the COROT community as a whole.

Thanks to the loA helpdesk, administrators, secretaries and support staff for their invaluable help and their friendliness, and to Max Pettini and Paul Hewett for being so helpful us graduate students. I am fondly indebted to Alan, Bunkie, Carine, Daniel, Edu, Elena, Estelle, Giuseppe, Justin, Lisa, Lola, Max, Michelle, Mike, Muon, Neil, Weimin and all the loA students and postdocs for making the loA such a great

place to work and be. A special mention for Simon who, besides sharing everything from cigarettes, coffee, books, music and films to dancing, holidays, the occasional work-related chat (honest), the panic ridden TAC proposal deadlines and some fun observing, graciously – though perhaps not always entirely willingly – provided me and those who visited my office with subject material for my white board cartoons. By the way, I challenge you to have your kitchen floor done by the time of my viva.

Thanks to Rosemary, Cathie, Robin, Rory, Sangam and everyone at 9 Adams Road, for a very special place to live, to Ed, Ze, Ben, Merav, Pete, Theo, Carl for dinners, dancing, parties and general misbehaviour (long live the bouncing pixies), and to Celine, Rory and Amy, Jago and Vanita, Tom Simon, Javi, Corin, Alistram, Jan, Sophie, for being friends and good times further afield.

Mireille et Philippe, Loulou, Michel et Geo, Mounette et Papy, pour la vie, l'appétit (de vivre, d'apprendre, des bonnes choses, culinaires et autres), la montagne, les livres, pour avoir répondu a mes questions, pour l'ouverture (d'esprit, de toutes les possibilités). Jude, Terry and Mary for being a second family.

And Tom, for being there, and being Tom.

Contents

1	Introduction	1
1.1	Searching for extra-solar planets	1
1.1.1	Context and motivation	3
1.1.2	Exo-planet detection methods	4
1.1.3	Properties of known extra-solar planets	12
1.2	Current and planned transit searches	19
1.2.1	Ground-based	19
1.2.2	Space-based transit searches	26
1.2.3	Terrestrial and habitable regimes:	27
1.3	The broader perspective	36
1.3.1	The time domain in astrophysics	36
1.3.2	Additional science from transit search data	37
1.3.3	Application of transit search tools to other data	38
1.4	Structure of the thesis	38
2	Transit detection algorithms	39
2.1	A Bayesian, step-function based algorithm	41
2.1.1	Derivation of the algorithm	42
2.1.2	Comparison with the original GL method	50
2.1.3	Performance evaluation in white Gaussian noise	53
2.1.4	Tests with solar micro-variability	63
2.2	A stripped-down box-shaped transit finder	68
2.2.1	Likelihood maximisation in Gaussian noise	69
2.2.2	Simplification of the algorithm	70
2.2.3	Improvements in the implementation	77
2.2.4	Performance evaluation	79
2.2.5	Discussion and future work	81
3	Characterising and simulating stellar micro-variability	83
3.1	Introduction	83

3.2	Clues from solar irradiance variations	85
3.2.1	SoHO/VIRGO total irradiance (PMO6) data	86
3.2.2	Modelling the 'solar background'	87
3.2.3	Evolution of the power spectrum with the activity cycle	88
3.3	Empirical scaling to other stars	90
3.3.1	The amplitude of the active regions component	91
3.3.2	Other parameters of the model	98
3.4	Testing the model	101
3.4.1	Mimicking the Sun	101
3.4.2	Trends with age and mass	102
3.4.3	Behaviour at high activity	103
3.5	Discussion	105
4	Variability filters	107
4.1	Introduction	107
4.2	Wiener or matched filtering approach	109
4.3	Least-squares fitting	113
4.4	Non-linear filtering	115
4.5	Light curve characteristics after filtering	118
4.6	Discussion	119
5	Impact of stellar micro-variability on transit detection from space	123
5.1	Introduction	123
5.2	Simulations to identify the best target stars for <i>Eddington</i>	124
5.2.1	Which are the best target stars for <i>Eddington</i> ?	125
5.2.2	Minimum detectable transit in the presence of variability	128
5.2.3	Relevance to <i>Kepler</i>	129
5.3	Performance of the non-linear filter plus box shaped transit finder	130
5.3.1	Preliminary tests	130
5.3.2	Monte Carlo simulations for <i>Eddington</i>	130
5.3.3	COROT simulations	134
5.4	Discussion	135
6	The COROT blind experiment	137
6.1	Introduction	137
6.1.1	Purpose of the blind experiment	138
6.1.2	Participants	138
6.2	The simulated light curves	139
6.2.1	General characteristics	139

6.2.2	Instrumental and photon noise	139
6.2.3	Stellar micro-variability	140
6.2.4	Transit signals	142
6.2.5	Stellar mimics and variables	143
6.2.6	Background pollution	143
6.2.7	The final set of 1001 light curves	143
6.3	The transit search	144
6.3.1	Method used by the loA team (team 5)	145
6.3.2	Methods used by the other teams	155
6.4	Discussion of the results	156
6.4.1	Relative performances of the various groups	157
6.4.2	Overall results and implications	159
6.4.3	Lessons learnt	161
6.5	Future prospects for COROT blind exercises	167
7	Application to ground-based data	177
7.1	The data	177
7.1.1	Telescope and instrument	177
7.1.2	Target	179
7.1.3	Data reduction and photometry	179
7.1.4	Post-processing	180
7.2	Transit search	183
7.3	Discussion	185
8	Conclusions	201
8.1	Data analysis tools for planetary transit searches	201
8.2	Applications	202
8.3	Future improvements	203
8.4	Additional tools	203

Chapter 1

Introduction

1.1 Searching for extra-solar planets

The search for, and study of, planets outside the solar system have developed into one of the most active topics in modern astronomy over the past decade.

Before any discussion can begin, it is necessary to adopt a definition of the word planet. Where does the boundary between brown dwarfs and planets lie? Is it purely a question of internal structure (i.e. mass), or does the formation process have to be taken into account as well? For the present work, I have used the working definition adopted by the IAU Working Group on Extrasolar Planets¹:

- 1. Objects with true masses below the limiting mass for thermonuclear fusion of deuterium (currently calculated to be 13 Jupiter masses for objects of solar metallicity) that orbit stars or stellar remnants are “planets” (no matter how they formed). The minimum mass/size required for an extrasolar object to be considered a planet should be the same as that used in our Solar System.*
- 2. Substellar objects with true masses above the limiting mass for thermonuclear fusion of deuterium are “brown dwarfs”, no matter how they formed nor where they are located.*
- 3. Free-floating objects in young star clusters with masses below the limiting mass for thermonuclear fusion of deuterium are not “planets”, but are “sub-brown dwarfs” (or whatever name is most appropriate).*

A variety of names and spellings for planets outside the solar system can be found in the literature: *extra solar planets*, *extra-solar planets*, *extrasolar planets*, *exoplanets*, *exo-planets*. . . For consistency, *exo-planets* has been adopted throughout the present thesis, though the choice is arbitrary.

¹<http://www.ciw.edu/boss/IAU/div3/wgesp/definition.html>

The first detection of an exo-planet orbiting a main-sequence star was made using the radial velocity (hereafter RV) method by Mayor & Queloz (1995), and was soon confirmed by Marcy & Butler (1995). Further detections rapidly ensued (Marcy & Butler 1996; Butler & Marcy 1996; Butler et al. 1997; Cochran et al. 1997, to cite only the first few), with the first exo-system (one star orbited by more than one known planet) discovered by Butler et al. (1999). These discoveries have provided the first hints of answers to a number of fundamental questions previously relegated to the realm of pure speculation, though only the very tip of the iceberg has been uncovered. We can now start to ask whether the solar system is a typical one, and are in a much better position to understand how it came to be what it is now: until 1995, planetary system formation scenarios had only one example to model themselves on.

More than a hundred stars are now known to harbour planets. A number of regularly updated websites act as useful information repositories on these systems: J. Schneider's Extrasolar Planets Encyclopedia², the California and Carnegie Planet Search Almanac³, as well as the list maintained by the IAU Working Group on Extrasolar planets⁴. Little by little, the range of parameter space which has been explored widens, the initial selection effects (such as the predilection for short-period planets induced by the limited duration of the RV search programmes) lessen, and the numbers become sufficient for significant conclusions to be drawn from the distributions of the observable characteristics of the planets (see for example Zucker & Mazeh 2002; Udry et al. 2003b; Santos et al. 2003; Eggenberger et al. 2004).

However, as the number of known systems and our knowledge of their properties has increased, more questions have been raised than answers found. The best known example of this is the question of the 'Hot Jupiters': the discovery of relatively large numbers of gas giant planets, orbiting close to their parent star. This poses a striking contrast to the solar system, where the gas giant planets are found beyond 5 AU, and has sparked intense and ongoing debate: can such large planets form that close to their star? If a migration mechanism is invoked, so that they could have formed further out and moved to their current positions later, what would stop the migration at the observed radius?

This is only one of the puzzles raised to date. Many more detections are expected in the next few years, with the advent of a number of ground-based searches using the transit method, which can be used to survey large areas of the sky simultaneously. Other methods, such as microlensing, probe totally different areas of parameter space, being capable of detecting lighter, smaller planets orbiting more distant stars with larger orbital distances. In all cases, two observational parameters are of

²<http://www.obspm.fr/encycl/encycl.html>

³<http://exoplanets.org>

⁴<http://www.ciw.edu/boss/IAU/div3/wgesp/planets.shtml>

key importance: precision (whether spectroscopic, photometric or astrometric) and time coverage (both high time sampling and long baselines are necessary).

This thesis summarises the author's attempts to contribute to the drive towards the next generation of planet discoveries via the transit method, with a particular emphasis on habitable planets. As the radii of interest are small (the planets have to be terrestrial rather than gaseous) and the periods of interest are a few months or longer (to ensure that liquid water can exist at the surface), the detection of habitable planets implies a significant step both in precision and in time coverage, and therefore dedicated instrumentation. It also requires a significant effort in terms of data analysis, which is the focus of the present work.

1.1.1 Context and motivation

The first RV searches for exo-planets were pursued for a decade before the first detection. There were compelling reasons to expect exo-planets to exist (and to be detectable), and therefore to pursue the projects further. Indeed, planetary formation scenarios in the early 1990s (as outlined for example by Udry 2000), though first developed centuries before the first exo-planet was known, predicted that they should be common.

The 'solar nebula' model, developed by Kant and Laplace as early as the 18th century, invokes mechanisms which are still considered important in star- and planet-formation scenarios today. Local instabilities in a gaseous cloud drive its collapse to a (number of) centrally concentrated 'nebula(e)'. As gas is accreted onto the proto-star at its centre, conservation of angular momentum induces the nebula to collapse to a differentially rotating disk, in which planets can form. This implies the coevality of a given star and its surrounding planets, an idea that is supported by the coincident ages measured for various components of the solar system (the Sun, the Earth, the Moon, meteorites) via independent methods. If planets are a natural by-product of star-formation, they must be nearly as pervasive as stars themselves. The discovery of exo-planets would confirm this prediction. Increasing the number of 'laboratories' in which to observe the products of planetogenesis would also undoubtedly lead to surprises, and thus changes to, and refinement of, the model.

More detailed observations could then be used to study the atmospheres and internal constitution (density, composition) of exo-planets, and compare them to what is known about solar-system planets. Finally, and perhaps most importantly for the wider public, it would provide insights into the ubiquity or otherwise of Earth-like environments, or of the different types of environments that may support life, and therefore of life as we know it.

1.1.2 Exo-planet detection methods

Numerous reviews have covered the variety of exo-planet detection methods over the last few years (see for example Marcy & Butler 1998; Perryman 2000). Although the following repeats what has been said many times before, a brief summary of the principles, characteristics and current prospects of these methods is nonetheless needed to place transit searches in context.

1.1.2.1 Direct observation

The problem with trying to directly image even nearby exo-planets arises from their small projected separation (50 mas for a Sun-Jupiter system at 100 pc) and high brightness contrast ($\sim 10^9$ in the optical, $\sim 10^6$ in the IR) to their parent star. This method is thus, for most systems, impractical with today's technology, though efforts are underway to develop instruments for this task using high Strehl ratio adaptive optics combined with coronagraphy, or interferometry. The problem of contrast is reduced in the infrared domain, especially for young planets which may emit significant amounts of intrinsic, as well as reflected thermal radiation. Recently, Chauvin et al. (2004) reported the detection with VLT/NACO of a candidate $5 M_{\text{Jup}}$ giant planet near a brown dwarf in the TW Hydrae association. A spectrum of the candidate exo-planet shows signs of water molecules, indicating it is a cool object, but further observations are required. Planned instruments both from the ground (e.g. VLT/Planet Finder) and from space (Darwin/TPF) make use of the advantages offered by the infrared domain.

Both contrast and projected separation issues are less critical for planets orbiting white dwarfs. Recent studies showed a planet could survive the late stages of stellar evolution, its orbital distance increasing proportionally to the amount of mass loss by the star, resulting in a potentially detectable system (Burleigh et al. 2002), though no detection has been announced so far.

Given the difficulty of obtaining direct observations, most exo-planet searches to date have concentrated on indirect methods, several of which rely on detecting the motion of star around centre of mass of the star-planet system.

1.1.2.2 Astrometry

The transverse component of the reflex motion of the star causes it to follow an elliptic trajectory on the sky, which can be measured astrometrically. This method is particularly sensitive to orbital planes perpendicular to the line of sight, and to planets in relatively long period orbits. This makes it complementary to, but difficult to

confirm by, other methods. Jupiter orbiting the Sun, when viewed from 10 pc away perpendicular to the line of sight, would result in an amplitude of 0.5 mas (milli-arc second) per year.

Measuring such a small displacement over a time baseline of years from the ground is very challenging and, though candidates have emerged over the last few decades, none have been confirmed. Ground based projects expected to deliver sub-mas precision in the next few years include VLTI and ALMA, but the potential of this method is much increased from space. Data from Hipparcos, which surveyed over 100 000 stars to mas precision, has already been used to constrain masses derived from other methods (Frink 2003). In the future, the pointed interferometric mission SIM will be ideally suited to providing tight constraints on the orbits of previously detected exo-planets, while the survey mission GAIA, designed to achieve μ as precision, is expected (under reasonable assumptions) to detect tens of thousands of Jupiter-mass planets astrometrically (see Perryman 2000, and references therein).

1.1.2.3 Pulsar timing

The motion of the central star about the centre of mass of the system causes variations in light-travel time across the orbit. However, only objects showing intrinsic, periodic brightness modulation, such as pulsars, provide a sufficiently precise frequency reference to measure this effect. In particular, precise timing of millisecond pulsars allows the detection of terrestrial mass planets, two of which were found by Wolszczan & Frail (1992) around PSR1257+12, the first detection of a planet outside the solar system. Additional planets were detected in that and other systems in the intervening years. These systems provide interesting constraints on planetary system formation and evolution, but the method only applies to this very distinct class of objects.

1.1.2.4 The radial velocity method

The most successful exo-planet detection method to date has been the radial velocity (RV) method, which consists in measuring periodic shifts in the wavelengths of spectroscopic lines due to the star's reflex motion. This method can be implemented from relatively small telescopes (~ 1 m and above) for the brightest target stars, though 8–10 m-class telescopes are needed for fainter targets and/or to resolve lower RV amplitudes. It requires high resolution spectra to be compared over a long period of time, implying a very stable spectrograph and accurately controlled temperature environment, as well as a very precise wavelength calibration, which sometimes achieved by inserting in the optical path a cell containing a gas produc-

ing absorption lines of well known wavelength (e.g. iodine cell). High throughput and sensitive detectors are also vital to gather a sufficient number of photons. These factors concurred to delay the use of the RV method for exo-planet searches until the mid- to late-eighties.

The RV method is most sensitive to high mass planets in close orbits, although the bias towards short period reduces as the baseline of observations of the existing programs increases. While RV measurements provide the period and eccentricity of the orbit, a significant drawback is that only a lower limit on the mass, $m \sin(i)$ where i is the inclination of the orbit relative to the line of sight, can be measured. This degeneracy can be lifted only through complementary observations, for example by the astrometric or transit methods.

One limitation of this method arises from the intrinsic velocity fluctuations at the surface of the star. These lead to an observed jitter in the RV measurements (Saar et al. 1998; Paulson et al. 2002), which could hinder the applicability of the method to active (very late type or young) stars. It also limits the achievable precision for Sun-like stars, making the detection of terrestrial mass planets with this method impossible. The lowest mass planet detected to date was a $14 M_{\oplus}$ planet, found with the recently commissioned HARPS instrument, which achieves sub- 1 m s^{-1} precision (Santos et al. 2004). The necessity to monitor a single star at a time, and the requirement that the star be bright enough to perform very high resolution spectroscopy, also limit the volume of space probed by the RV method to the immediate solar neighbourhood.

1.1.2.5 Microlensing

Microlensing programs aim to measure the temporary magnification of a background star (the source), as a foreground star (the lens) passes in front of it and its gravitational potential bends the light emanating from the source. If the lens is orbited by a planet, a distortion to the well known magnification profile will be observed. The amplitude of this distortion is maximum if the planet's orbital distance is close to the lens's Einstein radius, which for a late-type star is a few AU. This signal is brief and rare, but it is not particularly small, and even a moon can produce a detectable signal. The microlensing method thus requires high time sampling photometry of a large number of stars. A number of ground based microlensing programs have discovered candidate planets, but the sampling of the light curves is not sufficient to exclude other types of events. Going to space, as proposed with the GEST spacecraft, would improve the chances of a detection because it avoids interruptions in the observations.

This method is sensitive to relatively small, far-out planets orbiting distant stars, and as such is radically different from the others. While the planet mass and orbital

distance can be deduced from the light curves, measurements are not repeatable, and cannot be confirmed by other methods. The main potential for this method is thus to use the results to constrain period and mass distributions in a statistical manner, if sufficient numbers of candidate detections become available.

1.1.2.6 Reflected light

This method, principally applicable to giant planets, consists in measuring the modulation of the light emitted by the star and reflected by the planet as it follows its orbit. The amplitude of this modulation depends on the albedo of the planet, and may be detectable from the ground for 'hot', close-in planets in the IR, where it is strongest, for stars already known to harbour planets. One important advantage of this method is that, although the light variations are strongest for orbits close to the line of sight, the planet does not have to be transiting to be detectable. Given the low amplitude of the signal, its real potential lies in space-based observations, for example the MOST satellite (a small asteroseismic mission which has a number of planet host stars on its target list) or space-based missions primarily designed for planetary transit searches.

1.1.2.7 The transit method

If a planet's orbit is aligned with the line of sight, it will partially obstruct the disk of its parent star once per orbit. This will cause a periodic, temporary dip in the observed stellar brightness, known as a planetary transit.

A significant advantage of this method is that many thousands of stars can be monitored simultaneously, and (depending on the size of telescope used) relatively distant stars can be probed. Like the RV method, the transit method is most sensitive to large, close-in planets, though long time bases and high precision photometry reduce these biases. The geometry of a planetary transit is illustrated in Figure 1.1, and its implications in terms of transit probability, depth and duration are explored below.

Transit probability: First of all, for a transit to occur, its orbit must be aligned with the line of sight. To compute the probability of this occurring, we assume that planetary orbits are isotropically aligned, so that the inclination i (which is measured, in Figure 1.1, from the +z-axis), is drawn at random from a distribution uniform in $\cos i$ between 0 and 90° . A transit occurs if

$$b = a |\cos i| \leq R_{\text{crit}}, \quad (1.1)$$

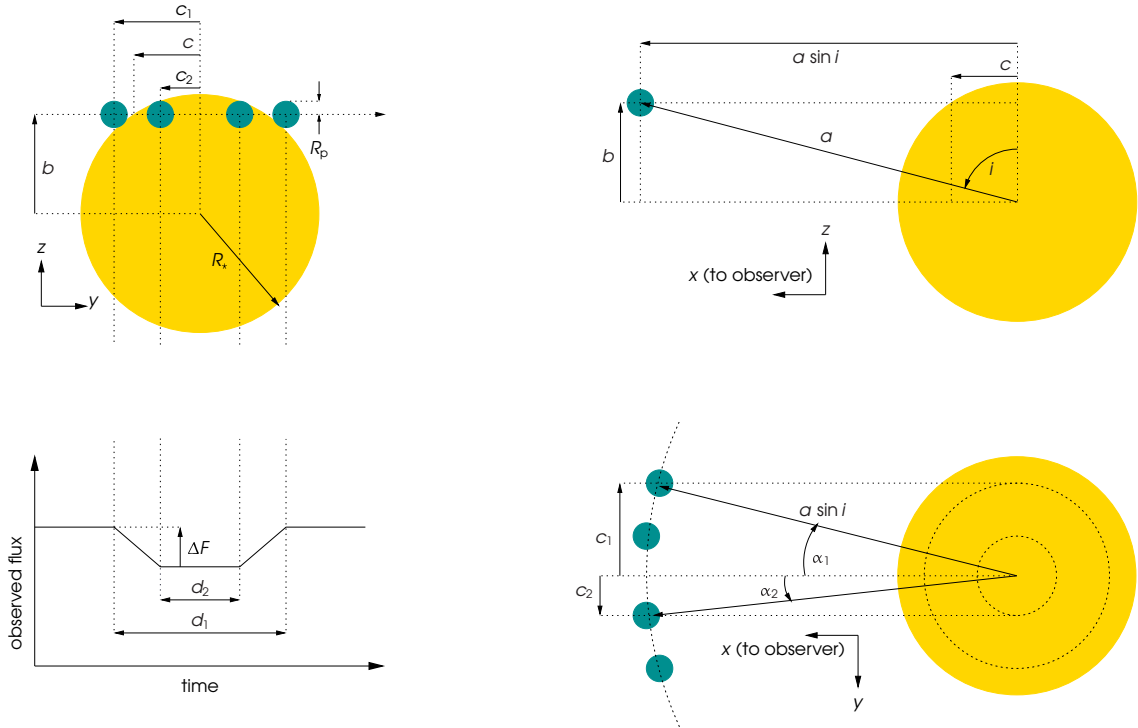


Figure 1.1: Geometry for a transiting planet with a circular orbit. Clockwise from top left: face-on view (as seen by the observer), side-on view, downward view and corresponding light curve. The star and planet radii are R_* and R_p respectively, the orbital distance a and the inclination i . The planet is shown at the four contact points (start and end of ingress and egress) in the top left and bottom right panels, and in mid-transit in the top right panel. The total transit duration is d_1 , the duration of totality d_2 and the transit depth ΔF . Note that the base of the 'observed flux' axis does not correspond to zero flux. The relative size of star and planet shown here would correspond, for a Sun-like star, to a $1.08 R_{\text{Jup}}$ planet. The effects of limb darkening are ignored here.

where b is the impact parameter, i.e. the closest distance between the star and planet centres, projected onto the y - z plane, a is the planet's orbital distance, and R_{crit} is a critical radius. Depending on whether one is considering only full transits, where the planet's disk fully overlaps with that of the star for some time, or whether grazing transits are also considered,

$$R_{\text{crit}} = R_* - R_p \quad \text{or} \quad R_{\text{crit}} = R_* + R_p. \quad (1.2)$$

Taking $R_{\text{crit}} = R_*$ provides a rough estimate which is suitable for most purposes, as usually $R_p \ll R_*$.

Assuming an isotropic distribution of inclinations, the probability for a transit to occur is then given by

$$P(\text{transit}) = \frac{R_*}{a}. \quad (1.3)$$

We can use Kepler's third law

$$\left(\frac{p}{2\pi}\right)^2 = \frac{a^3}{G(M_* + M_p)}, \quad (1.4)$$

where p is the orbital period, G is Newton's constant of gravitation and M_* and M_p are the star and planet mass respectively, to express Equation (1.3) in terms of period:

$$P(\text{transit}) = \left[\frac{R_*^3 4\pi^2}{p^2 G (M_* + M_p)} \right]^{1/3}. \quad (1.5)$$

The transit probability is thus a steeply decreasing function of orbital period. It can be as high as $\sim 30\%$ for a 3 Myr old system containing a Jupiter-mass planet in a 3 d orbit around a solar-mass star (the pre-main sequence stellar radius, which is significantly larger than that of a main-sequence star of the same mass, was taken from Rhode et al. 2001). On the other hand, an Earth-like planet orbiting a Sun-like star at a distance of 1 AU has a transit probability of 0.3%.

Transit depth: If the planet's disk is fully aligned with that of the star at any point during the transit (i.e. excluding grazing events), the transit depth is given by the ratio of the areas of the planetary and stellar disks (ignoring limb darkening), i.e.

$$\frac{\Delta F}{F} = \left(\frac{R_p}{R_*}\right)^2, \quad (1.6)$$

where F is the mean out-of-transit flux observed. The transit method thus provides a direct measure of the ratio of planet to star radius. This depth is generally small, of the order of 1 and 0.01% for Jupiter- and Earth- sized planets respectively (if orbiting a star of solar radius). Giant planets are thus detectable from the ground with this method, while terrestrial planets can only be detected from space, as atmospheric scintillation limits the achievable photometric precision (see Young 1967 and Gilliland & Brown 1992 for a formula to estimate scintillation noise for a given airmass, telescope diameter, observatory altitude and exposure time).

Transit duration: The total duration of the transit is:

$$d_1 = p \alpha_1 / \pi, \quad (1.7)$$

where α_1 , as illustrated in the bottom right panel of Figure 1.1, is the angle between the line connecting the centres of the star and of the planet at the first or fourth contact point (beginning of ingress or end of egress) and the x-axis (observer's di-

rection), projected onto the x - y plane. In turn,

$$\alpha_1 = \arcsin (c_1/a \sin i), \quad (1.8)$$

where

$$c_1 = \sqrt{(R_* + R_p)^2 - b^2}. \quad (1.9)$$

Therefore

$$d_1 = \frac{\rho}{\pi} \arcsin \left[\frac{1}{a \sin i} \sqrt{(R_* + R_p)^2 - (a \cos i)^2} \right] \quad (1.10)$$

$$= \frac{\rho}{\pi} \arcsin \left[\frac{R_*}{a \sin i} \sqrt{\left(1 + \frac{R_p}{R_*}\right)^2 - \left(\frac{a}{R_*} \cos i\right)^2} \right]. \quad (1.11)$$

Similarly, the duration of totality is given by

$$d_2 = \frac{\rho}{\pi} \arcsin \left[\frac{R_*}{a \sin i} \sqrt{\left(1 - \frac{R_p}{R_*}\right)^2 - \left(\frac{a}{R_*} \cos i\right)^2} \right]. \quad (1.12)$$

If we assume, as before, that $R_* \ll a$, α_1 is small and $\arcsin \alpha_1 \simeq \alpha_1$, Equation (1.11) becomes

$$d_1 = \frac{\rho R_*}{\pi a \sin i} \sqrt{\left(1 + \frac{R_p}{R_*}\right)^2 - \left(\frac{a}{R_*} \cos i\right)^2}. \quad (1.13)$$

For a central transit ($i = 90^\circ$), this further simplifies to

$$d_1 = \frac{\rho (R_* + R_p)}{\pi a} = (R_* + R_p) \left[\frac{4\rho}{\pi G (M_* + M_p)} \right]^{1/3}. \quad (1.14)$$

The transit duration is thus a slowly increasing function of orbital period. The total transit duration for a Jupiter-like planet orbiting a Sun-like star in a 3 d orbit with an inclination of 90° is ~ 2.6 hr. If the planet is Earth-like and has an orbital period of 1 yr, it is ~ 13 hr.

Practical implementation: From the ground, the transit method can be implemented using small telescopes which are no longer in heavy demand for other purposes, or very small telescopes, purpose-built using off-the-shelf, commercially available components. However, it is only recently that advances in detector technology and data reduction techniques (to obtain stable, high precision photometry over wide fields of view), and in data storage and processing capabilities (dealing with to the huge quantities of data to be stored and analysed) have made the implementation of

this method possible.

Available parameters: As shown by Equation (1.6), transit observations provide a direct measure of the ratio of planet to star radius. In a recent paper, Seager & Mallén-Ornelas (2003) show that three additional parameters can be derived directly from observables using Equations (1.4), (1.6), (1.11) and (1.12): the ratios of impact parameter b and orbital distance a to the star radius, and the stellar density ρ_* . They go on to show that, provided the transit period is known from repeat observations, the stellar mass is known (e.g. from spectroscopy), and given an assumed mass-radius relation for the star, one can solve for five physical parameters of the system: the stellar mass M_* and radius R_* , the planet radius R_p , the orbital distance a and the inclination i .

Without going into that level of detail, one can immediately see that the observation of a transit implies an inclination close to $i = 90^\circ$, particularly if the period is long. If combined with RV observations, a planet mass measurement free of the inclination degeneracy can thus be obtained. Small departures from strict periodicity can indicate the presence of additional (non-transiting) planets in the system, while detailed observations can reveal the presence of moons and rings.

Limitations: The transit method does suffer from a number of limitations. First, the combined effects of transit probability and duration imply that, with a given observation strategy (duration and time sampling of the observations), the probability of obtaining any in-transit observations of a given planet decreases rapidly with orbital period. This means that many stars must be monitored for long periods for a detection to be likely.

There are also numerous phenomena that can mimic transits, for example high mass ratio eclipsing binaries, or binaries whose eclipses are diluted by a third star, whether spatially close by (triple system) or fortuitously aligned ('blend'). Some of these effects can be ruled out by close examination of the light curve, others require radial velocity measurements to exclude stellar or sub-stellar mass companions. Intrinsic brightness variations of the the parent star can also hinder transit detection, or be a source of confusion, particularly for small transit signals.

Status: A large number of ground based projects are underway. The transit of a known RV-planet was first discovered by Charbonneau et al. (2000) and five *detections* via the transit method have been confirmed as of August 29th 2004: four were candidates announced by the Optical Gravitational Lensing Experiment (OGLE) collaboration (Udalski et al. 2002a,b, 2003; Konacki et al. 2003a, 2004; Bouchy et al.

2004; Pont et al. 2004), and one was found by the Trans-atlantic Exoplanet Search (TrES) network (Alonso et al. 2004a), the latter being the first exo-planet discovered by the transit method from very small (10 cm) aperture telescopes. The specificities of the various projects, as well as the lessons learnt from their operation so far, will be discussed in Section 1.2.1. Several dedicated satellites are also at the planning or construction stage (see Section 1.2.2).

1.1.3 Properties of known extra-solar planets

The properties of the exo-planets detected so far are starting to provide interesting constraints for planet formation and evolution scenarios. A brief summary is given here in order to put in context the expected impact of future detections via the transit method.

The first noticeable characteristic of the known exo-planets as a whole is their variety. Few of the regions of parameter space within the current limits of the various detection methods have been found empty (and consequently, those that have and are not due to selection biases are of significance). This state of affairs represents a marked change on the situation ten years ago, when the known planets neatly divided into relatively close-in, terrestrial planets, outer gaseous giants, and smaller icy planets in the outer reaches of the solar system, all of which had eccentricities lower than 0.25 (0.2 if Pluto is classified as a Kuiper Belt Object).

1.1.3.1 Orbital parameters and planet mass

There is much to be learnt from scatter plots and histograms of the orbital parameters (orbital period or semi-major axis and eccentricity) and masses (or minimum masses, as given by radial velocity measurements) of known exo-planets. A set of such plots is shown in Figure 1.2. These plots were produced from the catalogue of known planets orbiting main sequence stars kept by J. Schneider (Extrasolar Planets Encyclopedia), which was cross-checked against the list maintained by the IAU Working Group on Extra Solar Planets. All planets announced before August 29th 2004 are included, but detections which remained uncertain or controversial at that date (as indicated on either list) were left out. Similar plots were discussed in much greater detail in several recent articles and reviews (see for example Udry et al. 2003b and Eggenberger et al. 2004), and only the main points of interest are highlighted here.

Hot Jupiters: Many exo-planets have been found in very short period orbits (less than 100 d, or even than 10 d). These very close-in planets, now known as 'Hot Jupiters', were not at all expected. The formation of Jupiter-mass planets was gener-

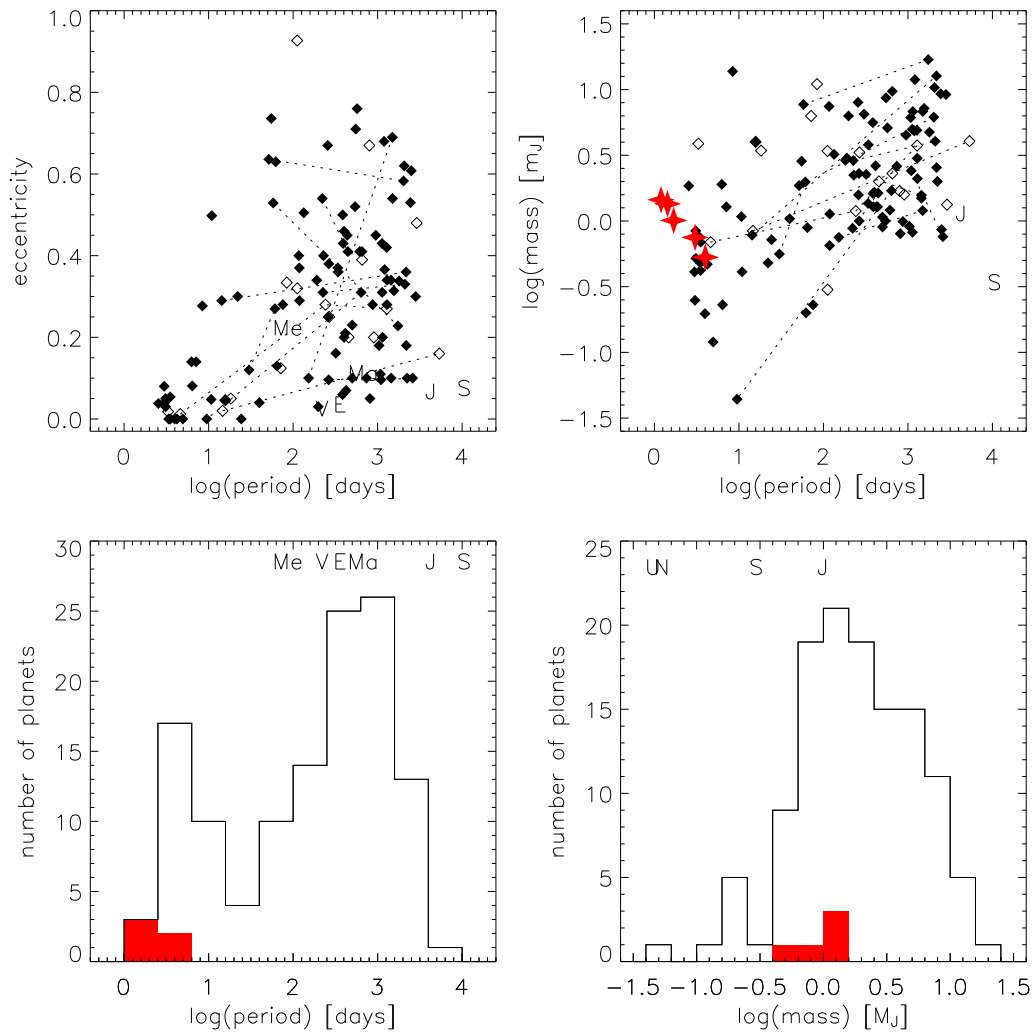


Figure 1.2: Scatter plots and distributions of eccentricity, orbital period and mass for all exoplanets known as of August 29th 2004 around main sequence stars (data from J. Schneider's *Extrasolar Planets Encyclopedia*). Filled symbols: planets orbiting single stars. Hollow symbols: planets orbiting multiple stars (from Eggenberger et al. 2004). Red symbols/area: planets detected via the transit method. Dotted lines link planets orbiting the same star. Letters mark the positions of solar system planets on each diagram.

ally thought to proceed via a two-step process (Pollack 1984; Lissauer 1993) of core formation (via collisions of smaller rocky or icy bodies) followed by runaway gas accretion once a critical mass has been reached. The core formation timescale is strongly dependent on the abundance of solid materials, which increases with distance from the central star, so that a core in the inner parts of a disk would not reach the critical mass before the disk evaporates (see Boss 1996, and references therein). It is therefore necessary to invoke a migration process to explain the existence of hot Jupiters: the planet forms at several AU from the star, but migrates inward due to

gravitational interactions with the protoplanetary disk it is embedded in (Goldreich & Tremaine 1979, 1980; Papaloizou & Lin 1984).

How migration stops before the planet is engulfed by the star is the subject of ongoing debate. There appears to be a critical period of 3 d, below which very few planets are found, and just above which a 'pile up' is observed. The recent discovery of a planet around HD 73256 with $P = 2.54$ d (Udry et al. 2003a) and of the three OGLE transiting planets (Udalski et al. 2002b, 2003) suggests that the cutoff may be less steep than previously thought, but it remains significant. This is not an observational bias: the planets on the shortest orbits should be the easiest to detect. It may be interpreted as a critical (stellar mass dependent) radius, where migration is halted. Several mechanisms for halting migration have been proposed (see for example Terquem 2004, and references therein). They include, amongst others, interaction with a magnetospheric cavity, reversal of migration in close-in eccentric orbits, interaction with a turbulent disk, tidal interaction with a rapidly rotating star, and mass loss through Roche lobe overflow – the latter an interesting possibility in view of the recent detection of an extended atmosphere of escaping hydrogen around the transiting planet HD209458b (Vidal-Madjar et al. 2003) – and the loss of the material in the disk – a scenario that would see many generations of protoplanets fall into the star and only the last 'lucky' ones survive. Again, how to remove the disk on a sufficiently short timescale for that to happen is an open question.

Eccentric orbits and multiple planet systems: The circular orbits of very close-in planets, as well as the range of eccentricities of planets in slightly longer orbits (periods between 6 and 21 days) can be explained by tidal interaction with the star (see for example Rasio & Ford 1996; Dobbs-Dixon et al. 2004). In the former, dissipation of the tidal disturbance induced by the planet within the star leads to circularisation and synchronisation of the orbit with the stellar rotation period, while the tides induced by the star within the planet lead to 'puffing up' of the planet which, in certain cases, can result in mass loss through Roche lobe overflow and the consequent expansion of the orbit (a possible explanation for the 3 d cutoff at the low end of the period range). In the latter, dissipation is shared between star and planet, leading to either damping or excitation of the eccentricity depending on the rotation period of the star.

The eccentricity distribution is uniform between 0 and 0.7 for periods longer than 21 d, a situation that is very different from the solar system. Planet-planet interactions within a disk can induce high eccentricity, specially if the planets' orbits are coupled through resonances (see for example Chiang 2003). This is supported by the fact that at least one planet of each multiple system (linked by dotted lines) has

significant eccentricity, and that several of the known systems have been found to contain planets in resonant orbits. Under certain circumstances, single planet-disk interactions can also excite eccentricity rather than damp it (Goldreich & Sari 2003).

Mass and distance dependence of migration: We have already established that migration is important to explain the properties of the known exo-planets, in particular those of hot Jupiters. The mass-period plot (top left panel of Figure 1.2) shows a lack of massive ($M_{\text{pl}} > 2 M_{\text{Jup}}$) planets on short orbits ($P < 100$ d), and of low mass planets ($M_{\text{pl}} < 0.75 M_{\text{Jup}}$) on long orbits ($P > 100$ d). Together with the valley in the period distribution for $10 < P < 100$ d, these suggest (in agreement with some theoretical predictions, see Udry et al. 2003b, and references therein) that migration rate increases with decreasing mass. Neither trend is an observational bias, but the first could also be explained if massive planets do migrate in but then disappear, either because their migration is not halted and they fall into the star, or because they lose a significant fraction of their mass. However, the second, with its sharp cutoff in mass, makes the mass-dependent migration explanation more compelling. Eggenberger et al. (2004) point out that outer low mass planets may yet be found at large radii or very low masses currently beyond the reach of RV searches.

Mass distribution: The mass distribution (bottom right panel of Figure 1.2) still suffers from incompleteness below $1 M_{\text{Jup}}$. Above that limit, a clear increase in frequency toward lower masses is observed, with a more or less linear trend in dN/dM . More precise measurements of this mass function, as the number of known planets increases and the completeness limit goes down, will provide important constraints for theoretical formation scenarios.

1.1.3.2 Properties of planet-host stars

Planets in binaries: Planets orbiting components of stellar binary systems are marked by hollow symbols in the top panels of Figure 1.2. As pointed out by Eggenberger et al. (2004), close inspection shows that their eccentricity is damped (< 0.1) out to longer periods ($P < 60$ d) compared to planets orbiting single stars (filled symbols), and that planets in binaries with periods > 100 d have moderate rather than high eccentricities. The most massive planets in close-in orbits also tend to be in binaries. This suggests that migration and/or mass accretion rates might be enhanced (and eccentricity damped) by the presence of a companion to the parent star, but the number of systems available for study is still relatively small (15).




Figure 1.3: Left: Metallicity distribution of stars with planets belonging to the CORALIE planet-search sample (shaded histogram) compared with the same distribution for a volume limited sample of ~ 1000 stars not known to host planets, observed with the same instrument. Right: Percentage of stars in the CORALIE planet search sample that have been found to harbour planetary-mass companions as a function of metallicity. Reproduced from Santos et al. (2003).

Host star metallicity: That planet-host stars appeared more metal rich than non planet-host stars was suspected early on (Gonzalez 1997, 1998), though small number statistics and the difficulty of obtaining an unbiased sample of metallicities for comparison have delayed the firm establishment of this result. However, recent studies (Santos et al. 2003; Fischer et al. 2004), where much attention has been paid to the removal of systematic effects and biases, seem to confirm that planet host stars are more metal rich than stars without planets, or conversely that more metal rich stars are more likely to harbour planets than more metal poor stars, as illustrated by Figure 1.3. One should of course keep in mind that this only applies to the type of planets currently detectable, i.e. gaseous giant planets in orbits less than ~ 5 AU.

This may be due to ‘pollution’ by infalling proto-planets at early stages, whose metal content is diluted into the stellar convection zone. This scenario is supported by the existence of binaries whose components appear to have different metallicities. On the other hand, it may be of primordial origin, planet formation mechanisms being more effective in metal-rich disks. No particular trend is observed in the metallicity of planet host stars with convective envelope depth, which supports the latter interpretation (Santos et al. 2003). In turn, the importance of metallicity in the formation process favours planetesimal formation through core accretion (see for example Pollack et al. 1996), rather than the alternative suggested by Boss (2000) of formation through gravitational instability.

No statistically significant trends between metallicity and orbital parameters, mass, or membership of a multiple planet or multiple star system have yet emerged, suggesting that migration is not strongly influenced by the metal content of the disk, though this may change as more planets are detected.

1.1.3.3 Where the transit method will make a difference

As shown through the very brief overview presented in Sections 1.1.3.1 and 1.1.3.2, many questions arise from the properties of the known exo-planets and their host stars. Plausible theoretical interpretations have been proposed for most of them, but there are very few cases where a single mechanism can fully explain a given aspect of the problem. Several processes can usually be invoked and often, the same process can have widely divergent consequences depending on small shifts in some parameters (for example, resonant planet-disk interactions can either damp or excite eccentricity). In the present section, we explore how detections via the transit method are particularly likely to contribute to the field.

Removal of the mass-inclination degeneracy: As pointed out earlier, the RV method only provides a minimum mass measurement for the planet, because of the degeneracy between mass and orbital inclination. All but 6 of the masses shown in Figure 1.2 are thus lower limits. Transit observations limit the inclination to a very narrow range close to 90° , thus breaking the degeneracy. Not only is this important to ensure the companion is really of planetary mass, it provides (given the radius measurements also available from transit observations) a means to determine the mass-radius relationship for exo-planets.

Spectroscopy of transiting planets: Transmission spectroscopy consists in taking spectra of a planet host star in- and out-of-transit and comparing them. During a transit, a small fraction of the light emitted by the star passes through the planet's atmosphere before reaching the observer. Absorption lines corresponding to elements that are particularly abundant in the planetary atmosphere will be stronger in the spectra taken during the transit than in those taken out-of-transit. This can be viewed as the planet having a wavelength-dependent radius, resulting in a wavelength-dependent transit depth. For close-in giant planets, this effect can be detectable.

If clouds are present in the planet's atmosphere, only the species prevalent at altitudes above the clouds, or those which absorb in spectral regions where the clouds are not opaque, will produce a detectable effect. The amount of additional absorption by a given element that is observed in-transit is thus determined not

only by the abundance of that element in the atmosphere but also by the altitude at which the absorption occurs (Brown 2001). Transmission spectroscopy therefore yields some information about the vertical structure, as well as the composition, of exo-planetary atmospheres.

A very weak detection of neutral Na in the atmosphere of HD209458b by Charbonneau et al. (2002) was thus interpreted as indicating that sodium was only present in significant amounts in the lower part of the atmosphere. Detection of more significant absorption by hydrogen (Vidal-Madjar et al. 2003) and oxygen and carbon (Vidal-Madjar et al. 2004) followed, and was shown to originate outside the Roche lobe of the planet, leading to an interpretation as an escaping comet-like tail. Taking into account the effect of evaporation of planetary atmospheres could significantly reduce their expected lifetimes compared to earlier theoretical predictions (Baraffe et al. 2004), planets with masses below a critical (orbital distance dependent) value being eventually evaporated out of existence.

In the longer term, low resolution spectroscopy of terrestrial exo-planet atmospheres is seen as the most promising method to detect life on another planet, concentrating on specific signatures such as the simultaneous presence of CO₂, O₃ and H₂O in the atmosphere, which are thought to be explainable only by the widespread presence of oxygen-consuming life on the planet (Selsis 2002).

Peculiarity of the first OGLE planets: As shown by the red symbols at very short periods on the top-right panel of Figure 1.2, the first three 'OGLE planets' (planets initially detected via their transits in OGLE data) occupy a distinct region of the mass-period diagram, at very low periods and fairly low masses. Their density is markedly higher than that of less close-in hot Jupiters (Konacki et al. 2003a; Bouchy et al. 2004; Konacki et al. 2004; Moutou et al. 2004b), and their predicted mass loss rates through evaporation up to four times higher than that of HD209458b (Moutou et al. 2004b), which may place them in a critical stage of their evolution.

The two more recently discovered transiting exo-planets (TrES-1 and OGLE-TR-111b, also shown in red on Figure 1.2) have orbital characteristics more similar to the hot Jupiters discovered via the RV method (with orbital periods of 3.03 and 4.0 d respectively), and have been presented as the missing link between RV and transiting planets. Importantly, their derived radii are closer to those of the three very close-in OGLE planets than to that of the first known transiting exo-planet, HD 209458b. These lower radii imply higher densities and may explain why so few planets to date have been found via the transit method: they may, on average, cause shallower transits and hence be harder to detect, than had initially been predicted on the basis of observations of HD 209458b's transits.

Improved statistics: We have seen that the detection of transits of a given planet paves the way for a wealth of detailed follow up observations, which have wide theoretical repercussions. The simple fact of increasing the number of known planets is also important. Many of the trends discussed in Sections 1.1.3.1 and 1.1.3.2 are still uncertain due to the low numbers involved. If the transit method lives up to its expectations, its ability to be applied to many thousands of stars simultaneously should lift the field of statistical studies of exo-planets out of its infancy: it will help place rigorous constraints on theoretical models by fitting the shapes of distributions of the various parameters, and break the potential degeneracies between the many parameters involved.

Exploring new regions of parameter space: Ground based transit and RV searches probe fairly similar regions of the mass-period diagram, but the parent stars of planets discovered via transits can be much more distant. The increased photometric precision achievable from space will also allow space-based transit searches to reach much smaller radii, hence masses. The absence of interruptions in the observations will also lessen the decrease in sensitivity towards longer periods, which is a problem for ground-based searches. The sensitivity of the method as implemented in various current and planned projects will be discussed in more detail in Section 1.2.

1.2 Current and planned transit searches

1.2.1 Ground-based

1.2.1.1 Projects and expected detection rates

The large number of ground-based transit search projects either recently completed or currently operating makes it impossible to discuss them all here in detail. Instead, those which have to date produced particularly important results, or with which I have had some interaction, are used to illustrate general characteristics of these projects as a whole. The hardware in use varies widely, with telescope apertures between 5 cm and 4 m, but as pointed out by Horne (2003) the projects can be crudely divided into two broad categories: 'wide' and 'deep'.

The 'wide' surveys make use of (often dedicated) small telescopes (0.5 m and below) with a very wide Field Of View (FOV) and large pixel scales to cover large areas of the sky. Given their small apertures, they can only detect giant planets out to a few hundred pc, but they can target any part of the sky while remaining within the galactic disk scale height. STARE, a 10 cm aperture telescope with a ~ 33 sq.deg. FOV, currently operating in Tenerife, was the first to observe a planetary transit (that

of the previously known RV planet HD209458b, Charbonneau et al. 2000). Two such projects to which I aim to apply the transit search tools developed in this thesis in the near future are: SuperWASP, consisting of eight cameras with 11 cm aperture, observing adjacent parts of the sky, each with a FOV of over 60 sq.deg., which started regular operations in La Palma in April 2004; and the University of New-South Wales planet search on the Automated Patrol Telescope at Siding Springs Observatory, a 0.5 m telescope with a 6 sq.deg. FOV, which has been monitoring 5 fields over the past three years. A large number of variable stars, periodic and aperiodic, were detected in light curves from 5 nights of APT data in the field of the open cluster NGC 6633 using software developed by M. Irwin and myself (Hidas et al. 2003).

The 'deep' surveys make use of larger telescopes (1 m and above) with less extensive fields of view (and usually available only for limited duration observations). They are capable of detecting hot Jupiters out to several kpc, so that they target specific areas of high density in the Galactic plane or in stellar clusters. Of note among these is the OGLE III project, which produced the first transit candidate to be confirmed as a planet by RV measurements (Udalski et al. 2002b; Konacki et al. 2003a). It uses the 1.3 m Warsaw telescope at the Las Campanas Observatory in Chile, and has a 0.35 sq.deg. FOV. A dataset obtained by the University of St Andrews Planet Search, which had already been analysed for variable stars and transits (Street et al. 2002, 2003), was used to train and refine variable star and transit search tools presented in this thesis. It consists of approximately 20 nights of observations of open cluster NGC 6819 using the Wide Field Camera (0.3 sq.deg. FOV) on the 2.4 m Isaac Newton Telescope in La Palma.

The expected hot Jupiter discovery rate of a number of these projects was compiled by Horne (2003), with results typically in the range 3–10 per month for both categories, though this estimate does not take into account the effect of 'impostors' (events that imitate planetary transits) or crowding (which limits the photometric accuracy).

1.2.1.2 The challenges ground based transit searches

Data storage and memory We have already emphasised the fact that, among the recent technological developments that have made planet-searching via the transit method possible, improvements in data storage and processing have a determining place. Despite these improvements, this area is still problematic. Data are usually stored onto tapes at the telescope, and reading the tapes onto disk is often the main bottleneck in the production of light curves (Mallén-Ornelas et al. 2003). Restricted disk space availability affects the production of light curves as well as their archival and retrieval for analysis. A significant part of the financial resources of many ground-

based transit searches have to be channelled into the acquisition of storage and computing capacity.

Achieving high precision photometry over very wide fields: Ultra-wide field detectors with large pixel scales exacerbate the problems of standard CCD photometry (see for example Bakos et al. 2004). The impact of variable extinction across the field and of crowding is much increased. The very short focal lengths of many of the small telescopes in use also imply significant vignetting and image distortion across the field. In these conditions, particular care must be paid to the extraction of flux measurements and the corrections applied for systematics.

Different teams have opted for different photometry techniques (PSF fitting, difference imaging, aperture photometry). The results are constantly improving, and several projects are now achieving relative photometric precision of 2–3 mmag over several nights of data at the bright end of the range of magnitudes surveyed (Hidas et al. 2003; Lister & the SuperWASP collaboration 2004).

Most of the data processing pipelines include a post-processing stage to minimise systematics in the light curves, due to the variation of image and extinction parameters (in time and across the field), and the difficulty of calibrating data from several nights or runs to the same magnitude system. This stage is sometimes considered separate from the processing pipeline and applied as a pre-processing step for the transit search (Kovács et al. 2004).

The transit detection itself: The automated detection of transits in light curves from ground based experiments has proven more difficult than initially expected, so much so that some groups have opted, at least initially, for the examination of the light curves by eye (see for example the EXPLORE team, Mallén-Ornelas et al. 2003). This is a very time-consuming process when tens of thousands of light curves containing hundreds of data points each are concerned, and one which is probably less sensitive than searches performed by algorithms, and affected by intractable biases.

The irregularity of the sampling is an unavoidable consequence of ground based observations. Unless they are coordinated from multiple sites around the world, which requires large, well organised collaborations and can only be achieved for limited periods of time, this problem complicates the analysis. The transit search algorithms developed in the present thesis were designed with irregularly sampled data in mind.

Another problem is the definition of a reliable threshold for detection. The most interesting events – those most likely to be of planetary origin – are also often those with the lowest significance, but the cost of follow-up observations means that false

alarms must be minimised. In white Gaussian noise, thresholds that give the desired false alarm rate can be estimated a priori (Jenkins et al. 2002) or found through Monte Carlo simulations. In the presence of unknown or complex noise characteristics, the best approach is to perform a posteriori completeness tests by inserting artificial signals into real datasets and attempting to retrieve them.

Follow-up of transit candidates: We have seen that the RV method can only monitor a limited number of stars. However, it is necessary to use this method eventually to confirm the planetary nature of transit candidates, as a number of more massive objects can mimic planetary transits. For example, brown and white dwarfs have radii similar to planets, and hence similar signatures. Stellar eclipsing binaries with a very large (early-type or giant) primary and a late-type dwarf secondary can also have eclipse depths of only a few percent, consistent with planetary transits, as can grazing binaries of higher mass ratios, or eclipsing binaries whose eclipses are diluted by a third star coincident with the line of sight. All these eventualities must be excluded, for example by measuring the mass of the secondary with the RV method. However, to avoid making excessive demands on highly sought after RV facilities, the identification of transit candidates must be as reliable as possible, and follow-up must be done in a rational and careful fashion. In some cases, the targets stars of transit searches based on relatively large telescopes may be too faint for the companion mass to be determined by RV observations, though upper limits may be placed. The ‘wide’ surveys which target brighter stars do not suffer from this limitation.

1.2.1.3 Published planetary transit candidates and their follow up

The present status of the many ground based projects in operation is hard to gauge, because many – though not all – choose to publish their transit candidates only after follow-up observations are complete.

Two significant non detections are worth mentioning first of all.

Gilliland et al. (2000) found no transit candidates in 8.3 d of Hubble Space Telescope (HST) observations in the field of the globular cluster 47 Tuc. This is, of course, a space-based project, but it is discussed in the present section because the range of parameter space probed corresponds to that of other ground-based searches rather than that of future space missions. The relative photometric precision, which reached 3 mmag at the bright end of the range of magnitude surveyed, was amply sufficient to detect close-in giant planets. The absence of detection implies that the frequency of hot Jupiters in 47 Tuc is at least an order of magnitude lower than that deduced from RV surveys in the solar neighbourhood. One possible interpretation

is that planet formation may be inhibited in low metallicity environments (47 Tuc has $(\text{Fe}/\text{H}) = -0.7$).

The low mass binary CM Draconis has its orbit almost aligned with the line of sight, and therefore makes a very favourable transit search candidate. Possible eclipses by a planetary mass objects were announced by Guinan et al. (1996), but detailed monitoring by the TEP (Transits of Extrasolar Planets) project (Doyle et al. 2000) ruled out the presence of planets orbiting the binary with radii larger than $3 R_{\oplus}$ and periods of 60 d or less at the 90% confidence level.

The ‘founding fathers’ of the field were the Vulcan and STARE projects, which were the inspiration for a plethora of small aperture wide angle searches. The Vulcan project uses a dedicated 5 cm aperture photometer equipped with a 49 sq. deg. FOV CCD detector. Monitoring a number of fields, each containing ~ 6000 stars, at the rate of 8 observations per hour for ~ 3 months each, they have detected and carried out spectroscopic follow-up on several tens of eclipses with depths of a few percent to date (Posson-Brown et al. 2000; Borucki et al. 2001; Latham 2003). Most of these were shown to be due to high mass ratio eclipsing binaries or triple systems where the light of the third star diluted the eclipses of the first two. The follow-up was inconclusive in a few cases, where no RV signal was detected – which leaves the possibility of a planetary mass companion open – but the transits were relatively low signal-to-noise ratio events and therefore need confirmation.

The STARE project, as already mentioned, was the first to detect a planetary transit, albeit in a system already known from its RV modulations (the transit was observed almost simultaneously by Henry et al. 2000a, and was found in Hipparcos data by Robichon & Arenou 2000). Since then, the STARE team have discovered a number of candidate δ -Scuti stars (Alonso et al. 2002), but until July 2004, they had announced no confirmed detections.

The most widely publicised project, of course, was the one whose candidates were the first to be confirmed. That example also serves to illustrate the complexity of the follow-up process.

The OGLE collaboration has announced 137 candidates to date, from the 2001 and 2002 seasons of monitoring in the Galactic centre direction and in the Carina region of the Galactic disk (Udalski et al. 2002a,b,c, 2003). Konacki et al. (2003b) followed-up the initial 59 galactic centre candidates, of which 8 showed clear signs of a luminous (hence stellar) secondary outside the eclipses, one was shown to be a duplicate entry, 4 had only one transit and hence no known period, and 7 were considered too faint for spectroscopic follow-up, leaving 39 for spectroscopic follow-up. Relatively low resolution spectroscopic observations with several epochs showed

that 8 were early type (A–F) primaries, ruling out a planetary radius for the transiting objects, and 25 showed radial velocity variations at the km s^{-1} level or double lined profiles which indicated they were grazing binaries. They obtained high precision radial velocity measurements with Keck/HiRes for 5 of the remaining 6 candidates, and concluded that two were likely to be binary star eclipses diluted by a third star (blends), including one which was suggested as a good planetary candidate by Dreizler et al. (2002) in a parallel investigation, but whose OGLE light curve showed hints of a secondary eclipse; two showed no radial velocity modulations and remained possible planetary candidates or blends; and one, OGLE-56-TR, showed significant RV modulations consistent with a planet (Konacki et al. 2003a). This last result was confirmed by additional Keck/HiRes observations by Torres et al. (2004). Similar follow-up observations were obtained for the Carina field candidates, in which Bouchy et al. (2004) reported two confirmed planets, OGLE-TR-132 and OGLE-TR-113 (the latter also independently confirmed by Konacki et al. 2004).

A number of other groups have presented transit candidates together with subsequent follow-up observations: an ADS search reveals results from the Survey for Transiting Extrasolar Planets in Stellar Systems (STEPSS, Burke et al. 2003), the EXPLORE project (Mallén-Ornelas et al. 2003), the University of St Andrews Planet Search (Street et al. 2003), a Danish group working with the NOT telescope (Bruntt et al. 2003), and the Berlin Exo-planet Search Telescope team (Rauer et al. 2004). However, none have confirmed detections: all candidates are either identified as stellar or substellar from the spectroscopic follow-up, or in need of additional observations to confirm the existence of the transits. It is likely that several more groups are currently investigating transit candidates, the chances of a true detection improving with time as observing, analysis and follow-up strategies are refined. For example, the UNSW transit search group are analysing spectroscopic data of 5 candidates in the field of NGC 6633 at the present time.

On August 25th 2004, the TrES network (to which STARE belongs) announced the first confirmed detection of a transiting planet with small aperture, wide field telescopes (Alonso et al. 2004a). A few days later, the planetary nature of another OGLE candidate, OGLE-TR-111, was confirmed by Pont et al. (2004).

1.2.1.4 Lessons from the results to date

Wide versus deep: Although the first confirmed detections have emerged from a ‘deep’ survey, a consensus seems to be emerging that the highest numbers of detections overall are to be expected from the ‘wide’ surveys using dedicated telescopes (Pepper et al. 2003). This can be explained in simple terms from the expected detection rates compiled by Horne (2003): the detection rates per month are similar

for both categories, but it is very rare to gain access to a 2 or 4 m telescope for a whole month, let alone the round-the-year access that is permitted by dedicated instruments. Indeed, the OGLE survey, the most successful to date, has exclusive access to the 1.3 m telescope it is operated on.

The role of the 'deep' surveys, in this context, is to test the frequency of close-in giant planets in specific high density environments such as stellar clusters, in which the high angular resolution achieved by larger telescopes is required. These targets constitute laboratories with relatively well-controlled parameters such as stellar age, metallicity, spectral type and density. The detection rates and the characteristics of detected planets in these environments can then be compared to those from RV surveys and the 'wide' transit surveys in the field.

Observation strategy and false alarm rates: The ratio of confirmed detections to announced transit candidates so far is rather low: 3 to 137 for OGLE, zero for many other projects. There have been suggestions that this is in part due to imperfections in the selection of transit candidates (Drake 2003; Sirko & Paczyński 2003), several of the OGLE candidates being single, partially sampled events at the beginning or end of a night, or showing clear signs of out-of-eclipse variations (the tell-tale sign of a stellar companion).

Sampling and phase coverage are therefore of prime importance. To exclude the effects of night-to-night variations and imperfect extinction correction, at least one of the ingress or egress plus (part of) the flat bottomed portion of a candidate transit event must be sampled. As pointed out by Seager & Mallén-Ornelas (2003) (and discussed in a simplified context in Section 6.4.3.3), it is possible to place limits on the density of the eclipsed star from the light curve alone (thereby excluding at least some stellar binary events) provided the time sampling is high, the photometric errors small and at least two complete transits are observed.

A fraction of the OGLE false alarms were due to high mass-ratio or grazing eclipsing binaries. Some of these might have been excluded on the basis of the eclipse shape, but the signal-to-noise ratio in the light curves was not always sufficient for this. Multi-colour observations during the eclipse would exclude all but the equal mass stellar binaries, and newly set-up transit searches as well as planned space missions are increasingly considering the use of at least two bandpasses as standard.

A significant fraction of the false alarms found by several projects were due to blends, as noted by Brown (2003); Alonso et al. (2004b). These cannot always be excluded even with multi-colour time series, and the probability of their occurrence increases with the number of background stars within the magnitude range of the

survey and the effective area of the aperture or PSF used to perform the photometry. Galactic plane or galactic centre surveys are thus likely to be prone to this type of false alarm to a greater extent than those pointing out of the plane.

Efficient follow-up strategies: Surveys can minimise false alarms by optimising their observation strategy – target field, time sampling, use of colour – or data analysis tools, but a carefully designed follow-up strategy can also weed out most of the non-planetary companions early on, thereby minimising the requests made of large facilities necessary in the last stages of confirmation via RV measurements.

Early transit candidate list announcements generated a lot of excitement in the exo-planet community, and a rush to RV facilities ensued. Experience has shown that many of the candidates could have been excluded on the basis of less time-consuming observations from smaller telescopes.

Given a transit candidate, there are a number of checks to perform. They are summarised in Table 1.1, together with the type of observations required. Any of the characteristics listed in the table can bring into question the planetary nature of the candidate, some ruling it out altogether.

Table 1.1: Tests of transit candidates and necessary observations.

Test	^a	Observations
Out-of-eclipse brightness variations	×	Existing light curves
Repeat transits at the predicted times	✓	New photometric observations
Different transit depth in different bandpasses	×	Multicolour transit photometry
Several stars within the original target image	×	Higher resolution imaging
Early-type primary	×	Single low to medium resolution spectrum or colours from existing sky survey images
RV modulations of several km s^{-1} (stellar companion)	×	Medium resolution spectra spread over orbital period
Double lined profile	×	Medium to high resolution spectrum
Very narrow line profile (giant primary)	×	Medium to high resolution spectrum
RV modulations of a few 100 m s^{-1} (brown dwarf companion)	×	High resolution spectra spread over orbital period
RV modulations of a few 10 m s^{-1} (absence can imply a faint background eclipsing binary rather than a planet)	✓	Very high resolution spectra with special wavelength calibration (iodine cell) spread over orbital period

^aTicks or crosses in the central column indicate whether the tests are meant in a positive or negative sense i.e. whether true planetary transits should (✓) or should not (×) exhibit the listed characteristics.

These checks can be performed in various orders, but it is most efficient to go from the smallest telescopes to the largest, and from imaging to low and then high

resolution spectroscopy. A 4 stage follow-up strategy has been proposed for Super-WASP (Horne, priv. comm.) including checks on the existing light curves, multicolour photometry from a 1–2 m telescope, low to medium resolution spectroscopy (with a few epochs) from a 2–4 m telescope, and repeated high resolution spectroscopy with a 4–8 m telescope. Some projects have opted for a small dedicated follow-up telescope to perform multi-colour photometry and check the repeatability of the transits rapidly (Kotredes et al. 2003).

1.2.2 Space-based transit searches

1.2.2.1 Why go to space?

Space-based observations are free from the noise induced by atmospheric scintillation. One of the best results in terms of time series photometry achieved from the ground was a coordinated monitoring campaign of M67 using several large telescopes (Gilliland et al. 1993). The precision reached over the entire monitoring time of 156 hrs was 20 ppm (parts per million). By contrast, the precision required to detect an Earth like planet, which causes 13 hr long transits of depth 84 ppm, with a signal-to-noise ratio of at least 6 over three transits, is 14 ppm over 39 hrs, i.e. 7 ppm over 156 hrs. The precision achievable from the ground is thus insufficient to detect terrestrial planets. Consequently, going to space is the only means of achieving the photometric precision required to detect terrestrial planets, which have much smaller radii than their gaseous giant counterparts – recalling that transit depth scales as the inverse of planet radius squared (for a given stellar radius).

Another important reason is the absence of significant interruptions in coverage, provided a suitably chosen combination of satellite orbit and target field is used. As the orbital period of a planet increases, transits become rarer for two reasons: first is the obvious one, that is that a transiting planet only causes one transit per orbit; the second is the decrease in the alignment probability for wider orbits (see Equation 1.3), which goes from $\sim 10\%$ at 0.03 AU to $\sim 1\%$ at 1 AU (for a Jupiter-sized planet orbiting a Sun-like star). The probability of observing more than one transit of a given system – a requirement for the reliable detection of shallow transits – is thus a very strongly decreasing function of orbital period if the observations are confined to observing runs of limited duration, themselves with daily interruptions. Uninterrupted runs of several months with continuous high time sampling, as can only be achieved from space, thus provide a much more favourable platform for the detection of planets orbiting at a greater distance from their star, i.e. planets in Earth-like orbits, opening the possibility of detecting *habitable* planets.

Existing space-based optical imaging facilities such as HST are primarily aimed

at high resolution applications and have fields-of-view that are too narrow for transit searches in all except the most concentrated environments (such as globular clusters). It is therefore necessary to launch dedicated transit search missions with wide fields of view and orbits that allow prolonged continuous viewing of a given field. The basic design of these missions involves no significant technological development: relatively small telescopes and wide field optical CCD detectors with low dark current and readout noise have been in existence for several years, so the cost of the missions can be kept reasonably low.

1.2.3 Terrestrial and habitable regimes:

At this stage, it is useful to establish a working definition of a terrestrial and a habitable planet. We follow here the reasoning of Borucki et al. (1997) – there are fluctuations in the limits of both regimes as quoted in the literature, but they do not significantly differ from those given here.

A terrestrial planet is one whose mass is insufficient to trigger runaway accretion of H and He. This is thought to occur for $M_{\text{pl}} \geq 10 M_{\oplus}$, which corresponds to an upper radius limit for terrestrial planets of $R_{\text{pl}} \sim 2.2 R_{\oplus}$.

For a planet to be habitable, it is additionally required to have a mass that is sufficient to sustain plate tectonics, as these are thought to have played an important role in the evolution of life on Earth, and to retain a significant atmosphere. This results in a lower mass limit of $M_{\text{pl}} \sim 0.5 M_{\oplus}$, i.e. a lower radius limit of $R_{\text{pl}} \sim 0.8 R_{\oplus}$. It must also lie in the ‘habitable zone’. This zone is usually defined to ensure the existence of liquid water on the surface of the planet, the inner edge being due to loss of water through photolysis and hydrogen escape, and the outer edge to the formation of CO_2 clouds in the upper atmosphere, which increase the albedo and thus lead to a drop in surface temperature (Kasting et al. 1993). Franck et al. (2002) compute age-dependent habitable zone limits, requiring the surface temperature to be in the range 0 to 100°C and taking into account the change in luminosity of a star while on the main sequence. They give the following orbital distance ranges for the habitable zones of 4.5 Gyr old stars: 0.5 to 1 AU ($0.8 M_{\odot}$, K2 star) and 0.9 to 1.3 AU ($1.0 M_{\odot}$, G2 star). A similar but simplified calculation (K. Horne, priv. comm.) using a single luminosity value, and converting to period using the mass range 0.5 to $10 M_{\oplus}$, yields the periods corresponding to the centre of the habitable zone as 1.2, 0.6 and 0.3 years for G2, K0 and K5 stars respectively.

A full definition of habitability must take into account many additional factors. To cite but one example, the interplay between stellar and planetary magnetic field (if the latter is present) is thought to play a major role in the survival of the planet’s atmosphere, and hence its habitability. However, in the first approximation the limits

established here provide a suitable framework to explore the detectability of terrestrial and/or habitable planets by the planned space-based transit search missions.

1.2.3.1 Challenges

Ultra-high precision photometry from space: Space missions are free from the effect of seeing variations (in space and time) and atmospheric extinction, which enables them to achieve much higher photometric precision than ground-based projects, reaching 15 to 20 ppm over the duration of an Earth-analogue transit for the larger projects. However, other noise sources come into play and must be taken into account and, where possible, corrected. The most significant are satellite jitter and the hostile radiation environment the detectors are submitted to.

On-board processing power and telemetry limitations: Telemetry limitations imply that it is impossible to download every image in full. Instead, some of the processing must be done on board and light curves rather than images are downloaded. On the other hand, on-board computing capabilities and power budgets are also limited, so that the processing steps must be kept as simple as possible, and only a limited number of targets per field can be monitored. These problems affect all the planned transit-search space-missions to a varying degree. For example, the telemetry limitations are even more stringent if the satellites are in distant (L2) orbits, which is necessary to allow continuous monitoring of a given field for more than 6 months.

Stellar micro-variability: With the improved precision achievable from space comes a drawback: increased sensitivity to small amplitude variations in the brightness of the parent star itself. The Sun is observed to vary on all timescales from minutes to decades with amplitudes of up to a few mmags, and every star is expected to display some degree of variability if observed with sufficient precision.

The lack of knowledge available about the observational characteristics – amplitude and frequency distribution for example – and physical mechanisms leading to this complex noise source, which is referred to hereafter as stellar micro-variability, makes it potentially very dangerous for transit searches for small planets. Figure 1.4 illustrates how, if left untreated, it can completely drown out the transit signal. It must therefore be taken into account at the observational strategy, target selection and data analysis stages.

Difficult follow-up: In addition to the relatively faint magnitude of the target stars, which already affects the follow-up of planet candidates from ground-based tele-




Figure 1.4: Top panel: simulated transits of an Earth-sized planet orbiting a Sun-like star with a 30 d period. Middle panel: a 60 d section of the Total Solar Irradiance (TSI) variations as measured with the PMO6 radiometer on SoHO/VIRGO near the peak of the solar activity cycle. Bottom panel: combined solar and transit light curves. The transits are no longer recognisable to the naked eye.

scopes, the lower masses and longer periods of the planets that can be detected via the transit method from space make them very difficult to follow up via the RV method. The fact that transits as shallow as a few parts in 10^5 are detectable means that a background binary can cause a detectable planetary transit-like signal even if it is fainter than the primary target star by 5 magnitudes or more: the necessity to follow-up candidates is thus even more acute than from the ground.

The best ground based RV facilities can achieve precisions of down to 1 m s^{-1} (HARPS, VLT+UVES, Keck+HiReS). This is enough to measure the signal induced by 'Hot Earths' – terrestrial planets in orbits lasting a few weeks or less – provided the target stars are bright enough – typically $V = 12$ to 14 with HARPS. For fainter targets or planets with lower masses or longer periods, only upper limits will be achievable: enough

to exclude stellar or substellar companions, but not to provide mass measurements or to fully exclude all possible stellar blend scenarios. A large fraction of the transit detections made by the larger space-based transit search projects – *Kepler* or *Eddington*– will thus not receive full confirmation from RV observations.

This has two important implications. The first is that the candidates must be thoroughly screened using all available information (in the light curve and from other photometric and spectroscopic observations taken as part of the mission preparation) before requiring additional follow-up observations. The other is that the scientific return of these missions will likely fall into two distinct categories: detailed studies of confirmed planet candidates on the one hand, and statistical studies of the entire sample of candidates, some of which may remain unconfirmed, on the other.

1.2.3.2 COROT

COROT (CONvection, ROTation & Transits) is a small space mission whose primary science goals are to perform asteroseismology of stars across the HR diagram and to detect exo-planets via the transit method. The mission design is described in an article by Baglin & the COROT Team (2003) and on the project website⁵. It is funded mainly by the French space agency CNES, with significant contributions from the European Space Agency (ESA) and from other individual countries (Austria, Spain, Germany, Belgium and Brazil). The third in the PROTEUS series of small CNES missions, it is on track for its planned launch date of June 2006.

The payload consists of a small (27 cm aperture) afocal telescope shielded by a very efficient baffle to minimise Earth-scattered light. The focal plane is occupied by four 2048×2048 pixel EEV CCDs covering a total field of view of $2.8 \times 2.8^\circ$, two of which are devoted to asteroseismology and two to exo-planet search (the two programs are carried out simultaneously). It will be launched on an almost polar near-Earth orbit with an altitude of ~ 900 km.

The seismology CCDs will be highly defocussed to allow the observation of bright stars ($6 \leq V \leq 9$), each projecting a ring-like image over ~ 400 pixels. They will be read once per second in frame-transfer mode. The exo-planet CCDs will be slightly defocussed, and the presence of an objective prism in the optical path of that part of the field of view will lead to tear-drop shaped images occupying between 50 and 100 pixels. The exo-planet CCDs will be read every 32 s, and the target stars in that program have $12 \leq V \leq 16$. For the brighter stars, the dispersed PSF will be split into blue and red regions to provide coloured light curves, for the fainter stars only white-light information will be collected.

During its 2.5 yr lifetime, COROT will carry out 5 long (150 day) observing runs,

⁵www.astrsp-mrs.fr/projets/corot/

designed to allow the detection of planets with orbital periods up to a few months, as well as a number of shorter (~ 20 day) runs devoted primarily to seismology or additional science. Due to its orbit, COROT can only follow a given field for up to 6 months, after which it has to be rotated to keep its solar panels facing the Sun. A given field can only be observed if it is at a safe distance from both the Sun and the bright rim of the Earth. This limits the observed fields to a region close to the ecliptic plane. Due to the alternation between one direction and the other every 6 months, they must be contained in two circles of $\sim 10^\circ$ radius on opposite sides of the sky within this region. These 'COROT eyes' were chosen to be centred on galactic coordinates $l = 35^\circ$, $b = 0^\circ$ (close to the galactic centre) and $l = 210^\circ$, $b = 0^\circ$ (close to the anticentre). The low galactic latitudes allow for large numbers of stars in the field, increasing the planet detection probability.

The small number of these long runs makes the choice of the target fields very important. This is a complex process, simultaneously driven by the requirements of seismology and exo-planet programs, and based on a very substantial ground based preparatory observation program. Spectra of every star with $V < 8.5$ in the COROT eyes have been taken and primary seismology targets proposed. An on-going multi-colour photometric survey of the potential exo-planet field area next to each proposed primary target (including all translations and rotations of the field keeping the primary target in the seismology CCDs) provides extinction maps and preliminary spectral types and luminosity classes for the potential exo-planet search targets. These are used to optimise the position of the exo-planet CCDs and to exclude a number of proposed primary targets surrounded by regions where the extinction is too variable. The final CCD positions on the sky also take into account secondary seismology targets. At the time of writing, the CCD positions have been chosen for at least two fields (one in each direction), and the final primary target choice and positions for all 5 long runs should be decided by the end of 2004.

Due to telemetry limitations, images from the exo-planet channel will be co-added to 8.5 min sampling, and photometry performed on board. Up to 12 000 stars can be monitored in each field. In the Galactic anticentre direction, the number of stars per field in the magnitude range of interest is lower than this limit, but the higher stellar density in the Galactic centre direction means that only some of the stars that fall in the appropriate magnitude range in a given field will be monitored in the Galactic centre direction. The preliminary spectrophotometric classification will be used to attempt to preferentially select late-type dwarfs in these cases. At the beginning of each run, a mask optimally encircling the tear-drop PSF (which varies according to brightness and temperature) will be chosen automatically for each star from a template set, and used throughout to perform white-light photometry by

summing all the pixels in the mask, or splitting the area within the mask in two in order to get two 'colours'. The resulting light curves will then be downlinked to Earth and noise corrections applied. A small number of stars (typically a few hundred per field) will be observed with higher (32 s) time sampling. The list of 'oversampled' targets will contain additional science program stars, but will be updated regularly to include stars showing suspected transits.

The exo-planet detection potential of COROT has been investigated by Bordé et al. (2003), who estimate maximum orbital distances at which planets can be detected as a function of stellar spectral type, and the number of expected detections as a function of planetary radius and surface temperature, showing that planets with radii of $1\text{--}2 R_{\oplus}$ are within COROT's reach provided their surface temperatures are $\geq 500\text{ K}$. COROT will thus be the first mission to probe the terrestrial regime, as well as detecting large numbers of gaseous giant planets and, if they exist, close-in giant planets whose atmosphere has been (partially) evaporated by strong high-energy flux from the star (Lammer et al. 2003) or ice-planets with liquid outer layers (speculated 'ocean' planets, Léger et al. 2004).

Although it is less ambitious than the larger missions that will follow it, COROT benefits from a significant advantage: it will be possible to follow-up the candidate planets it will detect with the most sensitive of the ground-based RV instruments, such as HARPS on the 3.6 m ESO telescope at La Silla (Pepe et al. 2004), which can achieve $\text{sub-}1\text{ m s}^{-1}$ precision, sufficient to study most of the COROT candidate planets.

COROT will also have another, very important role: it will be the first mission to measure stellar micro-variability – that is intrinsic variability with mmag amplitudes and timescales of minutes to weeks – for stars across the HR diagram. Although the smaller Canadian mission MOST, currently operating, is for the first time measuring this type of variations in detail for other stars than the Sun, it will do so only for a small, selected sample of stars. The COROT dataset will be invaluable for the study of the physical processes behind stellar micro-variability, a project which can be carried out at no extra cost as the exo-planet field observations can be used directly (Weiss et al. 2004). It will also be extremely useful to optimise the light curve filtering tools needed for later transit search missions concentrating on habitable planets, where stellar micro-variability will play a major role.

1.2.3.3 *Kepler*

Kepler is a NASA Discovery class mission designed to monitor 100 000 stars for 4 yrs in order to detect transits of Earth-sized and larger planets with orbital distances as large as 1 AU. It has a planned launch date of 2007 and will be launched into an Earth-trailing heliocentric orbit. The mission is described in detail in a number of pub-

lications including (Borucki et al. 2004) and on the project website⁶.

The design was guided by the requirement that a Sun-Earth system with $V = 12$ should cause at least an 8σ detection. For a 4 yr mission lifetime, this implies a 4σ event for a single transit.

The payload consists of a modified Schmidt telescope including flat-fielders near the focal plane. The corrector has a clear aperture of 0.95 m and the primary a diameter of 1.4 m. This aperture ensures a precision of 1 part in 10^5 for a $V = 12$ G2 dwarf, which satisfies the design requirement above. The detector is composed of 42 1024×2200 pixel CCDs, totalling a 100 sq. deg. FOV. Only white light observations are planned.

Telemetry limitations imply that only pixels illuminated by one of the 100 000 target stars are saved for transmission to Earth, before which images are co-added to produce one measurement per pixel every 15 min. A single field is monitored for the entire duration of the mission. Consequently, the choice of target field is of crucial importance. It is constrained by a 55° sun-avoidance angle. The *Kepler* team have opted to maximise the number of stars in the magnitude range of interest in the target field, while fine-tuning of the detector position on the sky minimises the number of excessively bright stars on the CCDs. A field centred on $l = 68.8^\circ$, $b = 6.5^\circ$, falling in the Cygnus constellation, yielded 450 000 stars brighter than $V = 15$. A ground-based observation program to classify the stars in the field and identify suitable stars to constitute the 100 000 strong target list is underway. The magnitude range of the targets will be similar to that of COROT ($12 \leq V \leq 16$), though the photometric precision at a given magnitude will be much higher.

Classical (matched filter) detection algorithms are envisaged to detect planetary transits in the *Kepler* data (Jenkins et al. 2002). A wavelet-based algorithm designed to adaptively measure and remove non-white noise such as that induced by stellar micro-variability has also been developed (Jenkins 2002). Follow-up of candidate detections will be problematic, as not even the best ground-based RV facilities will be able to measure the Doppler shifts induced by Earth-mass planets. The follow-up procedure will therefore proceed as for other transit search programs (light curve examination, high resolution imaging / centroid analysis, low then high precision RV measurements) but the RV measurements will serve to exclude stellar or sub-stellar objects, and only upper limits will be available on the mass of planetary objects. It may be possible to detect colour changes during the transits using future instrumentation on HST and JWST (James Webb Space Telescope).

To ensure any announced detections are genuine, *Kepler* data will be analysed and transit candidates followed-up internally before the data is publicly re-

⁶www.kepler.arc.nasa.gov/

leased, which means that release dates will vary depending on whether the light curves contain clear (deep, short period) transit signals (allowing quick confirmation and hence release approximately one year after the first few transits are detected), are too variable for transits to be detected (allowing release soon after) or contain shallow or long period events (released after the end of the mission only).

At the time of writing, *Kepler* has the best chance of being the first mission to detect habitable planets, and in particular Earth-analogues. It is expected (Borucki et al. 2004) to detect approximately 500 planets in the habitable zones of its target stars if every one of them has such a planet, with (under the same assumptions) 25 of these having the radii equal to or smaller than that of the Earth, and half of those orbiting single stars.

1.2.3.4 *Eddington*

Eddington resembles COROT in that it combines the two science goals of asteroseismology and planet finding via the transit method, though it is designed on a more ambitious scale. It was accepted as a reserve mission in ESA's science programme in September 2000, with two parallel primary science goals: asteroseismology and exo-planet detection via the transit method. The planned launch date was then very similar to that of *Kepler*. In December 2003, the mission was removed from ESA's Science program for budgetary reasons. Industrial and scientific studies related to the mission are still taking place and there is pressure from the exo-planet community to re-instate the mission should funds become available. Despite the uncertain status of *Eddington* at the present time, it is appropriate to describe it here as much of the work in the present thesis was carried out with applications to *Eddington* data in mind. Latest news on the project can be found on its website⁷.

The payload as planned at the time of selection is described in a report by Favata & the *Eddington* Science Team (2000) and consisted of a telescope with a 1.2 m diameter primary mirror, resulting in a collecting area of 0.6 m². The detector consisted of 20 740 × 2900 pixel CCD chips covering a circular FOV with 3° diameter. Hereafter, this is referred to as the old *Eddington* design. Following an industrial study, the planned payload was modified in early 2003, improving the photometric performance. The new design, described in Favata (2004), consists of three identical telescopes (pointing in the same direction) with a total collecting area of 0.764 m², each with identical detectors at the focal plane with a permanently covered circular FOV of 19 sq. deg. This modification allows not only shorter focal lengths and wider FOV, but the use of the already designed Herschel bus as the satellite platform. It also enables the use of colour information, which a number of studies recommended

⁷astro.estec.esa.nl/SA-general/Projects/Eddington

(see for example Kjeldsen & Tingley 2004), to be implemented in a simple way. Very broad bandpass filters, one excluding only the blue part of the spectrum and one excluding the red, are inserted in the optical path of two of the telescopes. A 'blue' and a 'red' light curve are then obtained for the bright stars by subtracting the light curve from the red-excluding or blue-excluding filter, respectively, from that from the third, white light telescope, while keeping the cost, in terms of photon noise, minimal for the fainter stars. The photometric performance resulting from the old design was used in the present thesis for work carried out in 2002 and before, while that resulting from the new design was used for work carried out in 2003.

The observation strategy was not significantly altered by the change in design. Of a planned lifetime of 5 yrs, 3 yrs are devoted to a single pointing for the planet-finding program, while the remaining two consist of several shorter runs primarily dedicated to asteroseismology (though also suitable to detect planets with relatively short periods). As for *Kepler*, the choice of the planet-finding target field is crucial and some of the simulations described in this thesis were designed to help place constraints on the most favourable stellar content of such a target field. The baseline Eddington orbit is around the L2 Sun-Earth point. Sun-avoidance constraints are similar to those of *Kepler*.

In the planet-finding mode, 8 s integrations are co-added to produce 10 min stacked images and photometry is performed on-board. System performance allows for up to 100 000 preselected planet-finding targets, which means that all the stars in the FOV which are bright enough to reach the photometric precision necessary to detect transits can be observed. The lack of target pre-selection may lower the detection efficiency, but allows for more serendipitous discoveries than the approach of the *Kepler* mission. The requirement of 80 ppm (parts per million) precision over 1 hr (necessary to detect an Earth-like planet with a 3 yr run) is achieved down to $V = 14$, while the magnitude limit of the targets is $V = 17$ (only larger or shorter period planets will be detectable around the fainter stars). The overall dynamic range is approximately $10 \leq V \leq 17$. The data release policy foreseen for *Eddington* was very open – immediate availability to all European teams having requested access in advance through response to an announcement of opportunity and public access one year later. All these characteristics make *Kepler* and *Eddington* complementary, despite having very similar photometric performances, time sampling and duration.

1.3 The broader perspective

The development of tools to search for planetary transits naturally leads one to consider other phenomena which can be detected in the same kind of data as that in

which the transits are being sought, as well as the possibility of applying similar tools to other kinds of data.

1.3.1 The time domain in astrophysics

The time domain can arguably be considered the last (relatively) unexplored domain in astrophysics. That is not to say that the study of variations in the observable characteristics – position, brightness, and later spectra – of celestial objects have not played a crucial role in astronomy to date. The motion of the stars and planets across the sky inspired early conceptions of the cosmos, which were soon developed into the capacity to predict specific events – as exemplified by many archaeological remains – and later deciphered in terms of the governing laws of physics which drive the observed movements – starting with Newton’s interpretation of Kepler laws in terms of an inverse square law gravitational force.

However, not until very recently have homogeneous times series – of any kind, but most importantly for the present discussion photometric – taken by the same instrument and treated in the same way, been available as a platform for *new discoveries*. There have been *targeted* observations of objects that were known in advance to be potentially interesting targets, but these do not truly enter the ‘exploratory’ regime which is of interest here. On the other hand, the variable star community have, throughout the 20th century, used repeated observations of a given field to find new variable stars, but technological and organisational limitations have made the acquisition of homogeneous, long duration, high sampling time series difficult.

Two factors have changed this situation relatively recently. One is the advent of high quality CCD detectors which can be coupled to wide field instruments, together with sufficient data storage and processing capacity. The other, I would argue, is the discovery of the first exo-planets. This might be a controversial point. There is a wealth of other fundamental science that requires the same type of data, and for which scientists had been attempting for many years to set up dedicated projects, with success in some cases. For example, a number of very fruitful ground-based microlensing projects have been in operation for several years. However, in particular for space projects, exo-planet science appears to have been a very significant, often decisive factor. Both COROT and *Eddington* were – sometimes under other names – originally asteroseismology projects. However, it is the potential discovery of hundreds if not thousands of ‘brave new worlds’, which brought this type of observations to the attention of the public eye and created the necessary good will for these projects to be funded.

1.3.2 Additional science from transit search data

High precision, high sampling time series photometry with relatively long coverage are a treasure trove for the study of near-Earth objects, variable stars, gamma ray bursts and variability of extragalactic origin. There are too many types of variable stars to list here. They are not only identifiable from the shapes of their light curves but also from other parameters besides photometry. It is commonly thought that, if observed with sufficient precision, all stars would show some degree of variability. Any sample of light curves which is used to search for transits is therefore bound to contain many forms of variability other than transits. We have already seen that under certain circumstances stellar variability can hinder transit searches; it is also, of course, of interest for its own sake.

It is therefore desirable to identify variable stars, on the basis of global light curve statistics, and to perform a period search. Developing tools for this purpose was a logical extension to the work I have been doing on transit search algorithms. Though they are used to complement transit search methods, they provide a first pass ‘filter’ that can be followed by more detailed study. The elaboration of automatic variability classification tools is a very active subject (Mizerski & Bejger 2002; Belokurov et al. 2003, 2004), but it is beyond the scope of the present work.

1.3.3 Application of transit search tools to other data

Planetary transit search programmes ground- and space-based are not the only source of large quantities of photometric time series data. Scanning astrometry missions, starting with Hipparcos – whose dataset already constitutes a formidable database – and continuing with GAIA, will provide repeat observations of an even larger number of stars, albeit with much sparser sampling. The variable star and transit detection tools developed in the present thesis were therefore designed to be compatible with, or adaptable to, much sparser sampling than that commonly available for transit searches.

1.4 Structure of the thesis

Chapters 2, 3 and 4 describe the tools I have developed for the purpose of analysing transit search data: transit search algorithms, an empirical model of stellar micro-variability, and micro-variability filters to reduce stellar noise before applying the transit search algorithms.

The rest of the thesis is devoted to the application of these tools. In Chapter 5, various simulations are used to identify promising types of target stars and set

rough detection limits for *Eddington* (and by extension *Kepler*) and COROT. In Chapter 6, I describe my participation in the first COROT blind transit detection exercise, where several groups independently attempted to detect transits in a set of simulated COROT light curves whose content was unknown to the testers. Chapter 7 summarises early experiments with ground-based data from the University of New South Wales transit search program. Conclusions and future perspectives are presented in Chapter 8.

Chapter 2

Transit detection algorithms

The present time is an exciting one for the transit searching community. After a few years of tuning both observation and data analysis techniques, the first ground-based programs (in particular the Optical Gravitational Lensing Experiment, or OGLE) are now yielding detections, albeit in smaller numbers than expected (Konacki et al. 2003a; Bouchy et al. 2004). Is the heralded 'landslide' of extra-solar planets detected via transits really just around the corner? Will the planned space missions really deliver as expected? Their success hangs, at least partially, on our ability to automatically search huge datasets for tiny, repetitive signals deeply buried in noise. The work presented here represents an attempt to contribute to this challenging task.

At the time of starting the present thesis in 2001, there was no widespread agreement as to the best approach for planetary transit detection. However, a small number of papers dedicated to the problem (see e.g. Jenkins et al. 1996; Doyle et al. 2000; Gilliland et al. 2000; Defaÿ et al. 2001) provided a natural starting point from which to develop new methods, or improve on existing ones. By contrast, at the time of writing, the number of published articles devoted to or dealing with this problem has vastly increased (see e.g. Aigrain & Favata 2002; Jenkins et al. 2002; Kovács et al. 2002; Koen & Lombard 2002; Street et al. 2003; Aigrain & Irwin 2004). The references listed here are by no means exhaustive, but provide a good overview of the variety of techniques commonly used today or under investigation for future missions.

The first comparative studies (Tingley 2003a,b, see also Chapter 6) represent an attempt to synthesise and sort through this zoo of methods. One subject on which there is widespread consensus, is that a diversity of available algorithms is a better guarantee of success than a single, uniformly used approach. Not only is it likely that different methods will perform better in different circumstances, but the independent detection of a given event using two algorithms might improve the confidence level in the detection. On the other hand, most published algorithms are based on the same underlying maximum likelihood principles. Clarifying this common case helps

to understand the apparently disparate menagerie in terms of a family of closely related siblings, differing mostly in their implementation details – which affect speed, robustness, and adaptability to specific circumstances.

The algorithms presented in this chapter were conceived primarily with upcoming space-based planetary transit search missions in mind, such as COROT, *Eddington* and *Kepler*, but they are equally applicable to data from ground-based programs. Both types of datasets, despite the differences highlighted in Chapter 1, place very similar requirements on the detection algorithms.

The space missions are expected to produce 10 000's of light curves per target field, each containing 100 000's of individual observations, as the time sampling is high and a single field is observed for months or even years. Most small aperture ground-based programs have similar fields of view, or larger fields of view but brighter magnitude limits, leading to similar numbers of stars per target field. While the number of observations of each field is much less than for the space missions (100's or 1000's), it is usual for a number of fields to be observed in a given observing season, thereby leading to similar sized datasets in a given year.

The time sampling of the data from the space missions is expected to be regular, and the duty cycle very high (up to $\geq 95\%$ for *Eddington* and *Kepler*), though a number of data gaps, e.g. due to telemetry losses, cannot be avoided. In ground-based data the sampling is, of course, very irregular, with interruptions between nights and between observing runs, as well as due to weather or technical problems. In order to keep the algorithms as general as possible, they were designed to be applicable to data with irregular sampling. In some cases regular sampling could allow for small alterations to the code, leading to slight gains in computing time.

To tackle the enormous databases involved, speed and automation are vital. This leads, at least at the detection stage, to a single motto: the simpler the better. This is consistent with maximising the statistical efficiency (see Section 2.1.1.3). The present chapter reflects a learning process in this respect, in that the first algorithm investigated, a Bayesian approach, is more sophisticated, and for detection purposes less effective, than the second, a stripped down version where the use of Bayesian priors was dropped and the assumptions about the shape of the transit signal further simplified. This does not mean that sophisticated methods, incorporating as much prior knowledge about the sought-after signal and the noise characteristics as possible, should be discarded, but they are expected to come into their own at the characterisation, rather than detection stage – though that stage is beyond the scope of the present work.

Section 2.1 describes the Bayesian algorithm, together with performance tests which were carried out on simulated light curves affected by photon noise as ex-

pected for *Eddington*. Section 2.2 describes the simpler, but derivative, box fitting algorithm.

2.1 A Bayesian, step-function based algorithm

Transit detection algorithms based on a Bayesian approach, already investigated in the context of the COROT mission (Defaÿ et al. 2001), constitute an interesting alternative to more conventional approaches, e.g. matched filters. They maximise the use of whatever information is available on the phenomenon one is trying to detect, and are relatively flexible, allowing the seamless incorporation of new information into the detection process as it becomes available. While a global ‘marginalised’ statistic can be used for the detection, information is directly available to reconstruct the detected signal if wanted, therefore providing a tool to discriminate between planetary transits and other types of periodic signals (Defaÿ 2001), as well as directly measuring additional parameters such as the planet’s radius.

After a brief investigation of the method of Defaÿ (2001), a decision was made (for reasons outlined below) to develop a novel algorithm, based on the GL method of Gregory & Loredó (1992) (hereafter GL92), which was also the starting point for Defaÿ et al. (2001). The GL92 method was developed for the search of periodic variations in emission from pulsars in X-ray data.

The approach of Defaÿ et al. (2001) was based on the expansion of the light curve into a truncated Fourier series. C. Defaÿ kindly provided a coded implementation of this method, but experiments with this code showed that performing the detection in the Fourier domain made the algorithm computationally sensitive to data gaps and discrete sampling rates. The *direct space* approach investigated here is expected to be more robust, though it does not have the advantage of providing a direct means of reconstructing the shape of the detected signal.

The GL method was initially developed for Poisson noise dominated light curves (as is the case for X-ray pulsars) and later extended to the Gaussian noise case (Gregory 1999, hereafter G99). At the flux levels of interest for transit searches, the photon shot noise per detection element (which is Poisson distributed) can be very well represented by Gaussian noise. The original formulation of the GL method made no assumptions about the shape of the variations, the model consisting of a step-function with m even duration, arbitrary level bins. The new algorithm was developed to be as ‘general purpose’ as possible, but makes intrinsic use of prior knowledge of the expected signal, by allowing one of the bins to have a variable width, to represent the out of transit constant signal level. This formulation also permits the phase of the transits, or epoch (start time) of the first transit, to be identified, a task the original GL




Figure 2.1: Schematic illustration of the type of model used in the GL method. There are m equal duration bins per period p . The (arbitrary) epoch e is defined as the time elapsed between the start of the light curve and the end of the first bin.




Figure 2.2: Schematic illustration of the type of model used in the modified GL method. There are n in-transit bins of duration d/n plus 1 out-of-transit bin of duration $p - d$ per period p . The epoch e is defined as the time elapsed between the start of the data and the start of the next transit.

method is not suited for (see Section 2.1.1.1). The fitted parameters are the period, duration and phase of the transit. The shape of the transit can then be reconstructed from the phase-folded light curve.

The new algorithm, referred to hereafter as the modified GL method, is derived from the GL method in Section 2.1.1. Using simulated light curves, the modified GL method is compared to the original in Section 2.1.2, and its performance in the presence of photon noise evaluated in Section 2.1.3. Section 2.1.4 then outlines tests on light curves containing, in addition to simulated photon noise and transits, solar micro-variability.

2.1.1 Derivation of the algorithm

2.1.1.1 The models

In GL92 and G99, the periodic hypothesis is represented by a class of periodic step-function models, resembling a histogram, with m equal duration bins per period (see Figure 2.1). The model level in each bin is allowed to vary, giving these models the ability to describe light curves of unknown shape. A given model \mathcal{M}_m , with a given value of m , has $m + 2$ parameters: period p , epoch e (or phase ϕ) and m bin levels $R = \{r_1, r_2, \dots, r_m\}$.

This *periodic* hypothesis is contrasted with *aperiodic* and *constant* hypotheses, which are special cases of the periodic model, both having a period equal to the total duration of the light curve and the latter having $m = 1$.

In the case of transit searches, some information about the shape of the sought after signal is available. The model should therefore consist of a long flat section followed by a dip. A family of models similar to those of the GL method but where one bin is much longer than the others (see Figure 2.2) is therefore used. There are now $n+1$ bins, n being the number of *in-transit* bins. A given model \mathcal{M}_n , with a given value of n , has $n+4$ parameters: period p , transit duration d , epoch e , and $n+1$ bin levels $R = \{r_0, r_1, \dots, r_n\}$. Models with lower n will incur a lower Occam penalty factor, as emphasised in G99.

Only two hypotheses (transit, denoted by \mathcal{H}_T and constant, denoted by \mathcal{H}_C) are considered: the constant hypothesis is a special case of the transit hypothesis with a period equal to the full duration of the light curve and $n=0$, but the adopted model is not suited to aperiodic variations. The special case where the period is equal to the total duration of the light curve corresponds to a single transit, which is considered to be part of \mathcal{H}_T .

2.1.1.2 Bayesian analysis

Just like any signal detection problem, transit searching consists in hypothesis testing. Bayesian analysis provides a framework for the a posteriori incorporation of new or additional information in a detection process. Given the rate at which our knowledge of the characteristics of extra-solar planets is increasing, such an approach has distinct advantages. Already, one can directly incorporate into the detection the fact that transits are rare events due to the low alignment probability. One can also incorporate the fact that this alignment probability is lower for the more distant, i.e. longer period planets, by setting the a priori probability distribution for the period, or period *prior*, to be lower for longer periods.

An excellent discussion of Bayesian detection and parameter estimation is given in GL92. However, given the fact that Bayesian methods are relatively rarely used in the present field, the process is described below in some detail.

The hypothesis under test is that the light curve $X = \{x_1, x_2, \dots, x_N\}$, contains (periodic), short and shallow dips against an otherwise constant background. This global transit hypothesis will be referred to as \mathcal{H}_T . It is generally relatively straightforward, under given assumptions about the noise characteristics of the data, to compute a *likelihood* for \mathcal{H}_T , that is a measure of the extent to which \mathcal{H}_T predicts X , $P(X|\mathcal{H}_T)$ (see Section 2.1.1.3). But the quantity of interest is $P(\mathcal{H}_T|X)$, the *posterior probability* for the model. The relationship between the two is governed by Bayes' theorem (Bayes 1764):

$$P(\mathcal{H}_T|X) = \frac{P(\mathcal{H}_T)}{P(X)} \times P(X|\mathcal{H}_T), \quad (2.1)$$

where $P(\mathcal{H}_T)$ encompasses a priori information about hypothesis and the underlying physical processes, measurement effects, etc., and is consequently known as a *prior*, and $P(X)$ is a normalisation constant.

In the case of transit searches, the prior for \mathcal{H}_T , $P(\mathcal{H}_T)$, represents any a priori information as to the probability that the light curve contains transits. In the range of planetary radii already probed by other methods, its value could be deduced from the frequency of planets observed to date, as well as from the alignment probability of the orbit with the line of sight (requiring some assumption about the distribution of orbital inclinations). This would require a relatively complex integration over several parameters, and the possibility that other variations in the light curve might dominate over any transits would still need to be accounted for. At first glance however, the alignment requirement alone suggests that $P(\mathcal{H}_T) \leq 0.1$.

$P(X)$, the normalising factor, is given by:

$$P(X) = \sum_i P(\mathcal{H}_i|X) \times P(\mathcal{H}_i). \quad (2.2)$$

where the \mathcal{H}_i represent each of the hypotheses under test. The need to evaluate $P(X)$ is circumvented by comparing the hypothesis under consideration, \mathcal{H}_T , to that for another, *null* hypothesis. In the present context, the null hypothesis consists of a constant light curve, and is referred to as \mathcal{H}_C . One computes the *odds ratio*, or ratio of the posterior probabilities for the two hypotheses:

$$\mathcal{O} = \frac{P(\mathcal{H}_T|X)}{P(\mathcal{H}_C|X)} = \frac{P(\mathcal{H}_T)}{P(\mathcal{H}_C)} \times \frac{P(X|\mathcal{H}_T)}{P(X|\mathcal{H}_C)}, \quad (2.3)$$

If \mathcal{O} is greater than 1, there is evidence for transits in the light curve. If, as in the example above, $P(\mathcal{H}_T) = 0.1$ (and assuming all light curves correspond to either \mathcal{H}_T or \mathcal{H}_C), then $P(\mathcal{H}_T)/P(\mathcal{H}_C) = 0.11$.

In practice, one does not compute a likelihood for \mathcal{H}_T as a whole, but rather for a single member \mathcal{M}_n of the class of models represented by \mathcal{H}_T , with a particular value of n , and for a particular set of parameters (period, duration, epoch, shape) represented by the $(n + 4)$ -element vector V . Likelihoods are computed for each value of V , and combined through a marginalisation process. Given a multi-dimensional array of likelihood values $P(X|\mathcal{M}_n, V)$, unwanted parameters are marginalised over. One marginalises over each parameter v_i in turn by multiplying the likelihood by a one-dimensional prior for that parameter and integrating over (or summing over all

sampled values) of v_i :

$$P(X|\mathcal{M}_n, V') = \sum_{v_i=v_i^{\min}}^{v_i^{\max}} \{P(v_i|\mathcal{M}_n) \times P(X|\mathcal{M}_n, V)\}, \quad (2.4)$$

where V' is the $(n+3)$ -element vector of all parameters contained in V except v_i . The choice of prior depends on pre-existing knowledge about each parameter. For example, in the case of transits the prior for period might be chosen to reflect the orbital period distribution observed to date. Specific forms of prior can also reflect common sense, for example a logarithmic prior in period or frequency ensures that the results are the same whether one works in period or frequency (see GL92 and references therein). The marginalisation is repeated until no parameters remain (in that respect n can be treated as a parameter), and the left hand side of Equation (2.4) is the global likelihood for \mathcal{H}_T . After computing the likelihood for \mathcal{H}_C , one then obtains the global odds ratio \mathcal{O} .

If $\mathcal{O} > 1$, i.e. there is evidence for the transit hypothesis, it is natural to ask which set of transit parameters best describes the data. A transit detector must provide some information about the location of the transits in the light curve to be useful, at least the period and epoch. The posterior probability distribution for each parameter v_i is obtained from the one-dimensional likelihood function, marginalised over all other parameters (including n), using Bayes' theorem:

$$P(v_i|X, \mathcal{H}_T) = \frac{P(v_i|\mathcal{H}_T)}{P(X)} \times P(X|\mathcal{H}_T, v_i). \quad (2.5)$$

Given that $P(X)$ is constant with respect to V , the best trial value of v_i is that which maximises $P(v_i|X, \mathcal{H}_T) \times P(X) = P(v_i) \times P(X|\mathcal{H}_T, v_i)$.

2.1.1.3 Likelihood calculation

The calculations are given below in sufficient detail to allow the reader to reproduce the algorithm, but some lengthy derivations which were taken from G99 have not been reproduced. The original GL method was implemented in parallel to the modified version to provide a benchmark for test purposes, but with the same simplifications as the modified method, that is no noise scale parameter. The large number of parameters in the original GL method is its very weakness: each of them implies a dilution of the likelihood function (whose peak is spread over more dimensions of parameter space). This leads to larger variances in the estimated parameters, and hence to smaller statistical efficiencies.

Each data value x_i , corresponding to time t_i , is decomposed as

$$x_i = r_i + e_i, \quad (2.6)$$

where r_i is the value predicted by the model at time t_i and e_i represents any variations in the data not accounted for by the model. In the present context e_i will contain, apart from true variations not accounted for by the model, a photon shot noise contribution, well approximated by a Gaussian distribution for high photon counts, plus other instrumental and astrophysical noise of unknown distribution. According to the Central Limit Theorem, the most conservative assumption for the distribution of e_i is a Gaussian.

The treatment of errors has been simplified relative to G99: we assume the noise standard deviation for each data point x_i has a known value σ_i . This assumption is justified in the context of transit searches where the noise characteristics should be well determined from the large numbers of simultaneous high-precision light curves, and removes the need for the rather confusing noise scale parameter b of Gregory (1999).

The likelihood for a given value of n (i.e. model \mathcal{M}_n) with a given set of $n + 4$ parameters (p, d, e, R) is then given by a product of Gaussian probability distributions (assuming the data points are independent):

$$P(X|\mathcal{M}_n, p, d, e, R) = \prod_{i=1}^N \left\{ \frac{\sigma_i^{-1}}{\sqrt{2\pi}} \times \exp \left[-\frac{(x_i - r_i)^2}{2\sigma_i^2} \right] \right\}, \quad (2.7)$$

where N is the total number of data points.

Before calculating this likelihood one must determine in which bin j of the model a given data point falls:

$$j(t_i) = \begin{cases} t_i^{\text{mod}} & : \text{ if } 0 < t_i^{\text{mod}} \leq n \\ 0 & : \text{ otherwise} \end{cases}, \quad (2.8)$$

where

$$t_i^{\text{mod}} = \text{int} \left(\frac{(t_i + p - e) \bmod p}{d/n} + 1 \right), \quad (2.9)$$

n is the number of bins per transit, $\text{int}(y)$ is the nearest integer lower than or equal to y and $a \bmod b$ is the remainder of a divided by b .

As a small aside, it is interesting to compare this expression with Equation (6) of Gregory (1999), which corresponds to the original set of models with m equal duration bins:

$$j(t_i) = \text{int} \{ 1 + m ((\omega t + \phi) \bmod 2\pi) / 2\pi \}, \quad (2.10)$$

where the angular frequency $\omega = 2\pi/p$ and ϕ is a phase parameter representing “the position of the first bin relative to the start of the data” and running from 0 to 2π . The definition and range of this parameter are inconsistent, unless one bin can somehow be identified as the first, which is not the case given the unknown shape of the model. As Gregory (1999) employed a uniform prior for ϕ , this inconsistency has no effect on the detection, and the period and shape determinations, which were the quantities of interest in that paper. However it explains some of the results presented in Section 2.1.2.

The likelihood can now be expressed in terms of the $n + 1$ bins of the model:

$$P(X|\mathcal{M}_n, p, d, e, R) = \prod_{j=0}^n \left[(2\pi)^{-(n_j/2)} \times \left(\prod_{i=1}^{n_j} \frac{1}{\sigma_i} \right) \times \exp\left(-\frac{\alpha_j}{2}\right) \right], \quad (2.11)$$

where n_j is the number of data points in bin j , and

$$\alpha_j = \sum_{i=1}^{n_j} \frac{(x_i - r_j)^2}{\sigma_i^2}, \quad (2.12)$$

r_j being the model value in bin j . For all purposes except the determination of the light curve shape inside the transit, the individual r_j 's do not matter. It is therefore desirable to marginalise over the r_j 's, that is to compute a combined likelihood for all possible values of the r_j 's within a range set a priori. According to Bayes' theorem, this is done by multiplying the likelihood by a prior and integrating over the range of r_j :

$$P(X|\mathcal{M}_n, p, d, e) = \prod_{j=0}^n \left[(2\pi)^{-(n_j/2)} \times \left(\prod_{i=1}^{n_j} \frac{1}{\sigma_i} \right) \times \mathcal{R}_j \right], \quad (2.13)$$

where

$$\mathcal{R}_j = \int_{r_{\min}}^{r_{\max}} dr_j P(r_j|\mathcal{M}_n) \exp\left(-\frac{\alpha_j}{2}\right), \quad (2.14)$$

r_{\min} and r_{\max} being the minimum and maximum value of the r_j 's, respectively. The distinct advantage of step-function models is that, as shown by Gregory, it is possible to perform this marginalisation analytically. Using a uniform prior for the r_j 's: $P(r_j|\mathcal{M}_n) = (\Delta_r)^{-1}$ where $\Delta_r = r_{\max} - r_{\min}$, and following the derivation of G99, we

obtain

$$P(X|\mathcal{M}_n, p, d, e) = \left(\frac{1}{\sqrt{2\pi}}\right)^N \left(\frac{1}{\Delta r}\right)^{(n+1)} \left(\frac{\pi}{2}\right)^{\binom{n+1}{2}} \left(\prod_{i=1}^N \frac{1}{\sigma_i}\right) \exp\left(-\sum_{j=0}^n \frac{\chi_{W_j}^2}{2}\right) \times \left\{ \prod_{j=0}^n W_j^{1/2} [\operatorname{erfc}(y_{j\min}) - \operatorname{erfc}(y_{j\max})] \right\}, \quad (2.15)$$

where the quantities W_j , $\chi_{W_j}^2$, $y_{j\min}$ and $y_{j\max}$ are taken directly from Equations (11) to (16) in G99, and $\operatorname{erfc}(y)$ is the complementary error function. In fact, it is shown in Section 2.2.2.1 that the bin levels are not independent parameters, and are fully determined by the data. This fact is exploited by the improved method that was later derived from the one presented in this Section.

2.1.1.4 Choice of priors

Following G99, we use a Jeffreys prior for the period:

$$P(p|\mathcal{M}_n) = \frac{1}{p \ln(p_{\max}/p_{\min})}, \quad (2.16)$$

where p_{\min} and p_{\max} are the limits of the period-space explored and $\ln(p_{\max}/p_{\min})$ is a normalisation constant to ensure that $\int_{p_{\min}}^{p_{\max}} P(p|\mathcal{M}_n) dp = 1$.

As pointed out by G99 (see references therein), this prior arises naturally from considerations of invariance with respect to changes in time scale and ensures that an investigator working in terms of p with this prior would obtain the same results as an investigator working in terms of frequency ν with a $1/\nu$ prior. In the present context it also reflects (though in a qualitative rather than quantitative fashion) the lower transit probability for longer periods that arises from geometric alignment considerations.

As already mentioned in Section 2.1.1.3, we use a flat prior for the r_j 's, introducing only a normalising factor $1/\Delta r$. Similar priors are also used for epoch and duration.

Provided the results of the individual likelihood calculation (Equation 2.15) are stored, the period and duration priors could be changed at a later date, for example once the observed distributions for these parameters in the case of planetary transits are better known.

Finally, all values of n are considered equally likely a priori, bearing in mind that higher values of n are automatically affected by an implicit Occam's razor penalty factor, as discussed by GL92.

2.1.1.5 Weighting factor to compensate for uneven distribution into the bins

When the number of periods is low such that one bin might be represented four times while another only three times, or if there are gaps in the data which may not be evenly distributed over the bins, GL92 noted that some of their initial assumptions may fail, leading to the appearance of an erroneous trend in the posterior probability for the period.

In an appendix to GL92, a solution to this problem was proposed. A weighting factor s_j is applied to each bin:

$$s_j = \left(\frac{n_j m}{N} \right)^{-n_j}. \quad (2.17)$$

It is important to note that this factor was derived in the context of Poisson statistics.

Despite the low number of periods in our light curves, we found that no weighting factor was required in the benchmark algorithm that reproduced the GL method identically. However, it is clear that the problem is more acute in the modified algorithm. The 'out of transit' bin contains many more data points than the others, and therefore has a much larger effective weight. A weighting factor is required to compensate for this problem. The expression given above for s_j is only appropriate in the photon count context in which it was derived, not in the Gaussian noise case relevant here. A different weighting factor can be heuristically derived by considering Equation (2.12). The contribution of each model level to the likelihood is a χ^2 sum. The variance of a χ^2 distribution is equal to twice the number of degrees of freedom \mathcal{N} . In each bin there are n_j data points and n_{par} parameters to adjust (in the modified GL method $n_{\text{par}} = n + 4$). As $n_j \gg n_{\text{par}}$, $\mathcal{N} = n_j - n_{\text{par}} \simeq n_j$. Weighting each bin by a factor $1/n_j$ is therefore equivalent to weighting proportionally to the inverse variance. In practice this is achieved by maintaining the expressions for $\chi_{W_j}^2$, $y_{j\text{min}}$ and $y_{j\text{max}}$ given in Gregory (1999) in terms of x_j and σ_j , but replacing W_j by W_j/n_j .

This modification was implemented in our algorithm and found to give more robust results.

2.1.1.6 Minimising the computing time

For a given set of parameters, the calculation of the likelihood involves summing over each element in each bin. The time required to compute the likelihood for a given set of p , d , e therefore scales linearly with the number of points in the light curve. It also increases with the number of bins, but this is a slow increase. It does not depend on the individual parameter values.

The overall computing time also depends, of course, on how tightly the param-

eter space is sampled. It is necessary to minimise the number of trial values for each parameter without missing potentially localised likelihood maxima. Because of the relative sharpness of the peak in the posterior probability for the period, the period increment needs to be kept fairly small (typically once or twice the time step between data points). Attention was therefore concentrated on what increment was suitable in terms of epoch. The results are not significantly worsened by increasing the *posital* phase ϕ_{pos} ¹ increment from $1/p$ (i.e. shifting the model by 1 sampling time at each increment) to $d/2np$ (i.e. shifting the model by half the duration of an in-transit bin at each increment). Further increase leads to sharp steps in the posterior probability distribution (analogous to Shannon’s sampling theorem).

However, the computing time is inversely proportional to the increment, and a posital phase increment of $d/2np$ is still prohibitively expensive. In practice, steps in the posterior probability distribution that result from a larger increment can be effectively removed by dividing it by the equivalent distribution for an entirely flat light curve with the same duration, sampling and data gaps as the light curve. We call this dividing function the ‘window function’². We therefore use a posital phase increment of $d/2p$ and perform the division before analysing the results. As the window function only needs to be calculated once per period and duration, this is much faster than using a smaller increment (see Section 2.1.3).

Note that due to the use of this window function one should not strictly speaking use the word ‘posterior probability’ when talking about the output of the algorithm. Hereafter, we will refer to ‘modified posterior probability’ to mean ‘posterior probability distribution divided by the window function’. This also implies that the global odds ratios mentioned in Section 2.1.1.2 cannot be used to directly measure the ratio of the probabilities for a periodic model compared to a constant model. Instead, we use bootstrap simulations (see Section 2.1.3.1) to set a threshold value of the detection statistic above which a detection is accepted.

2.1.2 Comparison with the original GL method

In order to establish a reference point and to gain a preliminary estimate of the modified algorithm’s performance, some qualitative tests were run on both the original and the modified version.

For this purpose light curves containing transits and photon noise were generated with the parameters of the *Eddington* mission in mind. We describe below a reference light curve simulated with one particular set of parameter values. Each

¹ $\phi_{\text{pos}} = \phi/2\pi$.

²This also has the advantage of ironing out any residual effects of the uneven bin duration not removed by the weighting factor.

parameter was then varied in turn over a small but representative range.

All light curves have a sampling time of 15 minutes. The total light curve duration is 4 months. To simulate the transit signal from a planet in a 1 year orbit observed for 3 or 4 years without the computational expense such a duration would imply, a 1 month orbital period was used.

Given the presence of limb darkening in stellar photospheres, planetary transits are not perfectly ‘flat bottomed’ (nor are they, strictly speaking, truly grey). To simulate transits in a realistic way, we use the Universal Transit Modeller (UTM) software written by H. J. Deeg (Deeg et al. 2001). UTM can simulate light curves from any number of luminous or dark objects, including stars, planets, rings and moons. Circular orbits are assumed, and a linear limb darkening law is adopted for the stars. We used limb darkening coefficients from Van Hamme (1993). The dark objects – planets, rings and moons – are assumed to have zero albedo, i.e. no reflected light is included. This is justified because, in white light, the amount of light reflected by planets is expected to be small compared to any transit signal. Even a large close-in planet (0.05 AU) with the size and albedo of Jupiter ($\sim 30\%^3$) would reflect approximately 8×10^{-6} of the light of its parent star, while it would cause transits of $\sim 1\%$ (for a Sun-like star).

The reference light curve corresponds to a $1 R_{\text{Jup}}$ planet orbiting a $1 R_{\odot}$ star with $V = 10$, resulting in a transit depth of $\sim 14\sigma$. The chosen orbital distance of $a = 15.3 R_{\odot}$, for a 1 month period, results in 4 transits lasting ~ 15 hours each, the ingress and egress lasting approximately 3.3 hours each. The posital phase was set to 0.25.

The level of photon noise in the light curve was computed from the photon counts expected for a G2V star, based on the throughput and aperture of the *Eddington* baseline design as described in Favata & the *Eddington* Science Team (2000), i.e. a collecting area of 0.6 m^2 and a total system throughput of 70%. Such an instrument would detect $\simeq 50$ photons per second from a $V = 21.5$ G2V star. The photon noise for each point in the light curves is then simulated by independent draws from a Gaussian distribution, with a standard deviation equal to the square root of the expected photon count per integration for that star.

Both versions of the algorithm were run on the reference light curve described above and the modified posterior probabilities were plotted as a function of period and as a function of phase. The number of bins used was $m = 10$ in the case of the GL method, and $n = 4$ in the case of the modified method. In order to sample the transit as well with the GL method as with the modified method, a much higher value of m

³<http://nssdc.gsfc.nasa.gov/planetary/factsheet/jupiterfact.html>




Figure 2.3: Comparison of the GL and modified GL methods for the case of a Jovian planet transiting across a 10^{th} magnitude star, with a period of 30 days (a) and a posital phase of 0.25 (b). Solid line: modified GL method. Dashed line: GL method. Both methods successfully detect the period of the transits although the peak is sharper with the modified method. The GL method is unsuccessful in the phase domain (the GL phase results are folded over the 10 bins). Note that the probabilities are expressed in arbitrary units.

would need to be used, but this would be too computationally expensive. Instead the values of m and n were chosen such that the computing times were similar. The results obtained for this benchmark case are shown in Figure 2.3.

Each of the parameters (be they associated with the light curve or with the model) was then varied over a small range of representative values. These one-off tests on a small parameter space confirmed some expected trends.

- For a given light curve duration the detection is less precise for longer periods as the light curve contains fewer transits.
- As expected, the original GL method is not well suited to detecting the phase, as there is no natural way of labelling one particular bin as the first one. On the other hand the phase is very successfully recovered with the modified version.
- The larger the value of m (GL method), the sharper the detection. However, $m = 10$ appeared sufficient for our purposes.
- Increasing the value of n (modified method) does not necessarily improve the detection ability since one starts to fit the noise inside the transits, which is not periodic. When fitting Gaussian profiles it is standard to require a minimum of 2 bins per FWHM. The shape of the transit is not Gaussian but it is relatively simple, hence we multiplied by a safety factor of 2, leading to $n = 4$ in further calculations. However when dealing with a particular value of d it is advantageous to choose n so that d is a multiple of it to avoid introducing extra noise by splitting individual data points across bin boundaries.
- Although the modified method should in principle allow us to determine the duration of the transit, in practice this is not successful. The program may be fitting a much wider region than the transit itself. In the GL method, as there

are only 10 to 20 bins per period, with p of order several hundred sampling times or more, the bin in which the transit falls is much larger than the transit itself. We have seen that the loss of information this implies does not prevent the detection of the period by the GL method. The modified algorithm is likely to overestimate the transit duration because fitting a region larger than the transit does not significantly reduce the likelihood. For now the duration of the transit was simply marginalised over; once the presence of a transit is asserted and its period known, phase folding should allow a fairly quick determination of the shape and duration;

- For a given set of parameters, with $m = 10$ and $n = 4$, such that both algorithms have similar computing times, the detection peaks are much sharper with the modified version.

2.1.3 Performance evaluation in white Gaussian noise

2.1.3.1 Method

As mentioned in Section 2.1.1.6, the use of a window function to remove the effects of under-sampling in the phase domain, while minimising the computing time, rules out the possibility of direct computation of a global odds ratio, whose value could be used to determine whether or not a given light curve contains transits. Instead, to evaluate the performance of the algorithm, it was run on simulated light curves with similar noise characteristics, some containing transits and some not, and the results were compared.

This method was previously used in a similar context by Doyle et al. (2000). For each set of trial parameters the algorithm was run first on a set of one hundred simulated light curves containing only Gaussian noise and no transits. Subsequently it was run on another set of one hundred simulated light curves containing Jovian-type planetary transits with the characteristics described in Section 2.1.2, with the same level but different realisations of the photon noise, and with uniformly distributed random phases

For each simulation, the modified posterior probabilities were plotted versus period and the value of the maximum was noted. This maximum is our 'detection statistic', on the basis of which we wish to determine whether there is a transit or not. We then plot a histogram of the detection statistics measured for all the light curves with transits and one histogram for all the light curves with noise only. In other words, one histogram corresponds to the cases where the transit hypothesis is correct and one to the cases where the null hypothesis is correct. Ideally, the two histograms should be completely separated, with no overlap, and choosing a detection thresh-




Figure 2.4: Schematic diagram of the performance evaluation method. Solid line: detection statistic distribution for the light curves with transits. Dashed line: idem for the transit-less light curves. Vertical solid line: threshold. Hashed area: missed detections. Filled area: false alarms.

old located between the two histograms would guarantee a 100% detection rate and a 0% false alarm rate. In practice, for the cases of real interest, close to the detectability limit, the two histograms will overlap. A compromise has to be found by choosing a threshold which minimises a penalty factor designed to take into account both false alarm and missed detection rates. This is illustrated in Figure 2.4.

Depending on the circumstances, it may be more important to minimise the false alarm rate than the missed detection rate. This is the approach followed by Jenkins et al. (2002), on the basis that detections from space experiments are hard to follow-up from the ground. An alternative view is any real transit that is rejected is a loss of valuable scientific information. As long as the false alarm rate is kept to a manageable level, further analysis of the light curves will prune out the false events. We have opted here for an intermediate position, and our penalty factor is simply the sum of the missed detection rate N_{MD} and the false alarm rate N_{FA} :

$$F_{\text{penalty}} = N_{FA} + N_{MD}. \quad (2.18)$$

After marginalisation over the other parameters, the detection algorithm yields modified posterior probabilities as a function of period and as a function of phase. The simultaneous use of the two detection statistics S_{per} and S_{ph} (plotting 2-D rather than 1-D distributions) increases the discriminating power of the algorithm, (as long as the two distributions do not have secondary maxima in 2-D space). This is shown when comparing the false alarm and missed detection rates obtained from period and phase information separately and together. The threshold in the 2-D case takes the form of a line: $S_{\text{ph}} = a + b \times S_{\text{per}}$. Here the optimal values of a and b were found by trial and error, although standard discriminant analysis techniques could be used to determine them automatically.

2.1.3.2 A preliminary test case

In Defaÿ (2001), analysis performed on the basis of 200 bootstrap samples for the COROT observations of a star with $V = 13$ and an Earth-sized planet, containing 6 transits lasting 5 hours each, yielded a probability of true detection of $\simeq 0.3$. We performed the simulations described in Section 2.1.3.1 for a similar case: an Earth-sized planet orbiting a K5V type star with $V = 13$, a period of 30 days and a transit duration of 5 hours. The sampling time was 15 minutes. The noise was different from the COROT case, as we concentrate uniquely on the photon noise expected for *Eddington*.




Figure 2.5: Detection statistic distributions for an Earth-sized planet orbiting a $V=13$ star with a 30 day period (light curve duration 120 days). Solid line: light curves with transits. Dashed line: transit-less light curves. Vertical solid line: threshold. a) Detection statistic obtained from the modified posterior probability distributions as a function of period. b) Idem as a function of phase. Over 100 realisations there were no false alarms and no missed detections.

The results are shown in Figure 2.5 for period and phase separately. As the distributions for the noise only and transit light curves are completely separated, each parameter alone is sufficient to determine a threshold ensuring null false alarm and missed detection rates.

2.1.3.3 Finding the magnitude limit of *Eddington* for Earth-like planets

Given that the key scientific goal of *Eddington* in the field of planet-finding is the detection of habitable planets, the performance of the algorithm was extensively tested for habitable planets at (or close to) the noise limit of *Eddington*. The case of an Earth-sized planet orbiting a K dwarf in a habitable orbit was used as benchmark. The light curve was simulated for a system with the following parameters:

- K5V star ($R_{\star} = 0.8 R_{\odot}$) with a range of apparent V -band magnitudes $V = 14.0$, 14.5 and 15.0;

- Earth-sized planet ($R_p = R_\oplus$) with an orbital period of 4 months, orbiting the star at a distance of 0.64 A.U. (leading to a transit duration of ~ 10.5 hours);
- light curve duration of 16 months;
- sampling time 1 hour.

An example of a light curve is shown in Figure 2.6. The resulting transit event has a depth $\Delta F/F = 1.4 \times 10^{-4}$. For the *Eddington* baseline collecting area a star at $V = 14$ will result in a photon count of 1.8×10^8 per hour, so that the Poisson noise standard deviation will be 1.34×10^4 . The S/N of the transit event in each 1 hour bin will thus be 1.88. Following the same reasoning for the $V = 15$ case, the S/N of the transit event in a single one hour bin is 1.19. As there are 4 transits lasting 10 hours each in the light curves considered, the overall transit signal has a S/N of $\sqrt{40} \times 1.19 \simeq 7.5$.

With the results of the simulations, an example of which is shown in Figure 2.7, the analysis described in Section 2.1.3.1 was performed for all three magnitudes, confirming that the combined use of the two statistics improves the results. This is illustrated for the $V = 14.5$ case in Figures 2.8 & 2.9 (for this particular case 1000 rather than 100 runs were computed to improve precision).

As illustrated in Figure 2.10, a mean error rate (the mean of the false alarm and missed detection rates) $< 3\%$ can be achieved down to $V = 14.5$. This magnitude is therefore taken as the performance limit for the algorithm for an Earth-sized planet around a K5V-type star. However this analysis is not complete enough to allow a precise determination of the magnitude limit. First the noise treatment is incomplete, photon noise only being considered. Second, one would need more runs per simulation to compute meaningful errors on the false alarm and missed detection rates. Sets of 1000 runs, as was done for the limiting $V = 14.5$ case, should be computed for all cases.

The asymmetric shape of the distributions shown in Figures 2.5, 2.8 & 2.9 implies that, even though the thresholds are chosen to minimise false alarms and missed detections equally, the optimal threshold results in more false alarms than missed detections. This could easily be avoided, if needed, by replacing Equation (2.18) by:

$$F_{\text{penalty}} = A \times N_{\text{FA}} + N_{\text{MD}}. \quad (2.19)$$

where A is a factor greater than 1. Alternatively one could keep the penalty factor unchanged but set a strict requirement on the maximum acceptable false alarm rate.

As in any unbiased search for periodicity in a time-series, the inclusion of a larger range of periods in the search will lead to a higher chance of finding a spurious (noise-induced) period signal in the data. The simulations used here to assess the




Figure 2.6: Example light curve containing 4 transits of an Earth-sized planet orbiting a K5V star with $V=14.5$. a) Full light curve. b) Portion around a transit. c) The four transits co-added.




Figure 2.7: Example detection statistic distributions arising from the light curve shown in Figure 2.6 (arbitrary units). a) Period – real value 2912 hours, error -2 hours. b) Phase: real value 0.885, error 0.005.

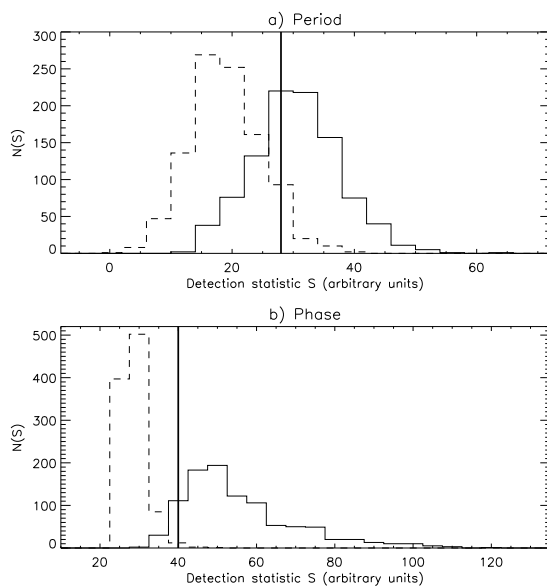


Figure 2.8: Detection statistic distributions for an Earth-sized planet orbiting a $V=14.5$ star with a 4 month period (light curve duration 16 months). Solid line: light curves with transits. Dashed line: transit-less light curves. Vertical solid line: threshold. a) Period: 190 false alarms and 185 missed detections over the 1000 realisations. b) Phase: 27 false alarms and 14 missed detections over the 1000 realisations.

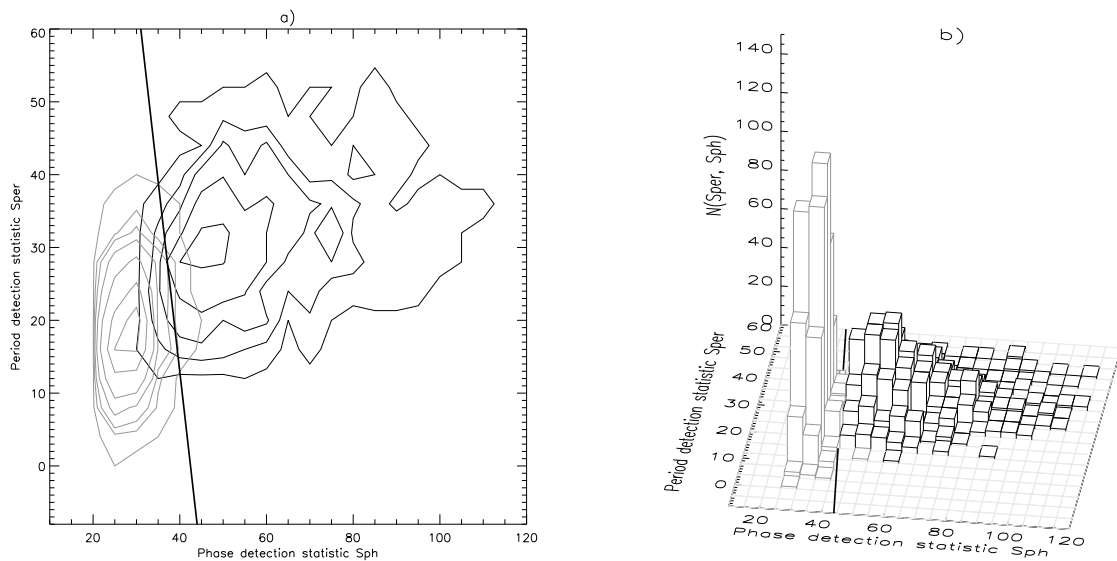


Figure 2.9: a) Contour plot and b) 3-D representation of the two-dimensional (period & phase) detection statistic distribution for an Earth-sized planet orbiting a $V=14.5$ star with a 4 month period (light curve duration 16 months). Black: lightcurves with transits. Grey: transit-less light curves. Solid line: threshold: $S_{ph} = 42.47 - 1.191 \times S_{per}$, yielding 29 false alarms and 9 missed detections over the 1000 realisations.




Figure 2.10: Evolution of the algorithm's performance (in terms of fractional error rates) with magnitude (period 4 months, light curve duration 16 months) a) using the period statistic only, b) using the phase statistic only and c) combining the two statistics. Dotted line: false alarm rate. Dashed line: missed detection rate. Solid line: mean error rate.

algorithm's performance are based on a search through a relatively small range of periods. In practice, lacking any a priori knowledge of the possible periodicity of planetary orbits around the star being observed, one will want to test a large range of periods, ranging from a few days (the physical limit of the period of planetary orbits) all the way to the duration of the data set (searching for individual transit events).

2.1.3.4 Data gaps

Any realistic data set will suffer from gaps in the data. While the orbits of both *Eddington* and *Kepler* have been chosen to minimise gaps, 100% availability is not realistic, and gaps will be present due to e.g. telemetry dropouts, spacecraft momentum dumping manoeuvres, showers of solar protons during large solar flares, etc... For this reason, any realistic algorithm must be robust against the presence of gaps in the data, showing graceful degradation as a function of the fraction of data missing from the time series.

We have therefore tested the algorithm discussed here using simulated light curves with 5%, 10% and 20% data gaps, randomly distributed in the data, i.e. 5% of the points in the time series are selected randomly with a uniform distribution and removed from the light curve. The gaps will probably not be randomly distributed in reality, but as the typical gap duration is expected to be of order 1 or 2 hours, simulated random gaps can already be used to test the algorithm's robustness. For reasons of computing time, to avoid having to recalculate the 'window function' at each run, the distribution of the data gaps is the same for all runs of a simulation. As the gaps are chosen one by one there are rarely gaps of more than two consecutive time steps, i.e. 2 hours. Note that e.g. the *Eddington* mission is designed to produce light curves with a duty cycle $\geq 90\%$, so that the case with 20% data gaps represents a worst case scenario.

The results are shown in Figure 2.11. There is visibly very little degradation up to 20% data gaps. When using S_{ph} alone or the two statistics combined there is no perceptible difference. We can therefore say this algorithm is robust at least for data gaps of the type likely to occur due to e.g. telemetry dropouts, which last only a few hours. One would also expect the algorithm to perform well in the presence of longer gaps: the effect of gaps is to render the number of samples per bin uneven, and this is already the case for this particular method with no gaps at all.

Note that the impact of gaps was tested using a configuration closer to the detectability limit ($V = 14.5$) with the second algorithm, the box-shaped transit finder (see Section 2.2.4).




Figure 2.11: Evolution of the algorithm's performance with data gaps (period 4 months, light curve duration 16 months, $V=14.0$) a) using the period statistic only, b) using the phase statistic only and c) combining the two statistics. Dotted line: false alarm rate. Dashed line: missed detection rate. Solid line: mean error rate.

2.1.3.5 Number of transits in the light curve

The planetary transits detection phase of the *Eddington* mission is planned to last 3 years with a single pointing for the entire duration of that phase. There will therefore be three or four transits in the light curve for a typical habitable planet. However, other missions such as COROT are planned with shorter (5 months) pointings and it is of interest for this type of mission to study the degradation of the algorithm's performance as the number of transits in the light curve reduces. If the algorithm performs well with 2 or less transits, in the context of *Eddington* it may also allow the detection of 'cool Jupiters', i.e. Jupiter-sized planets with orbits more similar to those of the gaseous giants in our solar system. This would be of relevance to the question of how typical our solar system is.

Sets of 100 runs with the characteristics specified in Section 2.1.3.3 for a star with $V = 14.5$ were computed for light curve durations of 4, 8, 12, 16 and 20 months, containing between 1 and 5 transits. The results are shown in Figure 2.12. The degradation only becomes significant when less than three transits are present. However, even mono-transits could be detectable for larger planets at that magnitude.

Defaÿ (2001) compared a matched filter approach with a Bayesian method based on the decomposition of the light curve into its Fourier coefficients. Their results suggest that the performance degradation in the low number of transits case is faster for the Bayesian method than for the matched filter. This is because the matched filter makes use of assumptions about the transit shape. It is also shown that when the Bayesian method fails to detect a transit, it can still reconstruct it if the detection is performed using a matched filter. Our algorithm has not been directly compared




Figure 2.12: Evolution of the algorithm's performance with the number of transits in the light curve (period 4 months, $V=14.5$) a) using the period statistic only, b) using the phase statistic only and c) combining the two statistics. Dotted line: false alarm rate. Dashed line: missed detection rate. Solid line: mean error rate.

to a matched filter. Its very design is based on the search for a short periodic signal in an otherwise flat light curve, which is itself an assumption about the shape of the signal. The matched filter makes use of more detailed knowledge of the transit shape and is therefore likely to perform better in the low transit number limit. However our algorithm with $n = 1$ may provide already a very good approximation to the relatively simple shape that is a transit, and may therefore perform nearly as well.

2.1.3.6 Differences in the two statistics

The two a posteriori probabilities show a different behaviour. In general the phase statistic is far more discriminatory than the period statistic. This is illustrated by contour plots of the likelihood as a function of trial period and posital phase, as shown in Figure 2.13. The period statistic's lesser effectiveness may be explained in the following way. If the phase is wrong, even if the period is right, it is likely none of the transits will be matched. If the phase is right, whatever the period, at least the first transit will be matched by the model. First we consider the likelihood distribution as a function of phase, normalised over all periods. For an incorrect phase the contribution from the correct period is nil as all transits are generally missed, but for the correct phase all trial periods produce a non-negligible contribution (the correct period of course contributing most). The likelihood distribution as a function of phase is therefore sharply peaked. Then we consider the likelihood distribution as a function of period, normalised over all phases. The contribution from the correct phase is non-negligible whatever the period. When the period is correct, the contribution from




Figure 2.13: Example contour plot of the likelihood as a function of trial period and posital phase for a simulated 16 month light curve containing photon noise as for a $V = 14.5$ star with 1 hour sampling and transits of a $1 R_{\oplus}$ planet with a period of ~ 4 months (121.33 days). Left: full parameter space explored. Right: zooming in on the true posital phase, which is 0.209. The diagonal lines correspond to pairs of trial periods and phases where at least one of the four transits is matched by the model.

the correct phase is washed out by the contributions from all the incorrect phases. The likelihood distribution as a function of period is therefore less sharply peaked. Additionally, the range of periods explored was, for computational reasons, kept relatively small, so that only a small region of parameter space is covered in that direction.

However, the combined use of the two parameters is more successful than the phase statistic alone. The reason for this is illustrated in Figure 2.9: in 2-D space the two distributions are aligned on a diagonal, such that no single value cutoff is optimal in either direction, compared to the line shown. The global odds ratio described in Section 2.1.1.2 could be used for such a purpose. We have noted in Sect. 2.1.1.6 that the global odds ratio for a given light curve cannot be used as an absolute statistic in the context of the present method. It can however be used as relative detection statistic, like S_{per} & S_{ph} , combined with bootstrap simulations.

The algorithm described in Section 2.2, which was derived using the lessons from the present one, directly combines phase and period information into a single statistic.

2.1.3.7 Discussion

Efficient data processing is one of the challenges for the upcoming generation of large scale searches for exo-planets through photometric transits. While radial velocity searches concentrate on limited number of stars, transit searches will investigate simultaneously large numbers of stars, and produce large amounts of data (photometric light curves) for each of them. A computationally efficient and robust algorithm for the processing of these data sets is necessary to make transit searches feasible. It is likely that the photometric time series which represent the observational product of the transit searches will be analysed in different stages, using more than a single approach. In particular, a first level of processing (after instrumental effects have been removed) should concentrate on singling out high-probability transit candidates, while efficiently pruning out the large number (more than 90%, even if all stars have planets, due to the low probability of transit events) of light curves in which no transits are present. In this first stage of analysis the ability to efficiently screen real transits in the data – even at the price of a moderate number of false alarms – is a key requirement for the algorithm. The candidate light curves in which a transit is suspected will then later be subject to a more detailed processing, which can then afford to be computationally less efficient (given it has to operate on a much smaller amount of data).

The modified GL method is able to detect transit events at the limit of the photon noise present in the light curve. It shows a graceful degradation of its performance as a function of different parameters of interest, e.g. the noise level in the data, as well as the presence of data gaps and the number of transits actually observed. Its strong sensitivity to the phase of periodic transits supplies significant additional information to be then used by further steps of processing for e.g. the reconstruction of the transit parameters. Thus, while little used in astronomy, Bayesian algorithms appear to be a powerful tool in the processing of transit data.

However, given the robustness and computational efficiency requirements, simplicity should be the guiding factor for subsequent work on transit detection. Simplifying assumptions will be tried one by one, and those that do not degrade the performance while improving the computing time will be incorporated in the algorithm for future use.

2.1.4 Tests with solar micro-variability

The performance of the modified GL method has been evaluated for simulated data from *Eddington*-like missions, containing simulated transits and photon noise. However, the influence of stellar variability induced by activity is likely to be the main

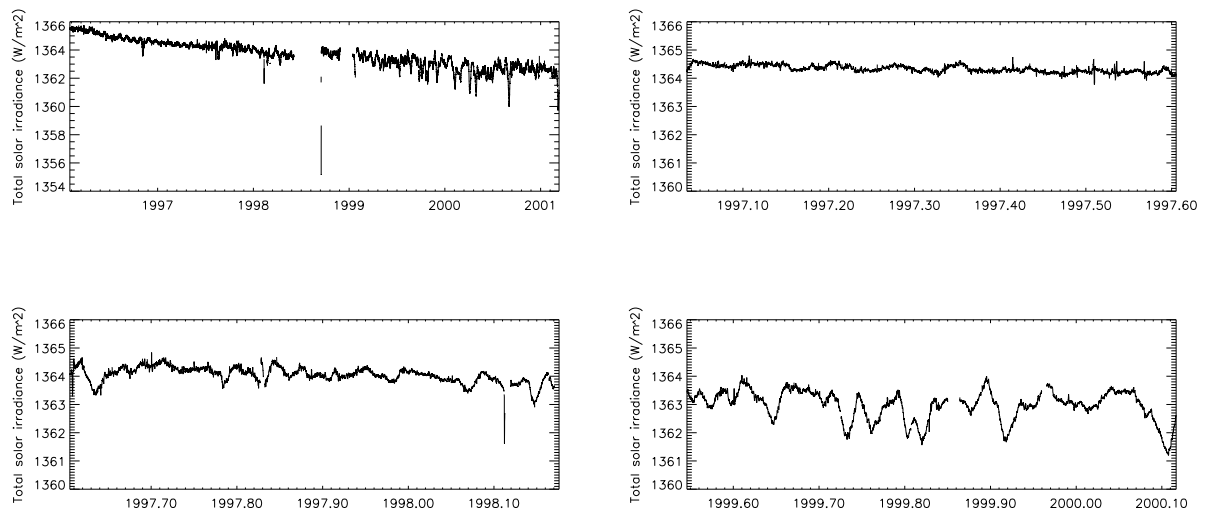


Figure 2.14: PMO6 light curves. Top left: full light curve, January 1996 to September 2001 (no correction for instrumental degradation). Top right: section used as low activity sample. Bottom left: *idem* for medium activity. Bottom right: *idem* for high activity.

limitation to *Eddington's* ability to detect planetary transits. The present section summarises the results of preliminary tests carried out by incorporating in the simulated light curves observed variations in the total irradiance of the Sun, as measured by the PMO6 radiometer, a part of the VIRGO experiment on the SoHO satellite, at solar activity minimum and maximum and at an intermediate phase.

2.1.4.1 Light curves

The transit light curves were simulated using UTM (see Section 2.1.2), with the parameters listed in Table 2.1. The full PMO6 light curve is shown in the top panel of Figure 2.14. This dataset is discussed in more detail in Chapter 3. Three 6 month long segments were selected at low, medium and high activity (bottom 3 panels of Figure 2.14). They were chosen to illustrate a particular feature of the Sun's variability, for example sun-spot signatures or modulation on the time scale of the solar rotational period. Care was taken also to avoid very long data gaps, although there are frequent short gaps (every few days, lasting a few tens of minutes to a few hours), and in the medium and high activity samples, some gaps last a few days. (N.B. This leads in some cases to the absence of one entire transit from the light curve, a situation which was not explored in Section 2.1.3, but with which the algorithm seems to cope reasonably well.) In each set of tests performed, the relevant segment of the PMO6 light curve was rebinned to 1 hour bins, normalised to a mean value of 1,




Figure 2.15: Example light curve with medium activity level, $R_p = 2 R_{\oplus}$. a) Full light curve. b) Portion around first transit.

repeated in order to cover the full duration of the simulated transit light curve and the product of stellar and transit light curves was taken. Care was taken to ensure that the 6 month repetition period of the PMO6 light curve was not too close to the transit period (4 months). Photon noise was then added as in the previous simulations according to the photon counts expected from the *Eddington* baseline design.

An example of the light curves produced is shown in Figure 2.15. Although individual transits are (except for the Earth-sized case) easily visible when zooming in on a portion of the curve, the most visible signal in the full curve is activity.

2.1.4.2 Results

The algorithm's performance with these light curves was tested using the method described in Section 2.1.3.1. However, only 10 realisations of the noise were computed in each case, these tests being intended as exploratory rather than systematic. As a result, the threshold analysis described in Section 2.1.3.1 would have been meaningless. Instead, results were recorded in terms of the number, if any, of detections which were not in fact transits but activity-induced features, i.e. the number of times the highest peak in the detection statistic distribution did not correspond to the tran-

Table 2.1: Parameters used for the simulated light curves with solar variability. Transit durations of 5 & 15 hours were also tested in the low activity case with a $2 R_{\oplus}$ planet.

Star:		Planet:	
spectral type	K5V	radius	10, 8, 6, 4, 2 & $1 R_{\oplus}$
radius	$0.74 R_{\odot}$	orbital period	4 months
activity level	low, med, high	orbital distance	0.43 A.U.
V-band magnitude	14.5	transit duration	10 hours

sit's period or phase. For each such case, the corresponding transit-less light curve (with identical noise and stellar variability realisations, but no transits) was used to check that the spurious detection was due to activity and not photon noise.

Table 2.2 gives an overview of the results in terms of the number of spurious detections for each set of 10 light curves with transits. We ascertain whether the detected peak corresponds to the correct transit period or phase by comparing with the distributions generated from transit-less light curves and by requiring that the error in the detected value be less than 2 hours (period) or 0.006 (posital phase).

As can be seen in table 2.2, the performance starts to degrade at 1, 2 and $2 R_{\oplus}$ for low, medium and high activity levels respectively. Noticeably, the breakdown is sudden, because it occurs when the transit induced peaks, whose height depends on the transit depth, i.e. on the planet radius, become smaller than the activity-induced peaks (which are of constant height and shape, as the same PMO6 sequence was always used for a given activity level). Although $V = 14.5$ was found to be the limiting magnitude for the reliable detection of an Earth-sized planet in the absence of activity (see Section 2.1.3), variability rather than photon noise is the limiting factor of the algorithm's performance, even in the low activity case. The results obtained here are therefore unlikely to improve much with increasing brightness. As the Sun is also a relatively quiet star for its type, and as other types of stars are likely to be generally more active, the problems outlined by these results are likely to have a serious impact on *Eddington's* performance and need to be addressed.

2.1.4.3 Example cases

It is helpful to highlight certain characteristics of the detection statistic distributions in a few representative cases in order to suggest ways to address the activity problem.

Figure 2.16 shows the distributions obtained for a planet twice the size of the Earth at the three activity levels. The middle row corresponds to the light curve illustrated in Figure 2.15. As activity increases, spurious activity-induced features appear in the distributions, and at high activity the peak corresponding to the actual transit is

R_p	low		medium		high	
$10 R_{\oplus}$	0	0	0	0	0	0
$8 R_{\oplus}$	0	0	0	0	0	0
$6 R_{\oplus}$	0	0	0	0	0	0
$4 R_{\oplus}$	0	0	0	0	0	0
$2 R_{\oplus}$	0	0	0	10	10	10
$1 R_{\oplus}$	10	9	10	8	10	10

Table 2.2: Number of cases where the highest peak in the posterior probability distribution is spurious (activity rather than transit induced) over each set of 10 noise realisations as a function of activity level (columns) and planet radius (rows). The first and second number in each cell relate to the period and phase detection statistic respectively.

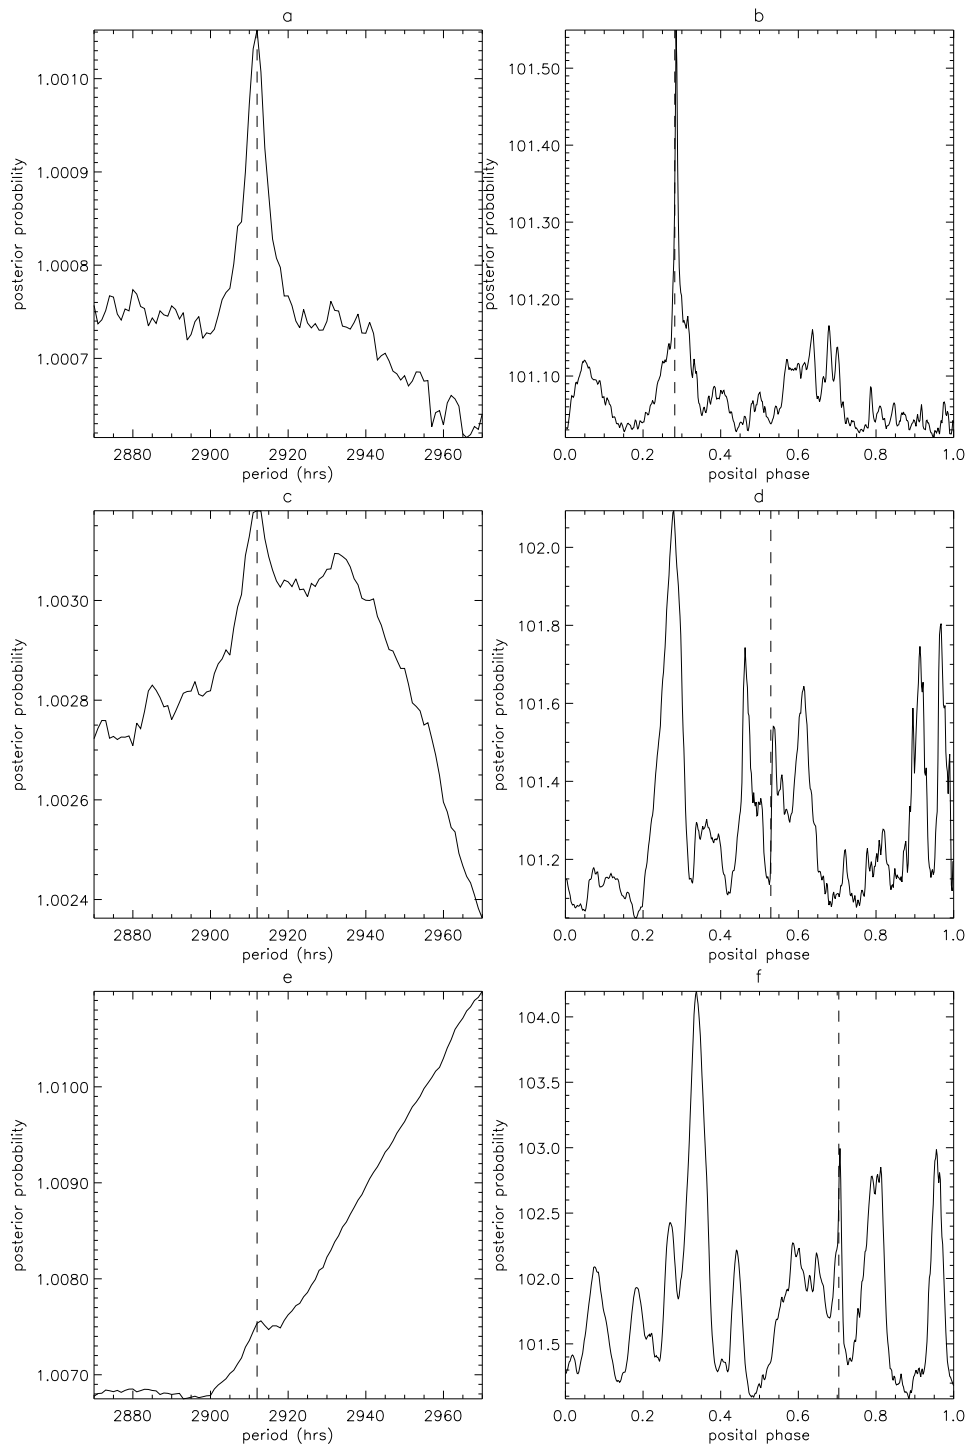


Figure 2.16: Example detection statistic distributions for period (left column) and phase (right column) at low (top row), medium (middle row) and high (bottom row) activity levels, with $R_p = 2R_\oplus$. In each panel the dashed line indicates the true value of the period or epoch.

no longer detected. The activity-induced features are particularly noticeable in the phase distributions, each peak corresponding to the first of the model light curve's transits being phase-folded onto (say) a sunspot transit. The period of repetition of the PMO6 sequences is well outside the range of periods tested, which avoids the detection of a peak at that period, but there is a discernible trend, which reduces the sensitivity of the period statistic.

Importantly, at the same magnitude, a similar planet was easily detected in the absence of variability. This implies that photon noise is not the limiting factor, hence that the same planet would not be detected around a brighter but active star.

2.1.4.4 Implications

The first step towards understanding the significance of the results presented here is to assess how typical the Sun's variability is, relative to other Sun-like stars and to stars of other spectral types and ages. Very little information on stellar activity on this kind of amplitude and timescale is available at present but there are a number of exploitable datasets either existing or expected within the next few years, which can be used to construct and calibrate a simple model of stellar micro-variability applicable to the variety of planetary transit search stars. A number of pre-processing techniques may be used to remove as much of the micro-variability signal as possible, before applying the modified GL method, or other methods designed primarily for light curves with white Gaussian noise. These topics are discussed in Chapters 3 and 4 respectively.

Given a more qualitative understanding of how micro-variability varies from star to star, and how it will impact transit detection, several aspects of the design of missions such as COROT, *Eddington* & *Kepler* may be affected. These include the choice of target field(s), which should be optimised to contain as many as possible of the least variable stars, or those whose variability can easily be filtered out. The observing strategy, for example the sampling time, may also be affected. Finally, it may be important to make use of additional information besides single bandpass photometry, for example to use the colour signature of the detected events to assess their planetary or stellar origin.

2.2 A stripped-down box-shaped transit finder

In Section 2.1, a dedicated Bayesian transit search algorithm was derived, based on the more general GL method for period finding. Here we develop this algorithm further and attempt to reconcile the apparent diversity of the extant transit algorithms.

Starting afresh from the original GL prescription, appropriate sequential simplifications can be made. We demonstrate that the levels of the step-function bins – which define the shape of the detected event – are not free parameters, their optimal values being fully defined by the data. The use of Bayesian priors can be dropped, given the lack of information currently available on the appropriate form for these priors. Finally, for detection purposes, the model can be simplified to an unequal bin duration square wave with only one out-of-transit and one in-transit value, where the out-of-transit section lasts much longer than the in-transit section.

After a brief aside on the close links between different families of detection methods in white Gaussian noise (Section 2.2.1), the algorithm itself is derived in Section 2.2.2, and its implementation is presented in Section 2.2.3. Its performance in terms of both detection capability and computational requirements is compared to that of the GL method in Section 2.2.4, and the results are discussed in Section 2.2.5.

2.2.1 Likelihood maximisation in Gaussian noise

Transit searches are generally performed by comparing light curves to a family of models with a common set of parameters, differing from each other according to the different values used for these parameters. A variety of methods exists to identify the best set of parameters. The most commonly used, in astronomy, is probably the matched filter, shown by Kay (1998) to be the optimal detector in the presence of white Gaussian noise. In the present section we show how the matched filter is derived from likelihood maximisation, and its equivalence to χ^2 minimisation and to a simple cross-correlation method. This is by no means a new result, but it helps to clarify the very close links that exist between the variety of transit search methods published in the literature, which are based on all of the above approaches, and to which the box-fitting algorithm will later be compared.

If the noise in each data point is assumed to be independently drawn from a Gaussian distribution (an assumption also valid for Poisson noise in the limit of large numbers of photons), the likelihood (or probability that the observed data is the result of adding noise to the model) can be written as the product of independent Gaussian probability distribution functions:

$$\mathcal{L} = \prod_{i=1}^N \left\{ \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp \left[-\frac{(x_i - r_i)^2}{2\sigma_i^2} \right] \right\}, \quad (2.20)$$

where x_i is the flux (or count, or magnitude) value at time t_i and r_i is the corresponding model value, N is the total number of data points and σ_i the expected error on

x_i . Equation (2.20) can be rewritten as

$$\mathcal{L} = \left(\frac{1}{2\pi}\right)^{(N/2)} \times \prod_{i=1}^N \left(\frac{1}{\sigma_i}\right) \times \exp\left(-\frac{\chi^2}{2}\right), \quad (2.21)$$

where

$$\chi^2 = \sum_{i=1}^N \left[\frac{(x_i - r_i)^2}{\sigma_i^2} \right], \quad (2.22)$$

so that likelihood maximisation, in the case of Gaussian noise, is equivalent to χ^2 minimisation, since the noise properties σ_i are assumed to be known, i.e. fixed.

Expanding Equation (2.22) gives

$$\chi^2 = \sum_{i=1}^N \left[\frac{x_i^2}{\sigma_i^2} \right] + \sum_{i=1}^N \left[\frac{r_i^2}{\sigma_i^2} \right] - 2 \sum_{i=1}^N \left[\frac{(x_i \times r_i)}{\sigma_i^2} \right]. \quad (2.23)$$

The first term is the error-weighted sum of all the data points, and is constant whatever the model. The second term is the equivalent sum for the model. In the case of shallow, short duration transits and provided the errors are relatively constant over the timespan of the observations, this term can also be considered constant. In such circumstances likelihood maximisation, or χ^2 minimisation, is therefore equivalent to maximising the third term, which is a zero-offset cross-correlation between the model and the data, i.e. a generalised matched filter:

$$\mathcal{MF} = \sum_{i=1}^N \left[\frac{(x_i \times r_i)}{\sigma_i^2} \right]. \quad (2.24)$$

In the case of transit searches, one of the parameters to be adjusted is the phase of the transit(s), i.e. the time at which the first model transit starts. Different phases are tested by simply introducing an offset between the model and the data, leading to formulation of the problem as a general cross-correlation:

$$\mathcal{CC} = \max \left\{ \sum_{i=1}^N \left[\frac{(x_i \times r_{i+n})}{\sigma_i^2} \right] \right\}_{n=1}^N. \quad (2.25)$$

where r is now defined for a unique reference epoch.

2.2.2 Simplification of the algorithm

2.2.2.1 Optimum χ^2 calculation for a generalised step-function model

We consider here a general periodic step-function model of the type used in the original GL method, characterised by the following parameters: number of bins m , period p , and epoch e (time elapsed between the start of the 1st bin and the start of the light curve) and bin levels $R = \{r_1, r_2, \dots, r_m\}$. By directly maximising the likelihood, or in this case minimising χ^2 , for such a model, it is straightforward to show that whatever the number and relative duration of the bins, the optimal values for the bin levels can be determined directly from the data given m , p and e . If we refer to the contribution from bin j to the overall χ^2 as χ_j^2 , and define J as the ensemble of indices falling into bin j , we have

$$\chi_j^2 = \sum_{i \in J} \left[\frac{(x_i - r_j)^2}{\sigma_i^2} \right]. \quad (2.26)$$

The value \tilde{r}_j of the model level r_j that minimises χ_j^2 is then simply given by the standard inverse variance-weighted mean of the data inside bin j , since by setting $\partial \chi_j^2 / \partial r_j$ to zero we have

$$\frac{\partial \chi_j^2}{\partial r_j} = 2 \sum_{i \in J} \left(\frac{x_i - r_j}{\sigma_i^2} \right) = 0, \quad (2.27)$$

hence

$$\tilde{r}_j = \bar{x}_j = \left[\sum_{i \in J} \sigma_i^{-2} \right]^{-1} \sum_{i \in J} x_i \sigma_i^{-2}. \quad (2.28)$$

Substituting into Equation (2.26), χ_j^2 now becomes

$$\tilde{\chi}_j^2 = \sum_{i \in J} \left[\frac{(x_i - \bar{x}_j)^2}{\sigma_i^2} \right], \quad (2.29)$$

where $\tilde{\chi}_j^2$ denotes the minimised value of χ_j^2 for a given period, epoch and number of bins. The contribution from each of the m bins can be simplified by expanding Equation (2.29):

$$\tilde{\chi}_j^2 = \sum_{i \in J} \left[\frac{x_i^2 - 2x_i \bar{x}_j + \bar{x}_j^2}{\sigma_i^2} \right]; \quad (2.30)$$

$$\tilde{\chi}_j^2 = \sum_{i \in J} \frac{x_i^2}{\sigma_i^2} - 2\bar{x}_j \sum_{i \in J} \frac{x_i}{\sigma_i^2} + \bar{x}_j^2 \sum_{i \in J} \frac{1}{\sigma_i^2}. \quad (2.31)$$




Figure 2.17: Schematic illustration of the type of model used in the box-fitting method. This is a single in-transit bin (cf. Figure 2.2, where the number of in-transit bins is 4) of duration d and one out-of-transit bin of duration $p - d$ where p is the period. The epoch e is defined as the time elapsed between the start of the data and the start of the next transit.

From Equation (2.28) we have

$$\sum_{i \in J} \frac{x_i}{\sigma_i^2} = \bar{x}_j \sum_{i \in J} \frac{1}{\sigma_i^2}, \quad (2.32)$$

so that

$$\tilde{\chi}_j^2 = \sum_{i \in J} \frac{x_i^2}{\sigma_i^2} - \bar{x}_j^2 \sum_{i \in J} \frac{1}{\sigma_i^2}. \quad (2.33)$$

The overall minimised χ^2 over all bins is thus

$$\tilde{\chi}^2 = \sum_{i=1}^N \frac{x_i^2}{\sigma_i^2} - \sum_{j=1}^m \left[\bar{x}_j^2 \sum_{i \in J} \frac{1}{\sigma_i^2} \right]. \quad (2.34)$$

The first term in Equation (2.34) is entirely independent of the model, and hence stays constant, so that only the second term needs to be calculated for each set of trial parameters.

2.2.2.2 Making use of the known characteristics of planetary transits

The Gregory-Loredo method makes no assumptions about the shape of the variations, and is fairly computationally intensive. However, when trying to detect planetary transits, most of the information is concentrated in a very small portion of the light curve. In Section 2.1, we adapted the Gregory-Loredo method to the planetary transit case by having one long out-of-transit bin (bin 0) and n short in-transit bins (see Figure 2.2). The value of n used was typically 4. For a given n , the parameters defining each candidate model are then p , e , and the transit duration d . The likelihood computation was carried out as described in G99.

This algorithm performed well when tested on simulated data (with photon noise only), but the likelihood calculation was still computationally intensive. A number of improvements were made following the completion of the simulations described in Section 2.1.3:

1. Given the current state of exo-planet research, the use of Bayesian priors is not expected to contribute significantly to the performance of the algorithm at the detection stage. The information available on period and duration distributions is relatively scarce for giant planets, and non-existent for terrestrial planets. The priors used in the modified GL method were generic and mostly identical to those used by G99 for X-ray pulsars, rather than specifically optimised for transit searches.
2. Using the χ^2 rather than the likelihood as a detection statistic, and implementing the calculation as outlined in Section 2.2.2.1, significantly reduces the computational requirements of the detection process.
3. The shape of most planetary transits is sufficiently simple that, for detection purposes (as opposed to detailed parameter estimation), a single in-transit bin, as illustrated in Figure 2.17, provides enough information. A significant advantage of this simplification is that it makes the method far more robust and capable of coping with real data, and all its concomitant problems, with negligible loss in detection efficiency.
4. Once a detection is made, a shape-estimation phase with either a large value of n , or by detailed model fitting of the phase folded light curve, can be implemented. As the dependency of transit shapes as a function of the stellar and planetary parameters is relatively well-known, Bayesian priors may have a part to play in this phase. This is, however, outside the scope of the present chapter.

2.2.2.3 χ^2 -minimisation with a box shaped transit.

The algorithm used in the present paper evolved from that of Section 2.1 taking into consideration the points listed in Section 2.2.2.2. The model therefore consists of one out-of-transit bin and a single level in-transit bin. (Although this simplification may seem disingenuous, by suitably pre-processing, or adaptively filtering, the signal to remove intrinsic stellar variability, this is a valid approximation to transit detection in practice.) All the data points falling into the out-of-transit bin form the ensemble O , while those falling into the in-transit bin form the ensemble I . No Bayesian priors are used. Adapting Equation (2.34) to this model gives

$$\tilde{\chi}^2 = \sum_{i=1}^N \frac{x_i^2}{\sigma_i^2} - \bar{x}_O^2 \sum_{i \in O} \frac{1}{\sigma_i^2} - \bar{x}_I^2 \sum_{i \in I} \frac{1}{\sigma_i^2}. \quad (2.35)$$

Provided the transits are shallow and of short duration (i.e. the most common case), the ensemble O contains the vast majority of the data points, so that $\bar{x}_O \approx \bar{x}$ (where

\bar{x} is the weighted mean of the entire light curve). Substituting this approximation into Equation (2.35):

$$\widetilde{\chi^2} \approx \sum_{i=1}^N \left\{ \frac{x_i^2}{\sigma_i^2} - \frac{\bar{x}^2}{\sigma_i^2} \right\} - \bar{x}_l^2 \sum_{i \in l} \frac{1}{\sigma_i^2}. \quad (2.36)$$

The first two terms in Equation (2.36) are constant. The minimisation of χ^2 is therefore achieved by maximising the statistic S^2 , given by

$$S^2 = \bar{x}_l^2 \sum_{i \in l} \frac{1}{\sigma_i^2}, \quad (2.37)$$

which can also be expanded as:

$$S^2 = \left(\sum_{i \in l} \frac{x_i}{\sigma_i} \right)^2 \left(\sum_{i \in l} \frac{1}{\sigma_i^2} \right)^{-1}. \quad (2.38)$$

If the light curve is robustly ‘mean-corrected’ prior to running the algorithm, such that x_i is replaced by Δx_i , \bar{x}_l becomes $\overline{\Delta x}_l$, the depth of the model transit. This results in a further simplification where the only free parameters are now the phase, period, and duration of the transit, since the depth is determined given the other three. It is also apparent that S^2 is simply equal to the square of the in-transit signal-to-noise ratio. This is easier to see in the case where $\sigma_i = \sigma$ for all i (a good approximation to the case for space data). Equation (2.38) then becomes

$$S = \text{SNR} = \frac{\sum_{i \in l} \Delta x_i}{n_l^{1/2} \sigma} = \frac{n_l^{1/2} \times \overline{\Delta x}}{\sigma}, \quad (2.39)$$

where n_l is the number of points in l , and $\overline{\Delta x} = \sum_{i \in l} \Delta x_i / n_l$ is the mean of the in-transit points, i.e. the model transit depth (the weighting being unnecessary in that case).

Equation (2.39) is used when the errors are constant, or when no individual error estimates are available for each data point. In the latter case, the Median Absolute Deviation (MAD) of the dataset is used to estimate σ , as this is more robust to outliers than a simple standard error estimate (Hoaglin et al. 1983). For a Gaussian distribution $\sigma_{rms} = 1.48 \times \text{MAD}$ and this factor is used throughout to scale the MAD sigmas. If individual error estimates are available, Equation (2.38) provides a more precise estimate of S at the cost of a slight increase in computation time.

If the noise is Gaussian, a theoretical signal-to-noise threshold (i.e. S threshold) can in principle be computed a priori to keep the false alarm rate below a certain value (Jenkins et al. 2002).

2.2.2.4 Detection of triangular or curved eclipses with a box-shaped model

The use of a single class of simple, box-shaped model greatly simplifies the problem, but how does it affect the sensitivity to highly triangular (grazing) eclipses? This can be quantified as follows.

Let $x(t)$ be a light curve which contains a single triangular eclipse of duration d and depth D starting at time e , and has a constant noise level σ . If we analyse this light curve using a matched filter⁴ with a model $r(t)$ of similar (triangular) shape, the maximum S_r in the matched filter statistic is obtained when the model parameters d_r , D_r and e_r are equal to d , D and e respectively. The expectation value of S_r is thus

$$\langle S_r \rangle = \int_0^T x(t) r(t) dt = \int_0^T r^2(t) dt = 2 \int_0^{d/2} r^2(t') dt'. \quad (2.40)$$

where $t' = t - e$. In the range $0 \leq t' \leq d/2$, $r(t') = -2Dt'/d$, so that

$$\langle S_r \rangle = \frac{8D^2}{d^2} \int_0^{d/2} t'^2 dt' = \frac{dD^2}{3}. \quad (2.41)$$

The variance of S_r is given by

$$V(S_r) = \int_0^T \sigma^2 r^2(t) dt = \sigma^2 \int_0^T r^2(t) dt = \sigma^2 \langle S_r \rangle, \quad (2.42)$$

so that the signal-to-noise ratio of S_r is

$$\text{SNR}_r = \frac{\langle S_r \rangle}{\sqrt{V(S_r)}} = \frac{d^{1/2} D}{3^{1/2} \sigma}. \quad (2.43)$$

Now let us perform the same analysis with a box-shaped dip model $b(t)$ of duration d_b and depth D_b starting a time e_b . The maximum in the detection statistic occurs when the centre of the dip in the model coincides with that of the eclipse in the light curve, i.e. when $e_b + d_b/2 = e + d/2$. One can see immediately that the optimal model duration and depth fulfil $d_b \leq d$ and $D_b \leq D$. The expectation value of the detection statistic for this model, S_b , is

$$\langle S_b \rangle = \int_0^T x(t) b(t) dt = 2 \int_{d/2-d_b/2}^{d/2} \frac{2Dt'}{d} D_b dt. = \frac{2DD_b}{d} \int_{d/2-d_b/2}^{d/2} 2t' dt' \quad (2.44)$$

$$\langle S_b \rangle = \frac{DD_b d^2}{2d} \left[1 - \left(1 - \frac{d_b}{d} \right)^2 \right] = \frac{DD_b d}{2} \left[2 \frac{d_b}{d} - \left(\frac{d_b}{d} \right)^2 \right], \quad (2.45)$$

⁴As outlined in Section 2.2.1, a matched filter is equivalent to a χ^2 minimisation approach.

while its variance is

$$V(S_b) = \int_0^T \sigma^2 b^2(t) dt = \sigma^2 \int_0^{d_b} D_b^2 dt = \sigma^2 D_b^2 d_b, \quad (2.46)$$

so that

$$\text{SNR}_b = \frac{\langle S_b \rangle}{\sqrt{V(S_b)}} = \frac{D d^{1/2}}{2\sigma} \left(\frac{d_b}{d} \right)^{-1/2} \left[2 \frac{d_b}{d} - \left(\frac{d_b}{d} \right)^2 \right]. \quad (2.47)$$

Maximising SNR_b with respect to d_b/d yields $d_b/d = 2/3$ and hence

$$\text{SNR}_b = \frac{2^{3/2} d^{1/2} D}{3^{3/2} \sigma}. \quad (2.48)$$

The difference in sensitivity is given by the ratio Equations (2.48) to (2.43):

$$\frac{\text{SNR}_b}{\text{SNR}_r} = \frac{2^{3/2}}{3} \approx 0.9428. \quad (2.49)$$

The loss of sensitivity is thus small: although the best-fit box-shaped model doesn't cover all the eclipse, most of the signal is concentrated in the central part which is covered. The expected loss for curved transits (where limb-darkening is important) is even smaller, as these resemble a box-shape more closely than triangles do.

2.2.2.5 Comparison with other transit search techniques

In following through the steps of the previous sections our prime motives were to modify a general purpose Bayesian periodicity estimation algorithm to make it simpler, faster and more robust. In so doing we have arrived at a very similar formulation to that developed by other authors, though the details of the implementation differ. For example, Kovács et al. (2002) derived and tested a box-fitting method (BLS) similar to the present algorithm on simulated ground based data with white noise, and showed that significant detections followed for in-transit signal-to-noise ratios (our S statistic) greater than 6.

Street et al. (2003) used a transit finding algorithm based on a matched filter technique. After identifying and removing large amplitude variable stars they generated model light curves consisting of a constant out-of-transit level and a single in-transit section. The models were generated for a series of transit durations and phases, and a χ^2 -like measure was then used to select the best model (indeed their Equation 3 is essentially a special case of the method derived in Sect. 2.2.1 for single transits).

Udalski et al. (2002b), who made the first direct detections of transiting planetary candidates later to be confirmed with the radial velocity method, also imple-

mented a version of the BLS algorithm and noted that it was much more efficient than their own algorithm based on “*a simple cross-correlation with an error-less transit light curve*” (Udalski et al. 2002a).

In a comparison of several transit finding algorithms, Tingley (2003a) found that matched filters and cross-correlation gave the best results compared with progressively more general methods ranging from BLS, through Deeg’s method (Doyle et al. 2000) to Defaÿ’s (Defaÿ et al. 2001) Bayesian approach. The fact that matched filters and cross-correlation methods give good results is hardly surprising, and can easily be deduced from the χ^2 minimisation developed in Section 2.2.1. The more general methods suffer from the added complexity of the underlying model, which through the Bayesian view of Occam’s Razor, reduces the tightness of the posterior probability distribution of the parameter estimation. What is however surprising, is that the BLS method did not give at least as good a result as the matched filter and cross-correlation methods. We would expect the BLS method to have similar performance to the matched filter as it is mathematically almost identical. In fact, the same author published more recently a revised comparison, in which he implemented the removal of the transit depth as a parameter in the BLS method, along the same lines as advocated here. After additional modifications to make the comparison method more rigorous, this modified BLS – now even closer to the present algorithm – compared well with the matched filter and cross-correlation (Tingley 2003b).

2.2.3 Improvements in the implementation

2.2.3.1 Optimised parameter space coverage

The formulation of the detection statistic presented in Section 2.2.2.3 is fully defined given only the dataset and the start and end times of each model transit. The model parameters are thus the duration d , period p and epoch e (defined for our purposes as the time at the start of the first transit in the dataset).

The range of expected transit durations is relatively small – from a few hours for close-in, rapidly orbiting planets, to almost a day for the most distant planets transiting more than once within the timescale of the planned observations. A simple discrete sampling prescription can therefore be adopted for the duration without leading to large numbers of trial values. One option is to choose the step δd between successive trial durations to be approximately equal to the average time step δt between consecutive data points. This ensures that models with the same period and epoch and neighbouring trial durations differ on average by ~ 1 data point per transit. However, if the observation sampling rate is high – such as the sampling rate of 10 min envisaged for most targets for *Eddington* in planet-finding mode (Favata

2004) – a larger step in duration can be used, provided it is smaller than the shortest significant feature in the transit, namely the ingress and egress, which have typical durations of ~ 30 minutes.

The period sampling prescription is designed to ensure that the error in the phase (or equivalently epoch) of the last model transit in the light curve is smaller than a prescribed value. Capping the error on the period (by using a constant trial period step) is not sufficient, as the error on the epoch of the n^{th} transit will be n times the error on the epoch of the first. This would lead to a larger overall error for shorter periods, where the number of transits in the light curve is large, thus introducing a bias in the distribution of detection statistic with period. This bias is not present if one uses a constant step in trial frequency. Defining the relative frequency $\nu_r = T/p$, T being the total light curve duration, the phase of an event occurring at time t is given by $\phi = 2\pi t/p = 2\pi t\nu_r/T$, so that for the last transit in the light curve $\phi \approx \phi_{\text{max}} = 2\pi\nu_r$. A fixed step in ν_r thus leads to a fixed error in ϕ_{max} . By trial and error, a value of 0.05 was found to be suitable for $\delta\nu_r$.

One caveat in the case of space missions with high sampling rates lasting several years, is that the above prescription can lead to very large numbers of trial periods. This implies that the overall algorithm must be extremely efficient. Some steps taken to optimise the efficiency are described below.

The phase, or epoch step interval, is set to the average sampling rate of the data since by so doing one can generate the phase information at no extra computational cost using an efficient search algorithm, detailed below.

2.2.3.2 A weighting scheme to account for non-continuous sampling

A further complication stemming from irregular sampling and from the finite duration of each sample, is that data points nominally corresponding to a time outside a transit may correspond partly to the out-of transit bin and partly to the in-transit bin. To account for this, the indices of points falling either side of the transit boundaries are also stored and included in the calculation of S , but with a weight which is < 1 and is inversely proportional to the interval between the time corresponding to the data point and the start/end time of the transit. This weighting scheme is particularly important for data with irregular sampling where transits might fall, for example, at the end of a night of ground-based observations, or even with space-based observations during a gap in the temporal coverage.

2.2.3.3 Speeding up the algorithm

By far the most time consuming operation in computing S and finding the set of parameters which maximises it, is the identification of the in-transit points, which must be identified for each model d , p and e . If one is dealing with a large number of light curves sharing the same observation times, it is more efficient to process many light curves simultaneously and compute $S(d, p, e)$ for the entire block of light curves for each set of parameters, as follows. For each trial period, the time array is phase-folded. At a given trial duration, the in-transit points are identified for the first trial epoch, by stepping through the folded time array one element at a time until the start time of the transit is reached, and then continuing, storing the corresponding indices, until the end time of the transit is reached. $S(d, p, e)$ is then computed and stored for each light curve. When moving to the next trial epoch, one steps backward through the folded time array from the end time of the old transit (which is stored between successive trial epochs) until the start time of the new transit is found. One then steps forward through the time array, storing the indices, until the end time of the new transit is reached. $S(d, p, e)$ is then computed and stored, and the epoch incremented, and so forth.

This minimises the overall number of calculations needed. As the number of in-transit points is the same for all light curves and σ only needs to be computed once per light curve (in the constant error case), this leaves only the sum of the in-transit points to be computed once per set of parameters and per light curve. The optimum number of light curves to process simultaneously depends on the amount of memory available.

A further speed increase is obtained by noting the redundancy within the computation of S for a range of phase/epoch and period trial values. Breaking down the search to a two-stage process consisting of a single transient event detector (essentially a matched filter stage) followed by a multiplexed period/phase search, removes the inner loop summation of data from the main search and gives a factor of ~ 10 improvement in execution time.

Example run-times computed using a laptop equipped with a 1.2 GHz Pentium IV processor with 512 MB of RAM are as follows. The light curves consisted of 157 680 floating point numbers, i.e. each was ~ 630 KB in size. The trial period and duration ranges were 180 to 400 days and 0.5 to 0.7 days respectively. These ranges are roughly appropriate to search for transits of planets in the habitable zone of a Sun-like star, and correspond to a total number of tested (p, d, e) combinations of $\sim 5 \times 10^7$. After finding the optimal number of light curves to search simultaneously, the runtime per light curve was ~ 4 seconds.

Note that close-in planets with periods below the range included in this sim-




Figure 2.18: Example detection statistic distribution for the $V = 14.5$ case, showing S as a function of trial epoch for the single event search (top), as a function of period for the multiple event search (middle) and as a function of epoch at the best period (bottom). The true epoch and period are 112 and 120 days respectively.




Figure 2.19: Example light curve corresponding to Figure 2.18. Top: full light curve with the positions and depths of the detected transits in red. Middle: phase-folded light curve (portion around transit). Bottom: full phase-folded light curve. Transits occur 112, 232, 352 and 472 days after the start of the light curve

ulation are, of course, of interest, so that lower trial periods (and hence lower trial durations) would also be included when searching for transits in real data, thereby increasing the runtime. As the trial period range is increased, the number of trial periods becomes prohibitively large due to the use of even sampling in frequency space (see Section 2.2.3.1): this leads to very small trial period steps at the low period end of the range if the steps are to be kept reasonable at the high period end of the range. This can be remedied by splitting the required range of trial periods and running the algorithm separately for each period interval. The runtime increases linearly with the number of trial durations.

2.2.4 Performance evaluation

Bootstrap simulations were carried out to evaluate the performance of the box-fitting algorithm in the same manner as for the modified GL method. We present here the evolution of the algorithm's performance as a function of magnitude, designed to verify that the simplifications which led from the modified GL method to the present one did not reduce the performance. Rather than going into more detailed simulations for different star-planet configurations at this stage, the box-fitting method will be tested more extensively in combination with simultaneously developed filtering tools to reduce the impact of stellar micro-variability (see Chapter 5).

To save time, a single transit duration value was used in the bootstrap simu-




Figure 2.20: Evolution of the box-fitting algorithm's performance (in terms of fractional error rates) with magnitude. Dotted line: false alarm rate. Dashed line: missed detection rate. Solid line: mean error rate.

lations (corresponding roughly to the FWHM of the input transits). Single tests on a given light curve with a range of trial durations did however show that, contrary to the modified GL method, the box-fitting method can be used to provide a rough estimate of the transit duration, which can then be refined using standard least-squares model fitting techniques on the phase-folded light curve.

Simulations were run for V-band magnitudes of 14, 14.5 and 15 for a system identical to that used in Section 2.1.3.3: a $1 R_{\oplus}$ planet orbiting a K5V star with a period of 4 months, the light curves lasting 16 months with 1 hour sampling. Even though the *Eddington* baseline design at the time these simulations were carried out had evolved from what it was at the time the modified GL method was tested, the photon counts expected with the older design were used to keep the comparison between the two algorithms fair. Figure 2.18 shows an example of the detection statistic distributions obtained for $V = 14.5$ for the single and multiple event search, while Figure 2.19 shows the corresponding light curve and phase-folded transits.

The results of these simulations are shown in Figure 2.20. Comparing with the bottom panel of Figure 2.10, a small improvement is indeed observed, the mean error rate being slightly lower at all three magnitudes (0, 0.5 and 5% are $V = 14$, 14.5 and 15 respectively, compared to 0, 2.5 and 8% with the modified GL method). Combined with the significant improvement in computing time, this makes the new algorithm decidedly more attractive for the detection stage. Running the modified GL algorithm on each set of 200 light curves (100 with and 100 without transits) took approximately 2 weeks on a Sun Sparc 5, while running the box-fitting method on




Figure 2.21: Evolution of the box-fitting algorithm's performance (in terms of fractional error rates) with the percentage of gaps in the data for $V = 14.5$. Dotted line: false alarm rate. Dashed line: missed detection rate. Solid line: mean error rate.

the same light curves took just over an hour on a laptop PC with 512 MB of RAM and a 1.2 GHz processor.

Tests were also conducted inserting short, random gaps (as was done for the modified GL method in Section 2.1.3.4) for the $V = 14.5$ case. The results are shown in Figure 2.21, showing a smooth, slow degradation of performance with increasing percentage of gaps. Given that the expected duty cycle for *Eddington* is over 95 %, short gaps of this type should not pose a problem for detection.

2.2.5 Discussion and future work

Through a process of simplification and consolidation, a robust least-squares box-fitting method for transit searching was derived from the modified GL method. Its performance is slightly improved compared to its 'ancestor', while the computational requirements have been vastly reduced.

In white noise, it is capable of reliably detecting periodic transits with a combined signal to noise ratio down to ≥ 6 , a limit similar to that found by the authors of its closest relative, the BLS of Kovács et al. (2002). As the *Kepler* mission was designed to produce a combined signal-to-noise ratio ≥ 8 for three transits of an Earth-analogue, the present algorithm should detect such events in *Kepler* data provided most of the stellar variability can be filtered out.

The fact that the detection statistic S is equal to the transit signal-to-noise ratio makes the interpretation of the distributions of the statistic with period and epoch, and thus the initial appraisal of potential candidates, relatively straight forward. The distribution of S with trial epoch at the best trial period (bottom panel of Figure 2.18) is, in simple terms, the convolution of the phase-folded light curve with a top-hat function of width equal to the trial duration, so that the actual duration (and to some extent shape) of the detected event, as well as the presence of other signal (secondary eclipses of a binary for example) can be assessed directly from that distribution. The reduction in sensitivity for triangular or curved eclipses due to the use of a box-shaped model is not expected to exceed 6 %.

The next step is to test this algorithm in a more realistic context, that is in the presence of data gaps and non-Gaussian noise. The two issues are linked: while non-Gaussian noise sources such as stellar micro-variability make a pre-processing stage necessary, the natural approach to filter out this noise involves Fourier domain decomposition. In the presence of data gaps, this is a non-trivial process, which is explored further in Chapter 4.

Following that, the algorithm has been applied blindly to simulated data containing a more complex realistic mix of noise sources, produced by a consortium of members of the COROT Exoplanet Working Group (see Chapter 6).

Chapter 3

Characterising and simulating stellar micro-variability

3.1 Introduction

The previous chapter focussed on the problem of transit detection, specifically in the white Gaussian noise case, isolating other issues affecting any planet search project, such as confusion and non-Gaussian noise sources. Tests of the algorithms presented therein on simulated data have shown they can perform very well in white Gaussian noise (as can other algorithms developed simultaneously or near-simultaneously by other authors), reliably detecting transits and quantifying their statistical significance.

However, rather than the detection of the transits themselves, the major difficulty for ground-based searches so far has in fact been distinguishing planetary transit-like events caused by stellar systems, such as eclipsing binaries with high mass ratios, or hierarchical triple systems (due to either a physical triple system or an eclipsing binary blended with a foreground star), from true planetary transits (Brown 2003).

In order to detect terrestrial planets, it is necessary to go to space, to avoid being affected by atmospheric scintillation and to monitor the target field(s) continuously, with minimal interruptions. With improved photometric precision comes an additional noise source, which is usually insignificant at the precisions achieved by ground based observations: the intrinsic variability of the stars, due mainly to the temporal evolution and rotational modulation of structures on the stellar disk. The Sun's total irradiance (see Figure 3.1) varies on all timescales covered by the available data, with a complex, non-white power spectrum (see Figure 3.3). The amplitude of the variations can reach more than 1 % when a large spot crosses the solar disk at activity maximum, compared to transit depths of tenths to hundredths of a percent. There is significant power on timescales of a few hours, similar to the typi-

cal transit duration. Untreated, solar micro-variability would significantly reduce the detection performance of missions such as *Eddington* or *Kepler* (see Section 2.1.4), while the variability levels of more active stars are expected to also affect COROT and even ground-based transit searches in young stellar environments.

However, it is possible to separate planetary transit signal and stellar variability because the former is of relatively well known shape and contains significant power at high frequencies, while the latter is concentrated mainly at low frequencies. Already, modified transit search algorithms, designed to distinguish between the transit and brightness variations of stellar origin, have been tested on simulated data including solar variability (Defaÿ et al. 2001; Jenkins 2002). Jenkins (2002) applied a simple scaling to the solar irradiance data to evaluate the impact of increased rotation rate. Nonetheless, a more physical model, in which the different phenomena involved can be scaled independently in timescale and amplitude for a range of spectral types and ages, is needed to simulate realistic light curves for stars other than the Sun. This will allow us to optimise, evaluate and compare different algorithms, but also different design and target field options for the space missions concerned.

The present chapter is concerned with the development of such a model. The philosophy adopted in the process is the following. Intrinsic stellar variability is by no means a well-understood process. Despite recent progress in the modelling of activity-induced irradiance variations on timescales of days to weeks in the Sun (Krivova et al. 2003; Lanza et al. 2003), the extension of these physical models to other stars remains problematic, due to the scarcity of information on how the timescales, filling factors of various surface structures, and contrast ratios, depend on stellar parameters. We have therefore adopted an empirical approach, using chromospheric flux measurements as a proxy measure of activity-induced variability. This step is possible due to the fact that a correlation between the two quantities is observed in the Sun throughout its activity cycle, as well as in other stars. Similarly, empirically derived relationships were used again to relate chromospheric activity, rotation, age and colour, rather than attempting to use models which make a number of assumptions about the physical process driving these phenomena, and generally depend on parameters which require fine-tuning.

The stellar micro-variability model is developed by extrapolating the results of a detailed Fourier analysis of total solar irradiance (TSI) variations (Section 3.2) to other spectral types and stellar ages through empirical scaling laws (Section 3.3). Tests of the model are presented in Section 3.4. The results are discussed in Section 3.5




Figure 3.1: PMO6 light curve, **a)** before and **b)** after the pre-processing steps described in Section 3.2.1.1. The data starts in January 1996.

3.2 Clues from solar irradiance variations

The Sun is the only star observed with sufficient precision and frequent sampling to permit detailed micro-variability studies, thanks to the recent wealth of data collected by the SoHO spacecraft, and particularly the full disk observations obtained by VIRGO (Variability of solar IRradiance and Gravity Oscillations), the experiment for helioseismology and solar irradiance monitoring on SoHO, (Frohlich et al. 1997).

Stellar micro-variability is difficult to observe from the ground due to its very low amplitude, except for very young, active stars – which are outside the main range of interest for planet searches. There is some information available on rms night-to-night and year-to-year photometric variability of a small sample of stars monitored over many years by a few teams (Radick et al. 1998; Henry et al. 2000b). We make use of these as they present the advantage of covering a range of stellar ages, but their irregular time coverage and limited photometric precision make them unsuitable for an in-depth study, and particularly for the detailed analysis of the frequency content of the variations.

A drastic improvement in our understanding of intrinsic stellar variability across the HR diagram is expected from the very missions this work is aimed at preparing. In the relatively short term, MOST will provide valuable information for a small sample of stars, but it is not until the launch of COROT, and later *Kepler* and *Eddington*, that a wide range of stellar parameters will be covered. In the mean time, we must make use of the detailed solar data, and make reasonable assumptions to extrapolate to other stars than the Sun.




Figure 3.2: Cumulative power spectrum of the PMO6 light curve after pre-processing, as a fraction of the total power at non-zero frequencies. The grey area shows the range of frequencies corresponding to typical transit durations (2 to 15 hrs). Over 95% of the power in the solar variations is below that range, which suggests that the two types of signal can be separated on the basis of their frequency signature.

3.2.1 SoHO/VIRGO total irradiance (PMO6) data

All SoHO/VIRGO data used in this work were kindly provided by the VIRGO team at ESTEC. The main instrument of interest was PMO6, a radiometer measuring total solar irradiance. The light-curves used here cover the period January 1996 to March 2001¹, which roughly corresponds to the rising phase of cycle 23.

3.2.1.1 Pre-processing of the data

The light curves were originally received as level 1 data, in physical units but with no correction for instrumental effects. Careful treatment was required to remove long term trends of instrumental and astrophysical origin. There was a difference of $\sim 0.24\%$ in the mean measured flux between the start and the end of the time series. Given that the observations roughly correspond to the interval between the minimum and the maximum of the Sun's activity cycle, one might expect to see a rise in the mean irradiance over that period. The instrumental decay may therefore be higher than the value quoted. However, the absolute value of the irradiance was of little interest for the present study, which concentrates on relative variations on time scales of weeks or less. Any long term trends in the data were therefore removed completely, regardless of whether they were of instrumental or physical origin². The decay appeared non-linear and there were discontinuities and outliers in the light curves, making a simple spline fit unsuitable.

The approach that was adopted consisted of a 5 step process:

¹Except for two interruptions roughly 1,000 days after the start of operations, corresponding to the "SoHO vacations", when the satellite was lost and then recovered.

²Note that fully corrected (level 2) data are now freely available from the World Radiation Centre in Davos, Switzerland for the period March 1996 to December 2002 (for data with minute sampling, more recent data is available with hourly sampling). Level 2 data is corrected for instrumental effects but contains the long-term trends due to the solar activity cycle. As only trends on timescales of 2 months or less were of interest here, the relatively crude corrections we applied to the level 1 data ourselves were sufficient.

- Visual inspection of the data was used to manually remove sections visibly affected by instrumental problems.
- Spline fits were performed on intervals chosen by visual inspection to start and end where discontinuities occurred. Each interval was divided by the corresponding fit, resulting in a normalised output light curve.
- A 5σ cutoff was applied for outlier removal.
- The sampling, originally 1 min, was reduced to 15 min to make the size of the light curves more manageable. This was done by taking the mean of the original data points in each 15 min bin, ignoring any missing or bad data points. It is unlikely any information on timescales shorter than 15 min would significantly impact the transit detection process, as the transits of interest here generally last several hours (corresponding to orbital periods of several months or years).
- Data gaps were replaced with the baseline value of 1.0, to allow the calculation of the amplitude spectra needed for the analysis³.

3.2.2 Modelling the ‘solar background’

Table 3.1: Typical timescales for the different components of the solar background.

Component	Timescale B (s)
Active regions	1 to 3×10^7
Super-granulation	3 to 7×10^4
Meso-granulation	$\simeq 8000$
Granulation	200 to 500
Bright points	$\simeq 70$

The power spectrum of the solar irradiance variations at frequencies lower than $\simeq 8$ mHz constitutes a noise source for helioseismology, usually referred to as the ‘solar background’. It is common practice to fit this background with a sum of powerlaws in order to model it accurately enough to allow the measurement of solar oscillation frequencies and amplitudes. Powerlaw models were first introduced by Harvey (1985). The most commonly used model in the literature today is that of Andersen et al. (1994), which is fairly similar: the total power spectrum is approximated by a sum of power laws, the number N of which varies between three and five depend-

³If d_i is a regularly sampled dataset, the dataset in which gaps have been replaced by 1.0 is $d'_i = d_i * w_i + 1 - w_i$, where w_i is window function, i.e. is 0 in the gaps and 1 elsewhere. The FT of d'_i is then, for all non zero indices k , $D'_k = D_k \otimes W_k - W_k$, where D_k and W_k are the FTs of d_i and W_i respectively.




Figure 3.3: **a)** Light grey: power spectrum of the PMO6 light curve (1996–2001). Dark grey: *idem*, smoothed with a boxcar filter. Thick solid line: multi-component powerlaw fit (see Section 3.2.2). Dotted lines: individual components of the fit. **b)** Solid line: amplitude of the low frequency component of the power spectrum (A_1 , computed as described in Section 3.2.3, using $L = 180$ days and $S = 20$ days) versus time, from 1996 to 2001. The gap at ~ 1000 days corresponds to a prolonged gap in the data. Dotted line: Chromospheric activity (BBSO Ca II K-line index) over the same period (arbitrary units), smoothed with a boxcar filter.

ing on the frequency coverage:

$$P(\nu) = \sum_{i=1}^N P_i = \sum_{i=1}^N \frac{A_i}{1 + (B_i \nu)^{C_i}} \quad (3.1)$$

where ν is frequency, A_i is the amplitude of the i^{th} component, B_i is its characteristic timescale, and C_i is the slope of the power law (which was fixed to 2 in Harvey’s early model). For a given component, the power remains approximately constant on timescales larger than B , and drops off for shorter timescales. Each power law corresponds to a separate class of physical phenomena, occurring on a different characteristic time scale, and corresponding to different physical structures on the surface of the Sun (see Table 3.1).

3.2.3 Evolution of the power spectrum with the activity cycle

In order to track the evolution of the solar background over the activity cycle, sums of power-laws – as given by Equation (3.1) – were fitted to the power spectrum of a section of data of duration L (typically 6 months). Such a fit is illustrated on the power spectrum of the entire dataset in the left-hand panel of Figure 3.3. The op-




Figure 3.4: Amplitude A_1 of the low frequency component of the solar power spectrum (computed as described in Section 3.2.3, using $L = 180$ days and $S = 20$ days), versus the BBSO CaII K-line index (arbitrary units) over the period 1996 to 2001.

eration is then repeated for a section shifted by a small interval S from the previous one (typically 20 days), and so on. Thus the evolution of each component can be tracked throughout the rise from solar minimum (1996) to maximum (2001) by measuring changes in the parameters defining each powerlaw.

A single component fit with parameters A_1 , B_1 & C_1 is made first. Additional components are then added until they no longer improve the fit, i.e. until the addition of an extra component does not reduce the χ^2 by more than 10^{-2} . The fit to the first section is used as the initial guess for the fit to the next section, and so forth. This method allows us to track the emergence of components corresponding to different types of surface structures throughout the solar cycle, as well as monitor variations in amplitude, timescale and slope for each component.

3.2.3.1 Results

The algorithm described above was run on the PMO6 data with $L = 180$ days and $S = 20$ days, and three components were found to provide the best fit in all cases. These components have stable timescales and slope, varying in amplitude only. The physical processes giving rise to each component are thus of a permanent nature. A number of points of interest emerge from the results. The first component, with $\tau \simeq 1.3 \times 10^5$ s (active regions) shows an increasing trend in amplitude which is well correlated with the CaII K-line index, an indicator of chromospheric activity. This is illustrated in the right-hand panel of Figure 3.3. The slope of the powerlaw is 3.8 (in good agreement with Andersen et al. 1998).

The observed correlation, which is further illustrated by the scatter plot in Figure 3.4, comes as no surprise. The passage of individual active regions across the disks of the Sun and other stars monitored by the Mt Wilson HK Project can be clearly seen in plots of the activity index S (from which R'_{HK} is derived) versus time⁴. On the other hand, the effect of the same type of event on the solar irradiance has been studied with a number of instruments, most recently VIRGO/LOI and PMO6 (Domingo et al. 1998). Recent models including contributions from faculae and sunspots of tunable size and number reproduce the PMO6 light curve to a high degree of precision (Krivova et al. 2003; Lanza et al. 2003).

However, observing and characterising a correlation throughout the Sun's activity cycle, between a chromospheric activity indicator which can be measured from the ground for a large number of stars, and total irradiance variations, whose amplitudes are so small they are only observable by dedicated high-precision photometric space missions, goes one step further. Most importantly for planetary transit searches, it implies that chromospheric activity indicators such as R'_{HK} can be used as a *proxy* to predict weeks timescale variability levels for a wide range of stars.

The amplitude of the second component also increases, but is not correlated to the Ca II index. This component corresponds to timescales corresponding to super- and meso- granulation, for which no detailed models are available to date. Our understanding of this kind of phenomenon is expected to improve dramatically when the results from space-based experiments designed for precision time-series photometry will become available.

3.2.3.2 Implications

The correlation between A_1 and the Ca II K-line index, although not extremely tight, is a clear indication that chromospheric activity indicators contain information about the variability level of the Sun on timescales longer than a few days. To establish more solidly a scaling law between photometric variability and chromospheric activity, we must use a wider stellar sample to constrain the relation over the entire expected range of activity levels (see Section 3.3.1.3).

Little useful information has been extracted from the solar data on what determines the parameters of the solar background other than A_1 . The few clues available from other sources will be presented in Section 3.3.2.

⁴See the Mt Wilson HK project homepage, http://www.mtwilson.edu/Science/HK_Project/.

3.3 Empirical scaling to other stars

The main aim of the present model is the simulation of realistic light curves of stars more active than the Sun, if possible as a function of stellar parameters such as spectral type and age, via observables such as $B-V$ colour and rotational period P_{rot} . The multi-component power-law model used in Section 3.2 to fit the solar background power spectrum now forms the basis of the simulation of enhanced variability light curves.

To simulate a light curve for a given star, the first task is to generate a power spectrum using this power-law model. When applying the component-by-component procedure described in Section 3.2.3 to fit solar power spectra, the optimal number of components was found to be three. We therefore use a sum of three power-laws to generate the stellar power spectra. The highest frequency component, a superposition of granulation, oscillations and white noise, has a characteristic timescale which is shorter than the typical sampling time for planetary transit searches, and thus can be replaced by a constant value (i.e. a random noise component). There are therefore a total of 7 parameters to adjust for each simulated power spectrum: three for each resolved power-law plus one constant.

The inputs of the model are the spectral type and age of the star. Starting from these theoretical quantities given, how are the power-law parameters deduced? Most of the information available concerns the amplitude of the first power-law, A_1 . The scaling of this parameter is described in detail in Section 3.3.1, while that of the other parameters of the model, whose treatment is much simpler, is discussed in Section 3.3.2.

3.3.1 The amplitude of the active regions component

We have established in Section 3.2 that there is a correlation between A_1 and the Ca II K-line indicator of chromospheric activity in the Sun. We will see in this section that such a correlation holds for a wide stellar sample. On the other hand, there is a well known scaling between rotation period, colour and chromospheric activity (Noyes et al. 1984). Provided one can estimate the rotation period of the star (Section 3.3.1.1), the activity level can be computed (Section 3.3.1.2), and from this one obtains A_1 (Section 3.3.1.3).

3.3.1.1 The rotation period-colour-age relation

As detailed in Section 3.3.1.2, it is possible to deduce the expected activity level for a star of known mass (i.e. colour) and rotation period. However, the number of stars

with known rotation periods is relatively small. In the context of the present work, it would thus be useful to be able to predict the rotation period for a given stellar mass and age.

Observational constraints on rotation rates for stars of known mass and age come from two sources: star forming regions and young open clusters, where one can measure photometric rotation periods or rotational line broadening ($v \sin i$); and the Sun itself. There is little else, as rotational measurements are hard to perform for all the quiet, slowly rotating intermediate age and old stars other than the Sun, except for relatively nearby field stars, for which little reliable age information is available. The status of observational evidence and theoretical modelling in this domain is outlined by Krishnamurthi et al. (1997), Bouvier et al. (1997) and Stassun & Terndrup (2003). Here we briefly sketch the current paradigm to set the context of the present work.

The observed initial spread in rotation velocities (from measurements of T-Tauri stars, with a concentration around $10\text{--}30 \text{ km s}^{-1}$ but a number of fast rotators, up to $100\text{'s of km s}^{-1}$) is attributed to the competing effects of spin-up (due to the star's contraction and accretion of angular momentum from the disk) and slowing-down mechanisms such as disk-locking (Königl 1991; Bouvier et al. 1997). This spread is observed to diminish with age (by the age of the Hyades, only some M-dwarfs still exhibit fast rotation, Prosser et al. 1995), leading to a dependency of rotation on mass only. Following this homogenisation, one observes constant spin-down in a given mass range. For example, Skumanich (1972), comparing rotational velocities Sun-like (G) stars in the Pleiades, the Hyades and for the Sun, found them to decay as the square root of age, implying that the rotation period of of main sequence Sun-like stars increases as $t^{1/2}$ where t is the age on the main sequence. Both homogenisation and power-law spin-down can be explained by the loss of angular momentum through a magnetised wind (Schatzman 1962; Weber & Davis 1967), a mechanism which is more effective in faster rotators.

Angular momentum evolution models have vastly improved recently, but they still rely on the careful tuning of a number of parameters, especially for young and low-mass stars. We have therefore chosen to use empirically derived scaling laws and to restrict ourselves to the range of ages (older than the Hyades) and spectral types (mid-F to mid-K) where a unique colour-age-rotation relation can be established. In this range the aforementioned parameters become less relevant and the models reproduce the observations fairly robustly.

A relationship between $B - V$ colour (i.e. mass) and rotation at a given age was empirically derived from photometric rotation period measurements in the Hyades (Radick et al. 1987, 1995). Only rotation periods were used, rather than $v \sin i$ mea-




Figure 3.5: Plot of rotation period versus $B - V$ colour for Hyades stars. Data from Radick et al. (1987, 1995). The thin parallel black lines correspond to 'rotational isochrones' from Kawaler (1989) with t as indicated next to each line (in units of 10^8 yr). The dashed line is a linear fit to the data with $0.6 \leq B - V < 1.3$. The thick solid line is a composite of two linear fits, one for $B - V < 0.62$ and one for $0.62 \leq B - V < 1.3$.

measurements, to avoid introducing the extra uncertainty of assigning random inclinations to the stars and having to assume theoretical radii to convert $v \sin i$ to a period. This relationship is valid for the range $0.45 \leq B - V \leq 1.3$. For stars bluer than $B - V = 0.45$, rotation rates saturate, but this is outside the range of spectral types of interest for the present work. Redder than $B - V = 1.3$, a significant spread is still observed in the rotation periods. The remaining range is divided into two zones, each following a linear trend. Redder than $B - V = 0.62$, the slope of the relation is quite close to the theoretical relation obtained by Kawaler (1989). We also estimated the age of the Hyades from the zero-point of the linear fit as done by Kawaler (1989), but incorporating the improved data from Radick et al. (1995). The age obtained in this manner is 634 Myr, consistent with recent determinations by independent methods: 655 Myr (Cayrel de Strobel 1990), 600 Myr (Torres et al. 1997), and 625 Myr (Perryman et al. 1998), thereby confirming the quality of the fit. The slope for stars bluer than $B - V = 0.62$ is much steeper, presumably due to the thinner convective envelopes of the stars in this range.

This can then be combined with the $t^{1/2}$ spin-down law into a rotation-colour-age relation:

$$\log(P_{\text{rot}}) - 0.5 \log\left(\frac{t}{625 \text{ Myr}}\right) = \begin{cases} -0.669 + 2.580 (B - V), & 0.45 \leq B - V < 0.62 \\ 0.725 + 0.326 (B - V), & 0.62 \leq B - V < 1.30 \end{cases} \quad (3.2)$$

A comment on the adopted value of 0.5 for n_t , the index in the spin-down law

is appropriate. It has recently been suggested that a value of 0.6 might be more appropriate (Guinan & Ribas 2002, on the basis of an updated sample of Sun-like stars with some new age determinations). However, the original value of 0.5 was kept for the present work. The change would not affect the predicted rotation rates significantly, and the errors on this new value of n_f (which depends, for example, on age determinations from isochrone fitting) are larger than the difference. We have therefore kept the lower value, as it leads, if in error, to overestimated rotation rates, hence more variability and on faster timescales, and eventually conservative estimates of transit detection rates.

3.3.1.2 The activity-rotation period-colour relation

The next step consists in estimating from the colour and rotation period the expected chromospheric activity level of the star. For this purpose, the scaling law first derived by Noyes (1983) and Noyes et al. (1984) is used. It relates the mean Ca II index $\langle R'_{\text{HK}} \rangle$ to the inverse of the Rossby number R_o , and can thus be understood in terms of the interplay between convection, rotation and the star's dynamo:

$$-\log R_o = 0.324 - 0.400 y + 0.283 y^2 - 1.325 y^3 \quad (3.3)$$

where $y = \log \langle R'_{\text{HK}} \rangle_5$ and $\langle R'_{\text{HK}} \rangle_5 = \langle R'_{\text{HK}} \rangle \times 10^5$. R_o is related to the rotation period P_{rot} and $B - V$ colour as follows:

$$R_o = \tau_c / P_{\text{rot}} \quad (3.4)$$

where P_{rot} is expressed in days, and the following (empirically derived) relation for τ_c , the convective overturn time, is used:

$$\log(\tau_c) = \left. \begin{array}{l} 1.361 - 0.166 x + 0.025 x^2 - 5.323 x^3, \quad x \geq 0 \\ 1.361 - 0.140 x, \quad x < 0 \end{array} \right\} \quad (3.5)$$

where $x = 1 - (B - V)$. Equation (3.3) was inverted (using an interpolation between tabulated values) to allow us to deduce the chromospheric activity index from the rotation period and $B - V$ colour of a star. For details of how the above relations, which are simply stated here, were obtained, the reader is referred to Noyes et al. (1984).




Figure 3.6: Photometric variability $[rms_{by}]_4$ versus chromospheric activity $\langle R'_{HK} \rangle_5$ for 72 field and Hyades stars. The colour coding refers to the source of the photometric variability data. Green: Henry et al. (2000b). Blue: Radick et al. (1998) (hollow: <10 observations per season). Grey: Lockwood et al. (1997) (hollow: $rms_{by} \leq 2$ mmag, upper limits only). Red: Hyades stars (filled: Radick et al. 1995, hollow: averages of rms_b & rms_y from Lockwood et al. 1984 & Radick et al. 1987). Solid line: 2nd order polynomial fit to the data. Dashed line: relation used in Aigrain et al. (2004), based on 29 stars only.

3.3.1.3 Active regions variability and chromospheric activity

Having noted that the ‘active regions’ component of the solar activity spectrum appears directly correlated to chromospheric activity (see Section 3.2.3), we turn to the small but valuable datasets containing both photometric variability measurements and activity indexes for a variety of stars. The data available on any one star is of course much less precise and less reliable than the solar data, but the dataset as a whole spans a much wider range of activity levels.

We use a sample of stars for which both $\langle R'_{HK} \rangle$ and photometric variability measurements are available in the literature. In Aigrain et al. (2004), we restricted ourselves to 29 stars whose variability levels and chromospheric activity indices were published in the same papers, namely Radick et al. (1998); Henry et al. (2000b). These data were then used to derive a quantitative relationship between $\langle R'_{HK} \rangle_4$ and the night-to-night rms variability in Strömgren b and y ($[rms_{by}]_4 = rms\{(b+y)/2\} \times 10^4$, in magnitude units).

Since then, a wider literature search has revealed a number of additional stars for which photometric variability and chromospheric activity information were available, but in separate sources: additional photometric variability data were taken from Lockwood et al. (1984); Radick et al. (1987, 1995); Lockwood et al. (1997), while activity data, when not found in the above sources, were taken from Barry et al. (1987); Duncan et al. (1991); Garcia-Lopez et al. (1993); Baliunas et al. (1995). The full sample, illustrated in Figure 3.6, now contains 72 stars.

At the same time as including the additional data in the calibration, additional care was taken to ensure consistency between datasets from different sources and

to minimise systematic errors.

Variability measurements

- Sparse time sampling: Some of the sources used, for example Radick et al. (1998), were mainly designed to study long term variations (similar to the Sun's activity cycle) and their observations were rather sparse, sometimes counting as few as 5 or 6 points per observing season. This could lead to an underestimate of, or in any case a larger uncertainty in, the short-term (night-to-night within 1 season) rms values, which were those used in Aigrain et al. (2004). In order to limit the impact of sparse time sampling, the average of night-to-night rms values from different observing seasons was used, including only seasons with more than ten observations whenever the number of observations per seasons was given. Lockwood et al. (1997) and Radick et al. (1998) gave only the total number of observations and the number of seasons. For those, stars for which the ratio of the two was less than 10 were flagged, to see if these appeared systematically less variable than stars observed more frequently.
- Exclusions: As was done in Aigrain et al. (2004), HD 95735, which was monitored by Henry et al. (2000b), was excluded from the analysis as it lies outside the range of $B - V$ colours considered in the model.
- Sensitivity limits: In the case of data from Lockwood et al. (1997), it should be noted that rms values ≤ 17 mmag are at the detection limit for variability and should be considered as upper limits only.
- Filters: Lockwood et al. (1984) and Radick et al. (1987) published b & y rms values separately. In later papers they noted the great similarity between the b & y light curves and averaged the two before computing rms values to improve precision. In the present work, when including data from Lockwood et al. (1984) and Radick et al. (1987), the average of the b and y rms values was used. These stars were also flagged, to highlight the fact that the rms values were not obtained in a totally consistent way with the others.
- Multiple appearances of a given star: When a given star was present in more than one source, the latest source with more than 10 observations per season was used.

Chromospheric activity measurements It is appropriate to mention here some precautions that were taken in the compilation of chromospheric activity index values. Two formulae for the temperature correction factor are commonly used (Middelkoop 1992; Rutten 1984), and published R'_{HK} values have been derived with both.

To ensure consistency, whenever possible S , the uncorrected quantity from which R'_{HK} is derived, was collected from the literature, and the correction factors of Noyes et al. (1984) were applied in all cases. They are somewhat intermediate between Middelkoop (1992) and Rutten (1984) and also include the photospheric correction. Several groups outside Mount Wilson Observatory have developed instrumental setups to measure the flux in the Ca II H & K lines that differ slightly from those of Mount Wilson, or setups to measure flux in other lines that are chromospheric activity indicators. In most cases, they derived conversions from their systems to S or R'_{HK} . These indexes were used for stars where no Mt Wilson system data was available, but care was taken to retain consistency in the temperature and photospheric correction factors used.

As illustrated in Figure 3.6, a 2nd order polynomial provides a good fit to the relationship between $[\text{rms}_{bV}]_4$ and $\langle R'_{\text{HK}} \rangle_5$. The equation of the fit is:

$$[\text{rms}_{bV}]_4 = 0.87 + 10.96 \langle R'_{\text{HK}} \rangle_5 + 2.06 \langle R'_{\text{HK}} \rangle_5^2 \quad (3.6)$$

The hollow symbols on Figure 3.6 indicate that, for one reason or another, the rms values are expected to be less reliable. As suspected, stars from Radick et al. (1998) with fewer than 10 observations per season appear on the lower envelope of the rms versus activity level trend (hollow blue circles). Also, some low activity Lockwood et al. (1997) stars that could be considered anomalously variable are in fact only upper limits (hollow grey circles). Finally, stars from Lockwood et al. (1984); Radick et al. (1987), for which the rms values were obtained by averaging single filter values, appear systematically high (hollow red circles). This could be explained if the individual rms values contain a significant contribution from measurement errors, as well as intrinsic variability. The rms of the average of the two light curves (the desired quantity) would then be systematically lower than the average of the rms of the two light curves. However, including or excluding these points does not affect the fit significantly, so they have been kept in Equation 3.6. The older fit, based on 29 stars only, is shown as the dashed line in Figure 3.6. There has been very little change over the range of chromospheric activity levels of interest, which bodes well for the robustness of the relation.

Note that one star, EK Dra, is significantly more active and variable than the rest of the sample – because it is younger. Including it in the fit thus gives it a disproportionate weight, but we have included it because it broadens the range of $\langle R'_{\text{HK}} \rangle$ covered. Furthermore, previous studies have shown (Messina & Guinan 2002) that it fits tightly on relationships between activity and stellar parameters derived from samples of solar analogues of various ages, suggesting its behaviour is representative of the mechanisms driving activity in general – and thus, in our reasoning, variability on

the timescales under consideration.

To be usable for our purposes, Equation (3.6) must be completed by a relationship between $[\text{rms}_{by}]_4$ and A_1 . The desired value of A_1 corresponds to white light relative flux variations, not b and y magnitude variations. To obtain the rms of the relative flux variations one must multiply the rms of the magnitude variations by a constant factor of $2.5/\ln(10) = 1.08$. This must then be converted to a white light flux rms value. This requires the definition of a reference point in the solar cycle, at which to compare the variability levels in the two bandpasses. Radick et al. (1998) quote a single $\log \langle R'_{\text{HK}} \rangle$ value of -4.89 for the Sun, which is an average of measurements performed over many years. This defines a reference solar activity level, corresponding roughly to a third of the way into the rising phase of cycle 23. The average night-to-night variability as measured in b and y by Radick et al. (1998) is $\text{rms}_{by} = 5 \times 10^{-4}$ mag. The corresponding white light variability level, measured from a 6 month long section of PMO6 data downgraded to 1 day sampling, centred on the date for which the measured activity level was equal to the reference level defined above, is $\text{rms}_{\text{white}} \approx 1.8 \times 10^{-4}$ mag.

This allows us to convert from rms_{by} in magnitudes to $\text{rms}_{\text{white}}$ in relative flux. Assuming a straightforward proportionality relationship the conversion factor is 2.78. This conversion introduces a significant error in the overall conversion, as we observed that the dependency of $\text{rms}_{\text{white}}$ on R'_{HK} (in the Sun) is slightly different in shape to that of rms_{by} on $\langle R'_{\text{HK}} \rangle$ (in the stellar sample), so a proportionality factor is inaccurate. However, until other stellar photometric time series that are sufficiently regular to perform the fitting process described in Section 3.2.3 are available, it is the best we can do. Finally one must convert from $\text{rms}_{\text{white}}$ to A_1 . As expected, a linear relationship between these quantities as computed for the Sun is observed, yielding an overall conversion between $[\text{rms}_{by}]_4$ and A_1 :

$$A_1 \times 10^5 = -0.24 + 0.66 [\text{rms}_{by}]_4 \quad (3.7)$$

Equation (3.6) thus becomes:

$$A_1 \times 10^5 = 0.33 + 7.23 \langle R'_{\text{HK}} \rangle_5 + 1.36 \langle R'_{\text{HK}} \rangle_5^2 \quad (3.8)$$

As space-based time-series photometric missions come online, we will be able to calibrate the relations above with more and more stars. Amongst these missions which have already provided useful data or will start doing so very soon are MOST (Micro-variability and Oscillations of Stars, Walker et al. 2003), a small Canadian mission which saw first light on 20 July 2003, and the OMC (Optical Monitor Camera, Giménez et al. 1999) on board ESA's new γ -ray observatory INTEGRAL. In the longer

term, COROT will also provide a wealth of information on stellar micro-variability.

These new data will allow us to calibrate the relations above with a wider variety of stars. In particular, chromospheric activity and rotational period measurements are available for a relatively large number of bright late type stars (Henry et al. 1996; Baliunas et al. 1996; Radick et al. 1998; Henry et al. 2000b; Tinney et al. 2002; Paulson et al. 2002). If any of these are observed by the missions listed above, yielding variability measurements, they will be incorporated in the present model.

3.3.2 Other parameters of the model

The previous sections were concerned with providing an estimate of one of the model's parameters, A_1 , given a star's age and colour. However, there are a total of 7 parameters to adjust. We have so little information on the super- and meso-granulation component in stars other than the Sun that we have chosen, for now, to leave it unchanged in the simulations, using the solar values.

3.3.2.1 The third component: timescales of minutes or less

The third (highest frequency) component observed in the Sun is a superposition of variability on timescales of a few minutes, which is thought to be related to granulation, and higher frequency effects such as oscillations and photon noise. The distinction between these effects is not resolved at the time sampling used for the present study. There is very little information at the present time on equivalent phenomena in other stars than the Sun (see below). Solar values were therefore used in all simulated light curves for this component. Given the low amplitude of this component, and the fact that it corresponds to timescales significantly shorter than the duration of planetary transits, it should not affect transit detection significantly.

Granulation can be traced by studying asymmetries in line bisectors, leading to the possible exploration of this phenomenon across the HR diagram (see for example Gray & Nagel 1989). Trampedach et al. (1998), who modelled the granulation signal for the Sun, α Cen A and Procyon, obtained very similar power and velocity spectra for the three stars, despite the fact that convection is much more intense on Procyon. Although very preliminary, these results suggest that the granulation power may change only slowly with stellar parameters, thus supporting the use of the solar values in the present model.

3.3.2.2 The second component: hours timescale

Again, solar values were used for all simulated light curves for this component. Work is underway to identify the types of surface structures giving rise to the hours-timescale variability in the Sun (Fligge et al. 2000), and the launch of the COROT mission will provide a dataset ideally suited to improving our understanding of this type of variability. Keeping the super- and/or meso-granulation component identical to the solar case is certainly an oversimplification, and it is the area where most effort will be focused in the future, as it is highly relevant to transit detection, being on timescales similar to transits.

The possibility of measuring the power spectrum of stochastic luminosity variations in stars other than the Sun on timescales of minutes to hours, using existing data from the star-tracker camera of NASA's WIRE (Wide Field Infrared Explorer), is also under investigation⁵. However, difficulties associated with non-Gaussian noise and frequent data gaps both short and long, have impeded such a measurement so far (see Section 3.4.3.4).

3.3.2.3 Timescale of the (first) active regions component

Two parameters remain for the active regions component. We have kept the slope of the power law, C_1 , unchanged from the solar case. Changing it slightly does not seem to affect the appearance of the light curve significantly. More crucial is the timescale B_1 .

If the active regions component of micro-variability is the result of the rotational modulation of active regions, we expect B_1 to be directly related to the period. However, the value obtained for the Sun is $B_1 = 8.5 \times 10^6$ s, i.e. 9.84 days, compared to a rotational period of ≈ 26 days. This suggests that the timescale is not (or not exclusively) dominated by rotational modulation of active regions, but by the emergence and disappearance of the structures composing the active regions. Individual active regions evolve on a timescale of weeks to months in the Sun (Radick et al. 1998), but spots and faculae evolve faster: observed sunspot lifetimes range roughly between 10 days and 2 months (Hiremath 2002). We therefore deduce that sunspot evolution is the process which determines B_1 in the Sun.

One can in fact interpret the observed timescale purely in terms of rotational modulation: the transit of a given spot across the solar disk would last half the rotation period only, so that, given the shape of a sunspot transit signature, pure rotational modulation could contribute to signal on timescales as short as a third of the rotation

⁵This data has already been used to perform asteroseismology on a number of bright stars (Schou & Buzasi 2001) following the failure of the main instrument shortly after launch.

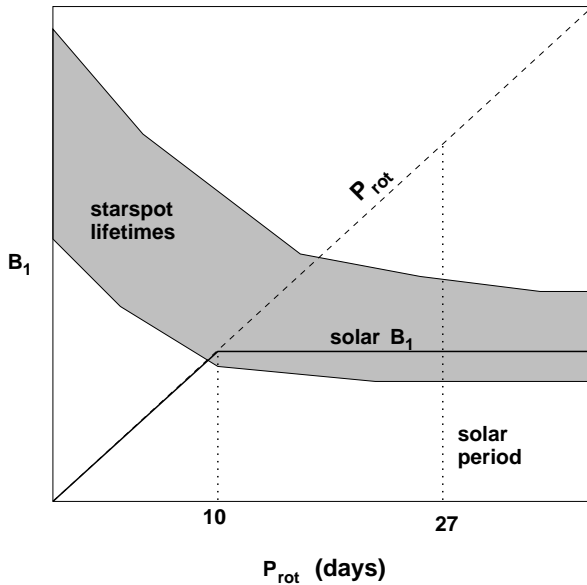


Figure 3.7: Schematic illustration of the expected dependence of the active regions component's characteristic timescale B_1 on rotational period.

period, approximately equal to the measured B_1 . This serves only to highlight the need for direct observations of this timescale in other stars.

Whatever the dominating factor in the Sun, the phenomenon with the shortest timescale (in the relevant range) is expected to determine B_1 for any given star. The present version of the model assumes that, in the Sun, this phenomenon is sunspot evolution. As spot lifetimes are, if anything, longer in faster rotators (Barnes et al. 1998; Soon et al. 1999), rotational modulation should take over below a certain period. As a rough estimate we have placed the boundary between the two regimes at $P_{\text{rot}} = 10$ days (see Figure 3.7). If on the other hand the important phenomenon is rotation, but a given rotation period gives rise to a timescale $B_1 \simeq P_{\text{rot}}/3$, then this rotationally dominated regime could dominate up to rotation periods of at least ~ 30 days. Such a modification would be trivial to implement should upcoming stellar data with sufficiently long time baselines, such as those from MOST and COROT, warrant it.

This completes, within the obvious limits of the assumptions used, the requirements for the simulation of white light stellar light curves with micro-variability, within the range of applicability of the scaling laws used: $0.45 \leq B - V \leq 1.3$, $t \geq t_{\text{Hyades}}$, and the star must still be on the main sequence. The stellar parameters required are age (or rotation period) and $B - V$ colour (or spectral type). It is also possible to supply R'_{HK} , or A_1 and B_1 directly. Once an artificial power spectrum is generated, phases drawn at random from a uniform distribution are applied before applying a reverse Fourier transform to return to the time domain.

Due to the use of randomly chosen phases, the shape of the variations does




Figure 3.8: Comparison of a portion of PMO6 data starting 1300 days after the start of the full light curve (top panel) with a simulated light curve generated using the Sun’s observed rotation period (25.4 days) and chromospheric activity index ($R'_{\text{HK}} = -4.89$, Radick et al. 1998). Both light curves have 15 min sampling, last 180 days and are normalised to a mean flux of 1.0.

not closely resemble the observed solar variability. It may be possible to characterise the sequence of phases characteristic of a given type of activity-related event, such as the crossing of the stellar disk by a star-spot, faculae or active region. This information could then conceivably be included, with appropriate scaling, in the model, in order to make the shape of the variations more realistic. However, how to do this in practice is not immediately obvious, and will be investigated in the future. As the model stands, the simulated light curves can be used to estimate quantities such as amplitude, timescale, the distribution of residuals from a mean level, but no conclusions should be drawn from the shape of individual variations.

3.4 Testing the model

3.4.1 Mimicking the Sun

A first check is to compare the predicted observables for the Sun to the measured values. The measured period and $\log \langle R'_{\text{HK}} \rangle$ are 25.4 days and -4.89 (Radick et al. 1998), in good agreement with the predicted values of 23.4 days and -4.84 . The measured rms_{by} is $\simeq 5$ in units of 10^{-4} mag. The value measured from a simulated light curve with 1 day sampling lasting 6 months, when allowing for the conversion between $\text{rms}_{\text{white}}$ and rms_{by} is 6.11. This slight over-prediction is attributable to the Sun’s slightly slow rotation and low activity for its age and type, and to the fact that it is slightly under-variable for its activity level (it falls slightly below the fit on Figure 3.6). Figure 3.8 compares a portion of the PMO6 light curve taken from a relatively high




Figure 3.9: Examples of simulated light curves containing micro-variability for Hyades age (left column) and solar age (right column) stars with spectral types F5 (top row), G5 (middle row) and K5 (bottom row). A single transit by a $0.5 R_{\text{Jup}}$ planet has been added to each light curve 75 days after the start (a transit by an Earth-sized planet would be ≈ 25 times smaller). The light curves have 1 hr sampling and last 150 days.

activity part of the solar cycle with a light curve simulated using the observed rotation period and activity index of the Sun. The amplitude and typical timescales of the variations are well matched.

3.4.2 Trends with age and mass

A set of six light curves have been simulated, corresponding to three spectral types (F5, G5 and K5) and two ages (625 Myr and 4.5 Gyr). Examination of the light curves and the various parameters computed during the modelling process can reveal any immediate discrepancies. The light curves are shown in Figure 3.9 and the parameters in Table 3.2. They follow the expected trends, variability decreasing with age and $B - V$ and increasing with P_{rot} , so that at solar age the least active star is the F star, while the most active is the K star, despite its long rotation period (due to the dependence of activity on colour).

Table 3.2: Parameters of the simulated light curves.

Age Gyr	SpT	$B - V$	P_{rot} days	$\log(R'_{\text{HK}})$	A_1 $\times 10^5$	B_1 days	$\text{rms}_{\text{white}}$ $\times 10^4$
0.625	F5	0.44	2.9	-4.64	19.27	2.89	22.0
0.625	G5	0.68	8.8	-4.44	38.16	8.80	46.5
0.625	K5	1.15	12.6	-4.42	41.40	9.84	69.5
4.5	F5	0.44	7.8	-5.23	4.66	7.87	13.3
4.5	G5	0.68	23.7	-4.80	11.95	9.84	13.9
4.5	K5	1.15	33.8	-4.67	17.53	9.84	14.4

3.4.3 Behaviour at high activity

Comparing the light curves shown in the left column of Figure 3.9 which have amplitudes of $\sim 0.5\%$, to published V -band amplitude measurements for Hyades stars (Messina et al. 2001, 2003), which are of the order of $\sim 3\%$, immediately highlights a discrepancy. This discrepancy could be explained by a number of factors:

3.4.3.1 Bandpass

Part of the difference is readily explained by the fact that the model produces white light flux variations, while the amplitudes reported by Messina et al. were V -band magnitude variations. As previously mentioned, white light flux variations in the Sun are observed to be ≈ 2.78 times smaller than b & y magnitude variations, and a similar effect is expected with V . To estimate the amplitude of such an effect requires the comparison of simultaneous light curves in V and either b & y or white light. S. Messina (priv. comm.) did this comparison and found near-identical variability levels both in amplitude and in rms. The bandpass effect could therefore account for no more than a factor of ~ 2.8 , which is not enough to explain the observed discrepancy.

3.4.3.2 Underestimated rms values at high activity

As outlined in Section 3.3.1.3, the activity-variability relation remains ill-constrained at high activity levels, with only one data point with $[R'_{\text{HK}}]_5 > 5.5$. Although many stars have been used to calibrate the relation over the range of activity levels typical of the Hyades – $1.6 < [R'_{\text{HK}}]_5 < 5$ – there is a lot of scatter over that range and different sources of data show different trends, indicating that systematics are still present. However, given that the inclusion or exclusion of such “dubious” datasets in the calibration hardly alters the scaling law, this is unlikely to be a significant factor.

3.4.3.3 Amplitudes versus rms

Simulated and observed light curves for the Sun were compared in Section 3.4.1, and both rms values and amplitudes are consistent. For younger stars, the calibration of the scaling law used should ensure that the rms of the simulated light curves approximately agrees with observations (see previous paragraph). However, if the amplitude scales differently from the rms as the activity level increases, the model could produce sensible rms values but underestimate the amplitude for active stars. This is possible if the statistical nature of the variability changes from relatively stochastic (dominated by the emergence, evolution and disappearance of spots, as we think is the case in the Sun) to close to sinusoidal (dominated by rotational modulation of small numbers of large, persistent active regions, as is generally thought to be the case in young active stars). In this case, a new component, more concentrated in period space around the stellar rotation period, could be added to the simulated power spectra.

3.4.3.4 WIRE time series data of α Centauri

The WIRE⁶ or Wide field InfraRed Explorer is a NASA satellite designed to perform sky surveys in the IR, whose original goals could not be fulfilled due to loss of detector coolant early in the mission. However, the rest of the satellite is in good condition and its 5 cm aperture star-tracker telescope has been used successfully for asteroseismic studies of a number of stars, including α Cen (Schou & Buzasi 2001). α Cen is a particularly interesting target to search for low-frequency variability such as is observed in the Sun, as it is a well studied binary whose primary has a spectral type and age close to that of the Sun. Recently, Kjeldsen et al. (1999) measured excess power in ground-based radial velocity observations of it the range 600 to 3000 mHz, to which they fit a powerlaw with a slope of -1.46 , and which they interpret as granulation.

An attempt was made to use the WIRE time series, which lasts 50 days with approximately 40% duty cycle (observations are taken during just under half of each 102 min orbit) to perform an analysis similar to that which was done for the solar PMO6 data, fitting a multi-component broken powerlaw model to the power spectrum to obtain further constraints on the various parameters of our micro-variability model. However, the complex window function could induce many features in the power spectrum. To check for these, a Fisher randomisation test was carried out: a 'scrambled' dataset was constructed by drawing samples in random order from the original dataset, keeping the same window function. There were no noticeable differences between the power spectra of the scrambled and unscrambled datasets.

⁶<http://www.ipac.caltech.edu/wire/>

This implies that any features in the power spectrum of the unscrambled dataset are in fact due to noise and the effects of the time sampling.

As α Cen was the WIRE target with the most complete light curve to date, this suggests that little useful information relating to stellar micro-variability on the timescales of interest for transit searches is likely to be extracted from WIRE data. The first measurements of this type of variability in other stars than the Sun (excluding active stars) are expected to become available in the near future thanks to the MOST (Micro-variability and Oscillations of Stars) satellite.

3.5 Discussion

A model to generate artificial light curves containing intrinsic variability on timescales from hours to weeks for stars between mid-F and late-K spectral type and older than 0.625 Myr has been presented. This model relies on the observed correlation between the weeks timescale power contained in total solar irradiance variations as measured by VIRGO/PMO6 and the Ca II K-line index of chromospheric activity. Except for the most active cases, the resulting light curves appear consistent with currently available data on variability levels in clusters and with solar data. Further testing and fine-tuning requires high sampling, long duration space-based stellar time-series photometry and will be carried out as such data become available.

The simulated light curves can be used to test the effectiveness of pre-transit search variability filters (see Chapter 4) and to estimate the impact of micro-variability on exo-planet search missions such as COROT, *Eddington* and *Kepler* (see Chapters 5 and 6).

All the simulated light curves produced so far, as well as the routines used to generate them, have been made available to the exo-planet community through the web page: www.ast.cam.ac.uk/~suz/simlc. Light curves with specific parameters can be generated on request. The model has been included in light curve simulation tools developed for the COROT mission.

It is important to stress that the present approach, which is highly simplified and empirically based (and which will be referred to hereafter as 'the simulator', for clarity), differs from, but complements, more detailed theoretical models of the complex physical mechanisms that give rise to the observed micro-variability of the Sun (here after 'detailed models'). There has been recent progress in the latter area: for example, Seleznyov et al. (2003) successfully model the full power spectrum of the VIRGO/PMO6 observations by combining the magnetic activity model of Krivova et al. (2003) (which relies on identifying active regions in resolved solar disk images and modelling their photometric signature) and a simple granulation model. Scaling

this approach to other stars is non-trivial, but there will be much to learn from detailed models about how the physical mechanisms at work vary from star to star when upcoming high precision stellar photometric time series data become available. On the other hand, our simulator has already been put to use to simulate light curves for a wide range of spectral types and ages, thereby providing insights relevant to the design of transit search missions and associated data analysis techniques. However, its simplicity naturally limits its information content. Its applications rest more on the statistical side: provided the overall trends are correct, it can be used, for example, to place planetary transit detection limits for a given type of parent star.

Future improvements of the simulator will be two-fold. Of course, direct fitting of the power spectra of other stars than the Sun, from MOST or COROT, will provide direct constraints. In parallel, regions of parameter space not yet covered by these missions can be explored using detailed theoretical models, to unveil the dependence of the power spectrum on the parameters of the models. If the dependence of those on stellar parameters is known, or if reasonable assumptions to that effect can be made, the simulator can be adjusted to reproduce similar trends in the power spectra it generates.

Another step currently under consideration is the extension of the simulator to produce light curves in specific spectral band passes. This would be very useful in the context of COROT (and potentially *Eddington*) whose design includes the use of colour information, both for target selection and the preparation of colour-discriminant light curve analysis tools.

Chapter 4

Variability filters

4.1 Introduction

We have seen that stellar micro-variability will be an important noise source for space-based planetary transit searches, and that it is vital to reduce its impact for *Kepler* and *Eddington* to achieve their stated goals. The basic idea behind the variability filters developed in the present chapter is the following: it is possible to disentangle the planetary transit signal from other types of temporal variability if the two have sufficiently different temporal characteristics.

The micro-variability simulator introduced in Chapter 3 provides us with the means of illustrating and testing the effect of the filters we develop as solutions to this problem. As an example, we use throughout the present section a light curve simulated according to the planned characteristics of the *Eddington* mission, containing stellar variability, planetary transits and photon noise. The light curve lasts 3 years and has a sampling time of 10 min. The transits were simulated using the Universal Transit Modeller (UTM) software of Deeg et al. (2001), while the photon noise was simulated as Gaussian distributed noise with a standard deviation equal the square-root of the expected photon count per integration given the collecting area (0.764m^2) and throughput of the December 2003 *Eddington* baseline design¹. The light curve contains transits of a $2 R_{\oplus}$ planet orbiting a G2V star ($R_{\star} = 1.03 R_{\odot}$), i.e. a radius ratio of 0.018, leading to a relative transit depth of 3.24×10^{-4} . The planet's orbital period is 1 year, and its orbital distance 1 AU, leading to a transit duration of ~ 13 hours. The epoch of the first transit is 1.5 day. The star's age is 4.5 Gyr and its apparent magnitude $V = 13$, leading to a photon count rate of 8.4×10^7 10 min integration. In this regime, the photon noise in each integration is well approximated by a Gaussian distribution with a standard deviation of 1.09×10^{-4} .

¹<http://astro.estec.esa.nl/Eddington/Tempo/eddiconfig.html>




Figure 4.1: Simulated Eddington light curve for a $V = 13$ solar-age G2V star orbited by a $2R_{\oplus}$ planet with a period of 1 year. Top panel: entire light curve. Bottom panel: first 30 days, with a transit 1.5 day after the start. The flux values shown have been normalised to a mean of 1.




Figure 4.2: Power spectrum of the light curve shown in Figure 4.1 (upper grey line). Lower grey line: stellar variability only. Lower black line: transits only (3 transits). Upper black line: photon noise. The power spectrum is dominated by stellar variability at low frequencies and by photon noise at high frequencies.

The power spectra of the different components of the light curve mentioned above are shown in Figure 4.2. Although the power contained in the transit signal is small compared to both stellar and photon noise components (and would be even smaller for the case of an Earth-sized planet), it retains significant power for frequencies higher than $\sim 1 \mu\text{Hz}$, where the stellar signal starts to drop off steeply. As long as the aforementioned condition that the planetary and stellar signal be sufficiently well separated in the frequency domain is fulfilled (i.e. if the stellar variability occurs on sufficiently long timescales), one should be able to separate and detect the transits. Furthermore, in the case of multiple transits, the regular period of the transits also helps constrain the Fourier space occupancy of the transit signal with respect to the stellar signal.

In the spirit of modularity adhered to throughout this thesis, the filters are developed as pre-processing tools, the output of which can be fed to a transit search algorithm. This differs from previous publications on the topic (Defaÿ et al. 2001; Jenkins 2002), which concern transit search algorithms specifically designed to detect signals buried in non-white noise. One advantage of our approach is that the filtered light curves can be searched for any kind of short-timescale event, not only transits.

The work presented in this chapter is the continuation of work carried out on this issue at ESTEC (Carpano et al. 2003). The latter article was a detailed exploration of a pre-whitened matched filter which we show in Section 4.2 to be closely related to a Wiener filter. This filter is then generalised to be applicable to data with gaps and/or irregular sampling in Section 4.3, while an alternative, the iterative non-linear filter, is presented in Section 4.4. This marks the difference between our approach and that of Defaÿ et al. (2001); Jenkins (2002) and Carpano et al. (2003): the filters are required to be *directly* applicable to data with gaps and irregular sampling, a problem that will affect any real dataset to some extent, whether ground- or space-based. Note that Jenkins (2002) does address this issue, but in a different way: by developing a method to effectively regularise the sampling before applying the algorithm. The characteristics of the light curves after application of the two filters are compared in Section 4.5, and their performance, particularities and potential improvements are discussed in Section 4.6.

4.2 Wiener or matched filtering approach

Carpano et al. (2003) demonstrated how use of an optimal filter can simultaneously pre-whiten and enhance the visibility of transits in data dominated by stellar variability. The Fourier-based method presented there is also closely related to a minimum mean square error (MMSE) Weiner filter. However, even for space-based missions uneven sampling of the data will occur. In these real-life cases, standard Fourier methods are no longer directly applicable and a more general technique is required.

To gain some insight to the problem consider the general case of intrinsic stellar variability, with the received signal $x(t)$ is composed of the three components:

$$x(t) = s(t) + r(t) + n(t) \quad (4.1)$$

where $s(t)$ is the intrinsic time variable stellar light curve, $r(t)$ is the transiting planet signal, and $n(t)$ denotes the measurement plus photon noise, which we can take to be random (and Gaussian in the cases of interest here)². Each component is statistically independent, hence the expected power spectrum $\Phi(\omega)$ of the received signal is simply given by:

$$\Phi(\omega) = \langle |S(\omega)|^2 \rangle + \langle |R(\omega)|^2 \rangle + \langle |N(\omega)|^2 \rangle \quad (4.2)$$

and in the case of random, or white, noise $\langle |N(\omega)|^2 \rangle$ is a constant, hence guarantee-

²Strictly speaking, the 1st two terms in Equation (4.1) should be multiplicative, but in the limit of low amplitude variability and shallow transits, an additive combination is a very good approximation.

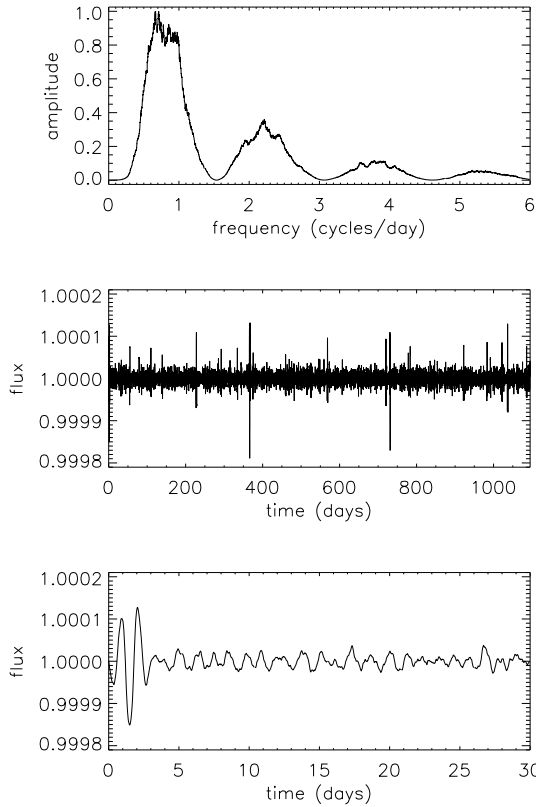


Figure 4.3: Top panel: Wiener filter constructed using the light curve shown in Figure 4.1 and a reference box-shaped transit of duration 0.65 day. Middle panel: filtered light curve. Bottom panel: idem, 1st 30 days, with a transit 1.5 day after the start.

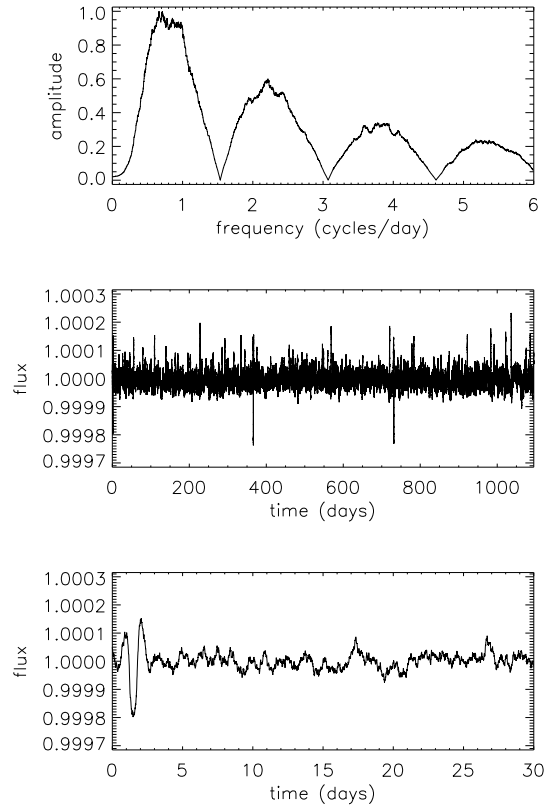


Figure 4.4: Top panel: matched filter constructed using the light curve shown in Figure 4.1 and a reference box-shaped transit of duration 0.65 day. Middle panel: filtered light curve. Bottom panel: idem, 1st 30 days, with a transit 1.5 day after the start.

ing positivity of the right hand term. This also highlights in a natural way a justification for the somewhat arbitrary constant in Equation (6) in Carpano et al. (2003) and how its value is related to the expected noise properties (although it would be more natural to implement it as a lower bound). However, as outlined below there is a simpler way to implement their technique without the need for the additional constant.

A standard MMSE Wiener filter attempts to maximise the signal-to-noise in the component of interest, in this case $r(t)$, by convolving the data with a filter, $h(t)$, constructed from the ratio of the cross-spectral energy densities between observation and target, such that:

$$x'(t) = h(t) \otimes x(t) \quad X'(\omega) = H(\omega) X(\omega) \quad (4.3)$$

and (using * to denote complex conjugate):

$$H(\omega) = \frac{\langle R(\omega)R(\omega)^* \rangle}{\langle X(\omega)X(\omega)^* \rangle} = \frac{\langle |R(\omega)|^2 \rangle}{\langle |X(\omega)|^2 \rangle} \quad (4.4)$$

for a long enough run (a fair sample) of observations. In practice the only example we have of $x(t)$ is often singular, implying that the best estimate of the denominator is simply the observed power spectrum $\Phi(\omega)$, subject to the constraint of positivity imposed by the implicit $\langle |N(\omega)|^2 \rangle$ term. Such a filter is illustrated in Figure 4.3: the top panel shows the filter, constructed using the Fourier transform of the light curve shown in Figure 4.1 and a box-shaped reference transit of duration 0.65 day, and the bottom two panels show the filtered light curve. Note that this filtering method does modify the transit shape. In particular, it induces positive deviations either side of the transit, which is effectively equivalent to removing some of the transit signal as well as the stellar and noise signal. However, the transit signal-to-noise ratio is obviously enhanced, and it becomes easily discernible even by eye.

This should be contrasted with the pre-whitened matched detection filter employed by Carpano et al. (2003), illustrated in Figure 4.4 (using the same layout as Figure 4.3), and which can be written in the form:

$$X'(\omega) = H(\omega) X(\omega) = \frac{X(\omega)}{\langle |X(\omega)| \rangle} \langle |R(\omega)| \rangle \quad (4.5)$$

and hence is equivalent to reconstructing the data using just the phase of the input signal Fourier transform modulated by the amplitude spectrum from the expected transit shape (see Figure 4.5). Viewing the problem in this way removes the need for the additional constant in their Equation (6) and emphasises the two stage nature of the filtering. The pre-whitening suppresses the stellar variability component, while the matched filter is directly equivalent to the transit search algorithm case presented in Section 2.2.2 with $n = 1$.

In practice, transit searching can be based directly on the output of the filtering, or preprocessing can be used to decouple the stellar variation estimation from the transit search phase, which then proceeds using algorithms optimised for white Gaussian noise. (In either case, detailed investigation of the transit depth and shape involves phase folding, unfiltered data, and local modelling.)

Either of these preprocessing filters works well in the case of regularly sampled data with no gaps and with a reasonable separation between the signatures of the Fourier components of the transits and the stellar variability. In Figures 4.3, 4.4 & 4.5, the transits are distinctly visible in the filtered light curve. The results in terms of transit detection performance using either method are very similar. For simplicity, the




Figure 4.5: As Figure 4.4, but the filtered light curve was obtained by modulating the phase of the Fourier transform of the data by the amplitude spectrum of the reference transit signal. The filter was omitted as it is effectively identical to that shown in Figure 4.4. Comparing, visually, the amplitude, shape and timescale of the variations in the filtered data with the bottom two panels of Figure 4.4 confirms that this gives very similar results to the matched filter approach.

matched filter approach, rather than the Wiener filter, is used in the remainder of this paper.

However, real data, even space-based, suffers from irregular sampling and the presence of significant gaps. Fourier domain methods cannot be directly applied to irregularly sampled data, but it is possible to treat regularly sampled data with gaps as a series of n independent time series, and to filter them separately. To test this, four arbitrarily chosen sections were removed from the light curve shown in Figure 4.1 (see Figure 4.6). The matched filter was then applied to the five unbroken intervals separately, and the results are shown in Figure 4.7. Though the filtering is effective on relatively long sections of data (bottom panel) it is not successful for short intervals (middle panel), even if they are significantly longer than the transit duration. This is because the power spectrum of the stellar noise is estimated from the data in order to construct the filter. For this to be successful, the data segment needs to be at least twice as long as the longest significant timescale in the star's variability, which is either the rotation period or the long end of the starspot lifetime distribution




Figure 4.6: Simulated light curve with data gaps. Four arbitrarily chosen sections were removed from the light curve shown in Figure 4.1. Note that the gaps were chosen to avoid the transit regions.

(Aigrain et al. 2004). In the case of the G2V star used in the simulations, the minimum data segment length for which the filtering was successful was ~ 60 days (last data segment in Figure 4.7), consistent with a rotation period of ~ 30 days for such a star.

It is therefore necessary to find other means of coping with this additional complexity. We have investigated two alternative approaches: one based on a least-squares generalisation of the Fourier filtering approach; the other based on a general purpose iteratively clipped non-linear filter. In both cases we use the preprocessing to attempt to remove the stellar signature, as much as possible, prior to invoking the transit search algorithms developed in Chapter 2.

4.3 Least-squares fitting

For a long run of regularly sampled data, a discrete Fourier transform asymptotically approaches a least-squares fit of individual sine and cosine components (see e.g. Bretthorst 1988). This naturally suggests an extension of the approach described in Section 4.2 to the case of irregularly sampled data. An analogous situation occurs in the generalisation of the periodogram method to Fourier estimation of periodicity; using generic least-squares sine curve fitting is a more flexible alternative (Brault & White 1971). This allows the case of gaps in the data, or more generally irregular sampling, to be dealt with in a consistent and simple manner.

The procedure is basically identical to that employed for the Wiener or matched filters described in the previous section, but the calculation of the Fourier transform, or power spectrum, of the received signal is replaced by an orthogonal decomposition of this signal into sine components whose amplitude, phase and zero-point are fitted by least-squares. Each of the components has the form:

$$\psi_k(t) = \alpha_k \sin(2\pi kt/T + \phi_k) \quad (4.6)$$

where T is the time range spanned by the data. The number of components to fit can be chosen such that the maximum frequency fitted is equal to some fraction of the Nyquist frequency, but for this one must define an equivalent sampling time δt . In the case of regular sampling with gaps, δt is simply the time sampling outside the gaps. In the case of irregularly sampled data the definition of δt is more open ended. However, provided that the sampling is close to regular, a good approximation will be the average time step between consecutive data points – keeping in mind that any significant gaps should be excluded from the calculation of this average. The potentially highest frequency component should then have frequency $\approx 1/(2\delta t)$, although in practice a much lower frequency cutoff for the components is all that is




Figure 4.7: Results of applying the matched filter independently to the 5 unbroken intervals of the light curve shown in Figure 4.6. Top panel: entire filtered light curve. Middle panel: 1st 30 days. Bottom panel: another 30 day section centred on the second transit (at 366.5 days). See text for an explanation.

Figure 4.8: Top panel: “power spectrum” (i.e. coefficients a_k versus frequency) obtained by the least-squares fitting method for the light curve shown in Figure 4.1. Middle panel: reconstructed light curve, obtained by summing over the fitted sine-curves up to a frequency of ~ 1.8 cycles/day. Bottom panel: 1st 30 days of the reconstructed light curve.

required.

Note that the first (zero-frequency) component is effectively the mean data value $\langle x(t) \rangle$ (which can be pre-estimated and removed in a robust way e.g. by taking a clipped median). The presence of gaps in the data provides us with a natural way of obtaining several independent estimates of $\langle X_{i_s}(w) \rangle$ by measuring it separately in each interval between gaps, or alternatively provides a natural boundary for doing independent light curve decompositions.

Figure 4.8 illustrates this least-squares fitting method, as applied to the light curve shown in Figure 4.1. The top panel shows the “power spectrum”, i.e. the coefficients a_k versus frequency, while the bottom two panels show the light curve re-

constructed by summing the fitted sine-curves. Note that high frequency variations are not reconstructed as only the first 2000 sine components were fitted (well below the Nyquist limit, but amply sufficient for the purposes of following the long timescale stellar variability).

The decomposition of the reference (transit) signal can usually be well approximated analytically. For example if a simple box-shaped transit of duration d is adopted as reference signal, the k^{th} coefficient is given by:

$$r_k = \frac{\sin(\pi k d / \delta t)}{\pi k d / \delta t} \quad (4.7)$$

However, this decomposition can also be performed in the same way as for the received data, for a reference signal of any given shape. The sets of coefficients a_k and r_k then define the filter h_k , which is equivalent to the filters of the previous section:

$$h_k = \frac{\langle |r_k|^2 \rangle}{\langle |a_k|^2 \rangle} \quad \text{or} \quad h_k = \frac{\langle |r_k| \rangle}{\langle |a_k| \rangle} \quad (4.8)$$

where the first expression corresponds to the standard Wiener filter, and the second to the pre-whitened matched filter used in Carpano et al. (2003).

Figure 4.9 illustrates this filtering method. Using the second expression in Equation (4.8) (equivalent to Equation 4.5), a 'matched filter' h_k is constructed from the coefficients a_k and r_k (the latter computed according to Equation 4.7). The filtered light curve, obtained by multiplying the a_k by h_k and reversing the 'transform', is shown in the middle panel, with a zoom on the first 30 days in the bottom panel.

Figure 4.10 shows the results of the matched filter constructed using the least-squares fitting method when the light curve contains gaps (as in Figure 4.6). The performance of the filter is generally not affected by the gaps, though artifacts near gap boundaries can sometimes be introduced.

The case of irregular sampling is not illustrated here, for practical reasons: if the sampling was allowed to vary, say, by $\pm 10\%$ of the normal sampling time in a random fashion, the effect is not visible in plots of such long light curves. In any case, we have found it to have negligible effect on the the least-squares filtering.

Note that the combination of the least-squares fitting method to construct power spectra with the pre-whitened matched filter (RHS of Equation 4.8) will, for conciseness, be referred to hereafter simply as 'the least-squares filtering method'.

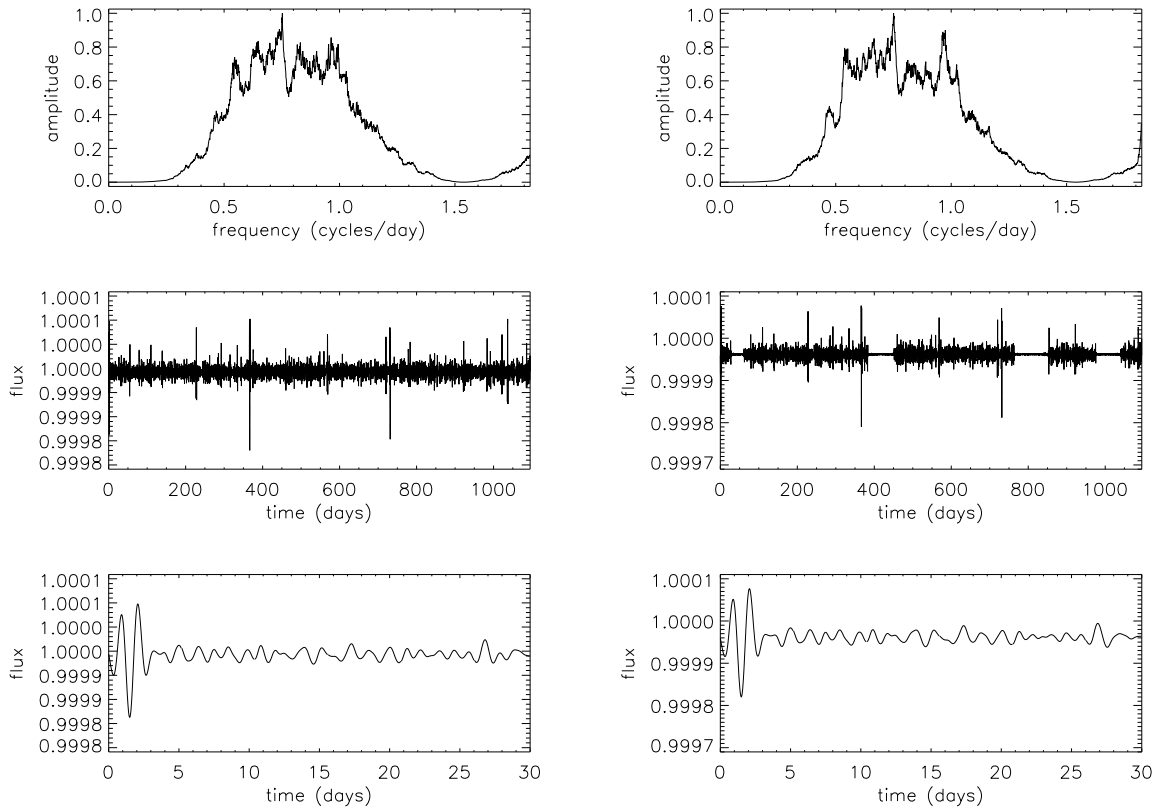


Figure 4.9: Top panel: equivalent matched filter constructed according to the 2nd expression in Equation (4.8), using the light curve shown in Figure 4.1 and a reference box-shaped transit of duration 0.65 day. Middle panel: filtered light curve. Bottom panel: *idem*, 1st 30 days, with a transit 1.5 day after the start.

Figure 4.10: As Figure 4.9, but the input light curve is that shown in Figure 4.6, with 4 significant data gaps.

4.4 Non-linear filtering

If the timescale of the transits is shorter than that of the dominant stellar variations, iterative non-linear time domain filters can pick out short timescale events. A standard median filter is a good starting point for this type of approach.

The data are first, if necessary, split into segments, using any significant gaps in temporal coverage to define the split points. These gaps, defined as missing or bad data points, or instances where two observations are separated in time by more than a certain duration, can be automatically detected.

Each segment of data is then iteratively filtered using a median filter of window ~ 2 to 3 times the transit duration, followed by a (small window) box-car filter




Figure 4.11: Light curve with data gaps filtered using the non-linear technique (black curve). The input data was the light curve shown in Figure 4.6. The window of the iterative median filter used was 3×0.65 days. The grey curve shows the same data with the residual noise level after filtering measured and artificial data with Gaussian distributed noise of the same standard deviation generated to fill the gaps. This illustrates the fact that, after non-linear filtering, the light curve (outside the transits) is well approximated by a constant level plus white noise.

to suppress level quantisation, which avoids excessive sensitivity to low-level dips (i.e. potential transits). Irregularities in sampling within a given segment are ignored, the filtering window being defined simply in terms of a number of data points. The difference between the filtered signal and the original is used to compute the (robust) MAD-estimated scatter (sigma) of the residuals. The original data segments are then k -sigma clipped (with $k = 3$) and the filtering repeated, with small gaps and subsequent clipped values flagged and ignored during the median filtering operation. The procedure converges after only a few iterations.

While applying this filter in the context of the COROT blind experiment (see Chapter 6), we found that the results were improved by the addition of a pre-filtering step, which consists in smoothing each interval using a median filter with a width of 2 or 3 data points before applying the iterative non-linear filter to construct the continuum (though the continuum is still subtracted from the original, unsmoothed data to give the filtered light curve). Because it smooths the sharp edges of the transits, this pre-filtering reduces the amount of transit signal (which we want to preserve) that is removed by the filter.

Break points and/or edges are dealt with using the standard technique of edge reflection to artificially construct temporary data extensions. This enables filtering to proceed out to the edges of all the data windows. The continuum obtained in this manner is then subtracted from the original to obtain the filtered light curve.

The main advantage of using a non-linear filter is that the exact shape of the transit is irrelevant and the only free parameter is the typical scale size of the duration of the transit events. The main drawback is that the temporal information in the

segments is essentially ignored. However, providing the sampling within segments is not grossly irregular this has little impact in practice. This filter is also relatively fast due to its simplicity: with the 512 MB RAM, 1.2 GHz processor laptop previously used to test the box-shaped transit finder (see Chapter 2), the running time for a transit duration of ~ 0.5 day is 4 seconds per light curve, about the same as the time required for the standard Wiener filter. The least-squares filtering method was significantly slower (requiring approximately 30 s when 1500 frequencies were fitted).

Figure 4.11 illustrates the non-linear filter as applied to the light curve with gaps shown in Figure 4.6. As with the indirect least-squares filtering, the high frequency noise remains, but this does not impede transit detection. Given the simplicity of this method and its good performance in the presence of data gaps, it appears to be the most promising, as long as the sampling remains relatively regular (if the sampling is significantly irregular, the least-squares filtering method, which takes the time of each observation into account directly, is likely to perform better).

4.5 Light curve characteristics after filtering

Several factors are to be taken into account when assessing the performance of the filters:

- How noisy is the filtered light curve?
- How Gaussian is the noise distribution in the filtered light curve?
- Has the transit shape and depth changed in the filtering process and how much?

Comparing Figures 4.10 and 4.11, one can readily see that, while both filtering methods suppress low frequency variations, the least-squares filtering method enhances any variations on the timescale of the reference transit (including the real transit) while suppressing high frequency noise. It also changes the shape of the transit significantly, as well as enhancing its contrast. On the other hand, the non-linear filter does not affect any variations on timescales shorter than two or three times the transit duration. This implies that, although the transit is more obvious to the naked eye after application of the least-squares filtering method, the noise distribution is close to Gaussian after non-linear filtering.

This is illustrated in a quantitative manner in Figure 4.12, which shows distributions of the deviations from the median before and after filtering (note that the transits were excluded from these distributions). In each case, a 1-D Gaussian was fitted to the distribution. The respective half-widths of the fitted Gaussian were $\sim 6 \times 10^{-4}$, 1×10^{-5} and 1×10^{-4} . The least-squares filtering method therefore removes more

noise. It also reduces the transit depth, though by a lesser factor: the approximate transit depth in each case was 3.2×10^{-4} , 1.4×10^{-4} and 3.2×10^{-4} . The ratio of the transit depth to the Gaussian width is thus higher after least-squares filtering than after non-linear filtering. However, the distribution after non-linear filtering is much closer to a Gaussian.

4.6 Discussion

The two filtering methods presented here share some advantages – both can be applied to data with gaps – but they also have different properties.

The least-squares filtering method is capable of making use of the time information in data with irregular sampling. It also allows a theoretically optimal filter (i.e. the Wiener or matched filter) to be combined with a pre-whitening filter, although from the point of view of detection, the matched filter is the main active component of any maximum likelihood-based detection algorithm. It is designed to be the method with the highest performance in terms of enhancement of the transit depth to noise ratio. As a by product of the filtering, the stellar signal can also be reconstructed. However, this is computationally intensive, particularly if one wishes to fit higher frequencies. Its performance also depends quite critically on concordance between the duration of the reference transit and that of any true transit. Its primary use in the context of space-based transit searches will therefore be the detailed investigation of borderline candidates, originally identified in light curves treated with the non-linear filter: there will be few of these so that the computing time requirements do not matter, and one will already have an idea of the approximate transit depth.

On the other hand, iterative non-linear filtering is simple to implement and fast, and produces nearly Gaussian residuals. This is a very important point because transit search algorithms in general – and in particular those developed in Chapter 2 – are optimised for white Gaussian noise. It is also less sensitive to the choice of reference transit duration, because it simply removes any signal on timescales longer than two or three times this duration, rather than applying a Fourier domain filter which has a complex structure over a wide range of frequencies. The non-linear filter is thus our filter of choice for space-based transit searches, where the time sampling is regular apart from the occasional data gaps. However, it ignores any local time information (except for the long gaps which are detected automatically). This means that its performance is likely to degrade if the sampling is seriously irregular, e.g. ground based transit searches, where one will have to resort to the least-squares filtering method.

Whatever the method used, there is a fundamental limit to what can be filtered

out. Stellar variability can only be filtered out if an orthogonal decomposition of the transit and stellar signal is possible, e.g. if the two signatures in the frequency domain do not overlap too much. Therefore, very rapidly rotating stars where the rotation period is close to the transit duration, or stars showing much more power than the Sun on timescales of minutes to hours (e.g. higher meso- or super-granulation) will be problematic targets – although perfectly periodic stellar signals, even if they have large amplitudes and periods close to transit timescales, are easily removed using e.g. the sine-fitting technique discussed in Chapter 6, Section 6.3.1.3.

Of course, these filters will not be used in isolation. The quantity we are really interested in is the performance of the transit search algorithms when applied to their output. This is the subject of the next chapter, where the filters are coupled with the box-shaped transit finder from Chapter 2 and applied to the output of the simulator from Chapter 3 for a wide range of input parameters, given the observational characteristics of the *Eddington*, COROT and *Kepler* missions. The same combination is also tested on COROT data simulated on the basis of inputs from a number of European groups as part of the blind experiment described in Chapter 6.




Figure 4.12: Distributions of the deviations from the median **a)** before filtering (in green), **b)** after applying the least-squares filtering method (in red) and **c)** after applying the non-linear filter (in blue). A Gaussian fit is shown in each case (black line). Obvious departures from Gaussianity were excluded from the fits (grey hashed regions). **c)** shows all three distributions, scaled and shifted so that each fit corresponds to a zero mean, unit variance, unit amplitude Gaussian.

Chapter 5

Impact of stellar micro-variability on transit detection from space

5.1 Introduction

This chapter summarises a number of simulations which were performed with the tools developed in the previous three Chapters. These simulations were run at different times during the period 2001–2004, so that both filtering and transit detection tools, and the mission parameters which enter into the simulations, change from section to section. The purpose of the simulations also varied, and this affects the choice of method in terms of parameter space coverage and level of detail and accuracy sought.

Section 5.2 describes simulations carried out in early 2003, in preparation for the 2nd *Eddington* workshop, which took place in April 2003. One of the main issues to be discussed at this workshop was the choice of target field for *Eddington*, and these simulations were aimed at identifying which types of planet-host stars would be favourable targets, and which would be too variable for transits of habitable planets to be detectable. At that time, the development of both variability filters and transit search algorithms was less advanced, in particular in terms of computational efficiency, so that full Monte Carlo simulations over a large grid of configurations were not feasible. A “quick look” approach was therefore adopted, in which the results from single realisations of each studied case are taken as indicative of the trends one is trying to identify, but no attempt is made to derive quantitative error rates.

Section 5.3 describes simulations carried out in late 2003 with the aim of assessing the performance of the box-shaped transit finder in combination with the non-linear filter, in the presence of stellar variability. Increased efficiency meant that Monte Carlo simulations were now feasible, but there was no need to repeat the

parameter space exploration of Section 5.2. Instead, a few targeted configurations, known to be close to the detectability limit, were used. A small number of similar simulations for the COROT case were also computed for the 5th COROT Week, which took place in December 2003, and are included in this section.

Photon noise was included (using a Gaussian approximation, which is appropriate in all the cases considered here) as expected for the missions concerned given their baseline design at the time of the simulations, and stellar micro-variability was included using the simulator developed in Chapter 3. No instrumental noise was considered. Transit signals were generated using the Universal Transit Modeller (UTM) of Deeg et al. (2001).

5.2 Simulations to identify the best target stars for *Eddington*

These simulations include:

- photon noise as expected for the *Eddington* baseline design as it stood in early 2003¹. The design has not changed since, but it differs from that used in older, photon noise simulations presented in Chapter 2;
- stellar variability light curves generated using the simulator presented in Chapter 3, but with an older version of the activity-variability relation (dashed line on Figure 3.6);
- transit light curves generated using UTM (Deeg et al. 2001).

For a given star-planet configuration, a single realisation of the light curve was generated. It was then filtered using the optimal filter of Carpano et al. (2003), which is also discussed in Section 4.2, and an early version of the box-shaped transit finder discussed in Chapter 2 was applied. This version was identical in its results, but computationally less optimised (the improvements to the implementation described in Section 2.2.3 had not yet been made). The planet was deemed detectable in that configuration if the peak in the detection statistic corresponded to the period and epoch of the input transit, and undetectable if not. In some cases, harmonics of the true period produced higher peaks than the period itself. These cases were also considered detectable.

¹The payload was constituted by a set of 4 identical wide field telescopes, with identical white light CCD cameras at the focal plane. In practice, this is equivalent to a single monolithic telescope with the same collecting area as the sum of the collecting areas of the 4 individual telescopes. The total collecting area of the baseline payload design was 0.764 m², and the field of view has a diameter of 6.7°. The CCD chips were from E2V, with the standard "broad band" response.

5.2.1 Which are the best target stars for *Eddington*?

The aim here is to identify, for example, those stars which are likely to be so variable that the detection of terrestrial planets orbiting them by *Eddington* (and by extension *Kepler*) will be seriously hindered, or on the contrary where the transits are easily recovered even in the presence of variability. This will be used to optimise both the choice of target field and the observing strategy, so that the range of apparent magnitudes containing most of the best target stars is well covered.

Light curves were therefore generated for a grid of star ages (0.625, 1.0, 2.0, 3.0 & 4.5 Gyr) and types (F5, F8, G0, G2, G5, G8, K0, K2 & K5). Planetary transits were added to the light curves for 1 & 3 R_{\oplus} and 1 R_{Jup} planets with periods of 30 days, 6 months, 1 year and 3 years, resulting in 37, 7, 3 and 1 transit(s) respectively. The light curves last 3 years and have a sampling of 1 hr. Two apparent magnitudes, $V = 13$ and $V = 15$, were used.

Note that the expected sampling rate is in fact closer to 10 min than 1 hr. However, this set of simulations was designed to rapidly explore the stellar parameter space to identify regions of interest. For this purpose, the light curves were generated using a longer integration time, thereby keeping them to a manageable size and maximising the contribution of stellar noise relative to photon noise on a given data point – the impact of photon noise having already been investigated previously. Later simulations, concentrating on the habitable zones of the more promising target stars, were made with 10 min sampling. The results are shown in Figure 5.1, using the notation described in Figure 5.2.

The main conclusions to be drawn from these results are:

- The detection of transits by planets with $R_{\text{pl}} \leq 3 R_{\oplus}$ around stars younger than 2.0 Gyr or earlier than G0 is significantly impaired, even for very high signal-to-noise ratio, and they are not good targets for the exo-planet search of *Eddington* or for *Kepler*.
- The small radius (hence increased transit depth) of K stars outweighs their relatively high variability levels, making them the best targets aside from effects not considered here, such as magnitude distribution and crowding. These will need to be assessed carefully. Whether this trend continues for M stars – recalling that they were not included in the simulations because of the significant number of fast M rotators present at Hyades age – is an open question (see Deeg 2004). Another complicating factor for these late-type stars is the need to include the effects of micro- and nano-flaring, an issue under investigation.
- Earth-sized planets are not detected correctly around G stars with 3 transits only. This is only an indicative result, but it demonstrates the need to increase

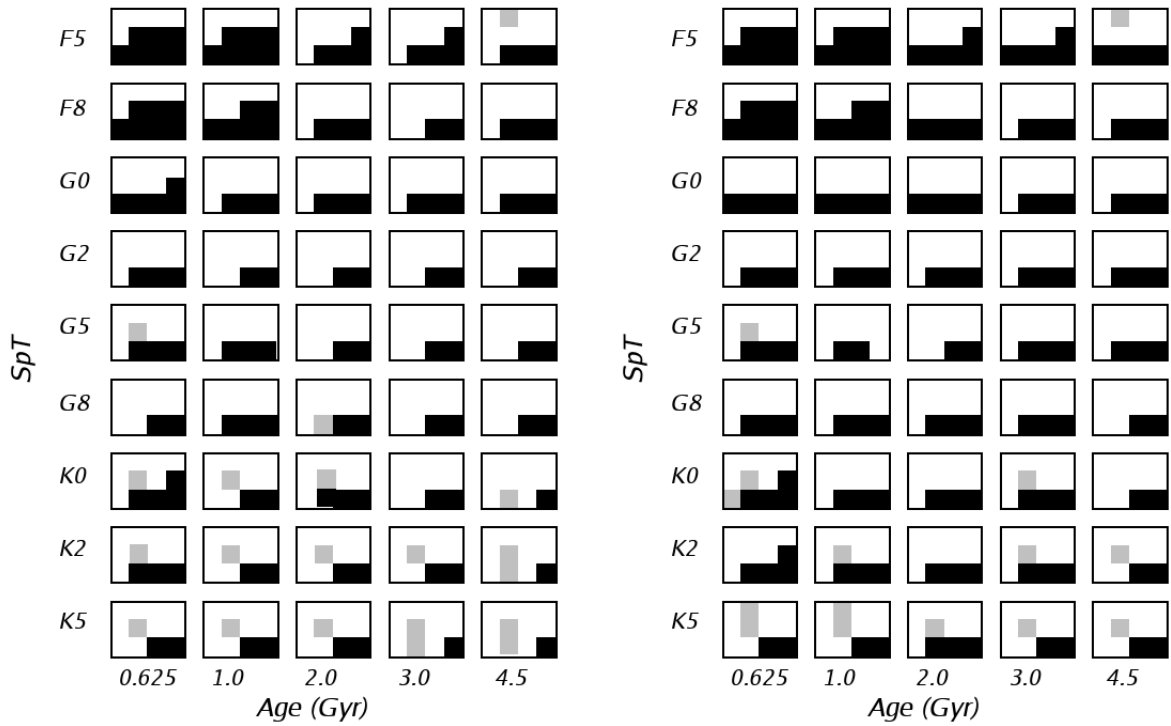


Figure 5.1: Results of the simulations with 1 hour sampling for $V = 13$ (left) and $V = 15$ (right). The notation used is detailed in Figure 5.2.




Figure 5.2: Notation used in Figure 5.1.




Figure 5.3: Results of the simulations with 10 min sampling for $V = 12$ (a) and $V = 13$ (b).

the transit energy or signal-to-noise ratio (our statistic S , cf. Section 2.2) for such systems, i.e. the square root of the number of in-transit points times the ratio of transit depth to per noise level per data point. As photon noise is not the limiting noise source at $V = 13$, increasing the collecting area would not significantly lower the noise level. Instead, one needs to increase the number of in-transit points, while keeping the noise level per data point constant. This can be achieved through longer light curve duration. Finer time sampling would increase the number of in-transit points, but also the photon noise per data point. It is thus expected to improve the results only for the bright, stellar noise dominated stars. The magnitude limit between the stellar noise and photon noise dominated regimes depends on the integration time, so that the optimum sampling rate is likely to be magnitude dependent.

- At $V = 15$ and with 1 hr sampling, photon noise has become the dominant factor for most stars.

The third point in the list above could have important consequences for the target field/star selection for *Kepler*. The current target field is centred on a galactic latitude of $\simeq 5^\circ$ in the region of Cygnus, and was chosen to maximise the number of stars in the field while being sufficiently far from the plane of the ecliptic to allow continuous monitoring (Borucki et al. 1997). Due to telemetry constraints, which limit the number of observed stars to $\simeq 100\,000$, and to crowding, which can have a very large impact in such high density regions, only stars with $V \leq 14$ are likely to be monitored in this field. However, with its larger collecting area (corresponding to a single aperture of 0.95 m), *Kepler*'s light curves will be star noise rather than photon noise dominated down to fainter magnitudes than in the case of *Eddington*. Given that, out of the spectral types tested here, K stars gave the best results, and that these stars tend to be fainter than earlier types, it may be desirable to choose a field at higher galactic latitude, combined with a deeper magnitude limit. It would likely be difficult to extend the magnitude limit while keeping the same target field due to increased crowding.

The final choice of target fields for both *Eddington* and *Kepler* will depend on many factors besides micro-variability and the change of stellar radii with spectral type, which are the only two effects taken into account in the present simulations. Other constraints are imposed by a number of factors. We have already discussed the impact of limited telemetry budgets, as well as that of the different apparent magnitude distributions of different stellar types, and the problems due to crowding in high density regions. Constraints arise from the luminosity class (i.e. radius) distribution of the stars in the field of view as well as from their spectral types (i.e. mass). Only fields with a large enough population of late-type (mid-F onwards) of dwarfs

are interesting for planet searches. Members of the COROT team, who have used spectrophotometry to determine the spectral type and luminosity class of fields in the potential COROT field, based on B , V , R , I , J , H & K information, have found that patchy reddening hampers the classification. Some candidate COROT fields in which the reddening was too spatially variable were excluded on that basis. Another important factor is the availability of ground-based facilities accessible to the scientific community involved with each mission, for further preparatory and follow-up observations. As the planet-finding target field will be close to one of the ecliptic poles, it will have a high declination and thus will be efficiently observable from either the southern or the northern hemisphere, but not both.

5.2.2 Minimum detectable transit in the presence of variability

Another question of interest is whether *Eddington* will really probe the habitable zone of the stars it targets, if micro-variability is taken into account. For this purpose, we have simulated light curves with increased sampling (10 min) for three 4.5 Gyr old stars (types G2, K0 & K5) orbited by planets with radii of 0.8, 1.0, 1.5 & 2.0 R_{\oplus} , with periods of 180, 360 & 365 days (corresponding to 7, 4 and 3 transits in the light curves respectively). The last two have very similar periods (i.e. virtually identical transit shapes and durations) but the extra transit in the 360 day case is due to the fact that the first transit occurs soon after the start of the observations. This is intended to test the effect of adding an extra year to the planet-search phase, in order to detect more transits. In white Gaussian noise, the dependence of performance on stellar and planet radius and apparent magnitude could be predicted analytically from the properties of the detection statistic employed by the transit search algorithm. However, in the presence of stellar variability, simulations are needed: one cannot assume that the noise is white and Gaussian even after filtering, particularly in the case of the more active stars.

The boundary between the stellar noise dominated and photon noise dominated regime, situated between $V = 13$ and $V = 15$ for 1 hour integration, will be approximately 1 mag brighter for 10 min integrations. In order to ensure that this second set of simulations still covered, at least in part, the stellar noise dominated regime, two values were used for the apparent magnitude of the stars: $V = 12$ and $V = 13$.

The results are shown in Figure 5.3. Generally speaking, they follow the expected trends. While the 0.8 R_{\oplus} planet is detected provided enough transits are present at $V = 12$, it is not detected around all but the smallest star at $V = 13$, as a result of the increased level of photon noise. The Earth-sized planet is only marginally detectable with 3 or 4 transits around a G2 star, even at $V = 12$. The fact that it is detected with 4 transits but not with 3 at $V = 12$, and is detected in both cases at

$V = 13$, highlights the need for full Monte Carlo simulations to obtain more reliable detectability estimates. It should serve as a reminder that the present work only aims to provide a global picture of the trends with star, planet and observational parameters, rather than quantitative results. Closer inspection of the light curve containing 4 Earth-sized transits around a G2 star at $V = 12$ showed that two of them were superposed on parts of the light curve where the noise was consistently positive, which impeded the detection. The rate of such coincidences can only be estimated from multiple realisations of the same system. To summarise, in the stellar noise dominated regime, the minimum reliably detectable radii for G2, K0 and K5 stars are roughly 1.5, 1.0 and $0.8 R_{\oplus}$ respectively with 3 or 4 transits, and $0.8 R_{\oplus}$ in all cases with 7 transits.

By comparison, ‘habitable’ planets are usually required to have radii in the range $0.8 R_{\oplus} \leq R_{\text{pl}} \leq 2.2 R_{\oplus}$. They are also required to lie in the habitable zone, given the following orbital distance ranges for 4.5 Gyr old stars: 0.5 to 1 AU ($0.8 M_{\odot}$, K2 star) and 0.9 to 1.3 AU ($1.0 M_{\odot}$, G2 star). Simplified calculations give the orbital period corresponding to the centre of the habitable zone, for the mass ranges of interest, as 1.2, 0.6 and 0.3 years for G2, K0 and K5 stars respectively, (see Section 1.2.3 for the rationale behind these limits).

The effect of the increased sampling rate is immediately visible. While a ‘true Earth analogue’ (Earth-sized planet orbiting a G2 star with a period of 1 year) may not be detected reliably, a good part of the habitable zone of the G2 star and all of that of the K0 and K5 stars are covered. This suggests that the primary goal of discovering and characterising extra-solar planets in the habitable zone will be achievable around intermediate age and old late-G and K stars with the current design. To push back these limits, modifications to the baseline design – such as the possible inclusion of colour information – have now been made. The fact that the detectability of both 1.0 and $0.8 R_{\oplus}$ planets increases significantly with increased number of transits, suggests that it may be desirable to increase the duration of the planet search stage, or to return to the planet search field for confirmation after a break (during which the asteroseismology programme, the other primary goal of *Eddington*, would be carried out).

5.2.3 Relevance to *Kepler*

As discussed in Sections 1.2.3.3 and 1.2.3.4, the photometric performances and time-sampling characteristics of *Eddington* and *Kepler* are very similar. The results of the present section can therefore be considered to apply to *Kepler* as well as *Eddington*. The epoch of 360 d period case in Section 5.2.2 was chosen so as to include 4 transits in the light curves, so that it corresponds roughly to a 1 yr period planet observed with *Kepler* (pointing duration 4 yr) while the 365 d period case corresponds to the

same system observed with *Eddington*.

5.3 Performance of the non-linear filter plus box shaped transit finder

The simulations presented in this Section were carried out using the micro-variability simulator, non-linear filter and box-shaped transit finder exactly as described in Chapters 3, 4 and 2 respectively.

5.3.1 Preliminary tests

To identify the zone of parameter space of interest, that is the limiting zone between detectable and undetectable transits, single realisations of individual star-planet configurations were tested one by one before embarking on more detailed Monte Carlo simulations.

The box-shaped transit search finder was first applied to the filtered light curve shown in Figure 4.11. The results are shown in Figure 5.4. The detection is unambiguous, and remains so for a $1.5 R_{\oplus}$ planet with otherwise identical parameters, though the detection is not successful for a $1 R_{\oplus}$ planet with only 3 transits (the star is a 4.5 Gyr old G2 dwarf in all cases). In the latter case, the 3 combined transits would have a total *SNR* of ~ 10 in white noise only, which should be detectable, but the combined *SNR* is only ~ 6 in the filtered light curves, which places it at the detectability limit. This is due to a combination of two effects. First, the filter does not completely remove the stellar variability (illustrated by a small departure from a Gaussian in the post-filtering noise distribution). Second, to mimic a realistic situation where the exact transit duration is not known before filtering, a guess duration of 50 data points, or ~ 8.3 hours, was used to filter the light curve, rather than the actual duration of ~ 13 hours. As a result some of the transit signal is filtered out, reducing the central transit depth by about one fifth. In a realistic situation, one would need to make a trade-off between computing time and sensitivity to choose the appropriate trial duration sampling.

5.3.2 Monte Carlo simulations for *Eddington*

The method employed was identical to that described in Section 2.1.3.1. The detection statistic (in this case the signal-to-noise ratio of the best candidate transit) is computed for N light curves with transits. All light curves have the same parameters, but different realisations of the noise and different epochs randomly drawn from a

Table 5.1: Light curve characteristics for each panel of Figure 5.5.

Panel	a)	b)	c)	d)	e)
Photon noise	✓	✓	✓	✓	✓
Stellar variability	×	✓	✓	✓	✓
Age (Gyr)	4.5	4.5	4.5	4.5	4.5
Spectral type	G2V	G2V	G2V	G2V	K5V
$R_{\text{pl}} (R_{\oplus})$	1.0	1.0	1.5	1.0	1.0
Period (yr)	1.0	1.0	1.0	0.5	1.0

uniform distribution (the epoch should not affect the detection process). The process is repeated for N transit-less light curves, which have noise characteristics identical to those of the light curves with transits. The chosen value of 100 for N is a compromise between accuracy and time constraints, and suffices to give a reasonable estimate of the performance of the method.

As the aim was to test the combined filtering and detection process, the light curves were subjected to the iterative nonlinear filter, before being forwarded to the transit detection algorithm. To avoid prohibitively time-consuming simulations, and thus to allow several star-planet configurations to be tested, a single transit duration value was used (corresponding roughly to the FWHM of the input transits). A detection threshold in terms of multi-transit SNR was automatically chosen to minimise the total number of false alarms and missed detection for each star-planet configuration.

Photon-noise only case As a benchmark against which to compare the performance in the presence of stellar variability, a photon-noise only case was tested first. Previous Monte Carlo simulations with white noise only involving this algorithm (see Section 2.2.4) were based on an older version of the *Eddington* design and did not include the non-linear filter.

The chosen configuration was a $1.0 R_{\oplus}$ planet orbiting a $V = 13$ G2V star with a period of 1 year, which as we have seen in Section 5.3.1 is at the detection limit when stellar noise is added. After a set of simulations was run for such a configuration, the maximum detection statistic from the noise only light curves was $SNR = 5.79$, while the minimum value from the light curves with transits was $SNR = 7.41$ (see Figure 5.5a). Any threshold in between would therefore allow the detection of all the transits where present, with no false alarms. As expected, the SNR limit of 6, established in Chapter 2, falls within the range of suitable thresholds.




Figure 5.4: Output of the box-shaped transit finder after non-linear filtering. Top: single transit SNR as a function of trial epoch. The signature of all three transits ($e = 1.5$, 366.5 & 731.5 days) is clearly visible. Middle: multiple transit detection statistic as a function of trial period. Bottom: idem, as a function of epoch at the optimal period of 365.0 days. The detected epoch (1.5 day) is correct. The x-axis for the top and bottom panels were shifted by 100 days for clarity.




Figure 5.5: Results of the performance evaluation for 5 star-planet configurations, as detailed in Table 5.1. Solid histograms: distributions of the maximum multi-transit SNR computed by the transit detection algorithm after non-linear filtering for 100 light curves containing transits. Dashed histograms: idem for 100 light curves containing stellar variability and photon noise only. Thick vertical lines: optimal detection threshold.

Photon noise and stellar variability

- **1.0 R_{\oplus} planet orbiting a G2V star**

This configuration is identical to that explored in the photon noise only case, but with stellar variability added. It is also similar to the case illustrated in Figures 4.1 to 5.4, but with a smaller planet. The results are shown in Figure 5.5b. The distributions of the detection statistics from the light curves with and without transits overlap almost entirely, i.e. the performance is poor. The threshold that minimises the sum of false alarms and missed detection leads to 56 of the first and 26 of the second.

Assuming that the sampling rate, light curve duration, and stellar apparent magnitude are fixed, there are three factors which should lead to better performance: a larger planet, a shorter orbital period (i.e. more transits) or a smaller star. Each of these options in turn is investigated below.

- **1.5 R_{\oplus} planet orbiting a G2V star**

The histograms are relatively well separated (see Figure 5.5c), with only a small overlap, so that the optimal threshold of $SNR = 7.85$ leads to one missed detection and no false alarms.

It is interesting to note the similarity between the results of this simulation and the requirements used for the design of the *Kepler* mission, which was to detect planets given a signal-to-noise ratio totalling at least 8 for at least three transits².

- **1.0 R_{\oplus} planet orbiting a G2V star with 6 transits**

The aim of this set of simulations was to investigate the effect of increasing the number of transits in the light curve by a factor of two by reducing the orbital period to 182 days. This is equivalent to increasing the overall duration of the observations. As expected, this leads to higher SNR values and hence better performance, with only 13 false alarms and 16 missed detections (see Figure 5.5d).

- **1.0 R_{\oplus} planet orbiting a K5V star**

A K5 star is smaller than a G2 star, leading to deeper transits, but also more active, leading to more stellar variability. Parallel studies (see Section 5.2.1) suggest that the former effect prevails over the latter, and that K or even M type stars might make better targets for space missions seeking to detect habitable planets than G stars, but these are based only on results from a few individual light curves, rather than Monte Carlo simulations. We saw in the previous chapter (Section 2.2.4) that such a configuration, in the presence of photon noise only could be detected down to $V = 14.5$ with the older (single telescope) *Eddington* design. Can it be detected at $V = 13$ with stellar variability added, using the (improved) photometric performance of the new, multiple telescope design?

²See www.kepler.arc.nasa.gov/sizes.html.




Figure 5.6: Results of the simulations for the COROT case. Top: $2.0 R_{\oplus}$ planet, bottom: $1.5 R_{\oplus}$ planet. The star is a 4.5 Gyr old G2V star with $V = 14$ and the planet's orbital period is 30 days. The light curves last 150 days with 16 min sampling. Solid histograms: distributions of the maximum multi-transit SNR computed by the transit detection algorithm after non-linear filtering for 100 light curves containing transits. Dashed histograms: idem for 100 light curves containing stellar variability and photon noise only. Vertical lines: optimal detection threshold.

The present tests confirm that K-stars make good targets: the separation between the with- and without transit distributions is wider (see Figure 5.5e) than in the previous case, though the best-threshold false alarm and missed detection rates remain high at 13 and 25 % respectively.

Note the higher SNR values for the transit-less light curves compared to the G2 case, which suggests the presence of more residual stellar variability after filtering, as would be expected.

5.3.3 COROT simulations

Two star-planet configurations were investigated in a similar manner for the case of COROT. COROT's exo-planet program is designed to detect relatively close-in (periods of days to a month or so) terrestrial planets around main-sequence stars with magnitudes in the range $12 \leq V \leq 15$. A solar-age G2V star with $V = 14$ was therefore used here, with a $2 R_{\oplus}$ or $1.5 R_{\oplus}$ planet in a 30 day orbit, leading to 5 transits lasting ~ 5 hours each. The first configuration is easily detected, as indicated by the top panel of Figure 5.6, but the detection is unreliable for the second.

These results are in approximate agreement with the design requirements of COROT, namely to detect 'big hot Earths', i.e. planets with radii above $2.2 R_{\oplus}$ in orbits lasting a month or less. Earth-sized planets might be detectable around smaller, brighter stars given even shorter orbital period (of order 10 days). No further simulations were carried out here, because the case of COROT is investigated in a much more detailed manner in Chapter 6.

5.4 Discussion

Light curves were simulated for single realisations of a large grid of star-planet configurations, including stellar micro-variability and photon noise as expected for *Eddington*. After optimal (or matched) filtering, the results of transit searches on these light curves suggest that stellar micro-variability, combined with the change in stellar radius with spectral type, will make the detection of terrestrial planets around F-type stars very difficult. On the other hand, K-stars appear to be promising candidates despite their high variability level, due to their small radius.

At $V = 12$ and with 10 min sampling, the smallest detectable planetary radii for 4.5 Gyr old G2, K0 and K5 stars, given a total of 3 or 4 transits in the light curves, were found to be 1.5, 1.0 and $0.8 R_{\oplus}$ respectively. This result was obtained in stellar noise rather than photon noise dominated light curves, and therefore also applies to lower apparent magnitudes or larger collecting areas. The magnitude limit beyond which photon noise would start to dominate, thereby increasing the minimum detectable radii for a given star and observing configuration, depends on the collecting area and sampling time, but the effects of increased photon noise are detected at $V = 13$ for 10 min sampling time and *Eddington*'s collecting area.

Full Monte Carlo simulations were run for a smaller number of configurations with the aim of assessing the performance of the box-shaped transit finder of Chapter 2 combined with the non-linear filter. From these it can be concluded that, if all or most of the stellar noise can be filtered out (so that the distribution of the detection statistic maxima in the noise-only light curves resembles that obtained for the photon-noise only case), any transit signal leading to a combined $SNR \geq 6$ can be detected. If residual stellar signal is present, the minimum detectable SNR is proportionally increased.

Most – though not all – of the stellar signal can be removed for G2 solar age stars, but a significant fraction of the stellar signal escapes the filtering for a solar age K5 star. In the latter case the detection performance could be improved by selecting transit candidates not only according to their combined SNR , but according to the ratio of combined to single-transit SNR , as the spurious SNR peaks due to stellar variability mostly arise from single events. The filter also does remove a small portion of the transit signal. These Monte Carlo simulations confirm that a $1.0 R_{\oplus}$ planet orbiting a solar age, $V = 13$ G2 star will be difficult to detect for *Eddington*. Improvements on this result can be sought along several directions: improved filtering, better candidate selection, and increases in the telescope collecting area (to reduce photon noise) and/or in the duration of the mission (cf. the 4 year mission lifetime for *Kepler*).

Full Monte Carlo simulations, using the photometric performance of the new *Eddington* design, combined with more recent versions of the filtering and transit

detection algorithms, confirm these results (though over a much narrower range of parameter space). They also provide a testbed for the performance of the new algorithms, highlighting for example the presence of residual stellar micro-variability after filtering for more active stars. This shows that light curve filtering is an area where more progress could potentially be made. The results of similar Monte Carlo simulations for COROT confirmed the detectability of 'big hot Earths' with that mission. The case of COROT is explored in more detail in Chapter 6.

Chapter 6

The COROT blind experiment

6.1 Introduction

This experiment, which took place between January and May 2004, was organised by a network of research groups across Europe who are involved in the preparation of the Franco-European transit search mission COROT, in particular its exo-planet search program.

At the 5th COROT week in December 2003, the need to test the capabilities of various detrending and transit detection algorithms on realistic light curves emerged. The satellite and instrument design phases were complete, and the instrument model was at a sufficiently advanced stage to provide realistic simulated light curves at least in white-light. An experiment, carried out in a concerted way between the different groups involved in the preparation of the data analysis, would minimise duplication of work while highlighting any areas where difficulties arose, and test some of the details of the observing strategy. It was therefore agreed to carry out such an experiment within the COROT Exo-planet Working Group (EWG), the results of which were to be announced at the 6th COROT Week in May 2004. The light curves from this experiment would then be made available to the wider transit search community. This was foreseen to be the first in a series of such experiments, subsequent generations drawing on the lessons from the first.

A collaborative publication in which all the participants will describe their contributions is currently in preparation (Moutou et al. 2004a). The purpose and organisation of the experiment and the generation of the light curves (Sections 6.1.1 to 6.2), whose descriptions are mainly drawn from the draft of this publication, are summarised here in order to set the stage for the transit search I performed in collaboration with M. Irwin (Section 6.3.1). A brief description of the methods used by the other participants is also given, though I have only limited information on the algorithms some of the groups used at this stage. The lessons learnt from the exercise are

discussed in Section 6.4, and future prospects in Section 6.5

6.1.1 Purpose of the blind experiment

Rather than a competition between the groups involved, the experiment was designed to identify the algorithm(s) most suited to the COROT data, and to estimate preliminary detectability limits in the presence of all the noise sources currently foreseen. It was designed along a 'hare and hound' scenario: light curves containing noise and stellar variability were built by a single 'game master' from components provided by various contributors, and transits and mimics of stellar origin were inserted. The final set of 1001 light curves – their contents known only to the game master – was then made available to the members of the EWG, who attempted to detrend the light curves and detect the transits by the methods of their choosing.

From my point of view, this was a golden opportunity to put the tools developed in Chapters 2 and 4 to the test and to compare their performance to that of others. It was also an opportunity to apply the simulator described in Chapter 3, which was designed exactly for this type of purpose, for the benefit of the wider exo-planet community.

6.1.2 Participants

The light curves included instrumental and photon noise simulated by D. Blouin from the Laboratoire d'Astrophysique de Marseille (LAM), stellar micro-variability simulated independently by A. Lanza from the Osservatorio Astrofisico di Catania and myself, planetary and stellar eclipses and variable star light curves simulated or collected from the literature by C. Moutou and F. Pont (LAM). The different components were combined to form light curves by C. Moutou, who coordinated the experiment.

Although any members of the EWG were invited to take part in the detection process, only five groups were able to send their results to C. Moutou by the deadline of May 1st: a team from the Institute of Planetary Research in Berlin-Aldershof, led by H. Rauer, two independent teams from LAM (P. Guterman and V. Guis), one team from Geneva Observatory led by S. Zucker and the IoA team composed of M. Irwin and myself.

6.2 The simulated light curves

6.2.1 General characteristics

The simulated light curves correspond to a COROT long run (lasting 150 days with 8.5 min sampling), and to the characteristics of the exo-planet field (see Section 1.2.3.2 for details of the COROT design and the distinctions between asteroseismology and exo-planet program). All light curves were simulated as though they had arisen from the same observation run, i.e. they have a common set of observation times and common time-dependent systematics.

Only light curves in white light were generated, as neither the COROT instrument model nor the stellar micro-variability simulating tools were at the time capable of producing coloured light curves. Work is now underway to upgrade the instrument model, and Lanza et al. (2004) have developed a technique for including colour information in stellar light curves. I am also planning to work on this question in the near future.

6.2.2 Instrumental and photon noise

The instrument model works by taking as input a number of simulated PSFs and a list of star positions and magnitudes, generating images by addition of the appropriate noise sources, computing appropriate masks, performing aperture photometry within each mask (applying the appropriate noise sources again), and then correcting for the known noise sources as well as we expect to be able to correct them in the real data.

25 PSFs were used, corresponding to stellar effective temperatures between 4500 and 6750 K and magnitudes between 12 and 16. Photon noise was added following Poisson statistics. A flat-field non-uniformity of $\sim 1\%$ was introduced. The read-out noise is $10 \text{ e}^- \text{ pixel}^{-1}$ and the jitter noise is negligible in white light. Zodiacal light is uniform over the CCD and constant along the orbit (with a value of $12 \text{ e}^- \text{ pixel}^{-1} \text{ s}^{-1}$), and is thus fully corrected for, resulting only in additional photon noise. On the other hand, the Earth scattered light is more difficult to correct, because it varies over the CCD, as well as along the orbit. This noise source was thus corrected only to first order, in a deliberate attempt to simulate a worst-case scenario. A scattered light variation component with a maximum amplitude of $1 \text{ e}^- \text{ pixel}^{-1} \text{ s}^{-1}$ was simulated and added to each light curve. A correction was then computed by multiplying the scattered light component by a factor between 0.9 (underestimation) and 1.1 (overestimation), and applied to the light curve, leaving residuals of $\pm 10\%$. As illustrated by Figure 6.1, residual scattered light is the dominant signal in




Figure 6.1: Example of an instrumental light curve before (top) and after the partial correction of the scattered light variations, which are either underestimated (middle) or overestimated (bottom). The sharp peaks in the upper panel are due to the South Atlantic Anomaly crossing, and correspond to gaps in the lower two panels. Plot provided by C. Moutou. .

the COROT instrumental light curves.

The large numbers of impacts by high energy particles as COROT passes through the South Atlantic Anomaly (SAA) makes the exposures taken during the SAA crossings unusable. This leads to semi-periodic gaps in all light curves lasting ~ 10 min, occurring ~ 10 times per day. The gaps were not exactly synchronous in all the simulated light curves because an upper flux threshold was applied to each light curve individually to determine where the gaps would be.

6.2.3 Stellar micro-variability

Two sets of light curves containing stellar micro-variability were simulated independently.

One set was produced by myself, using the model described in Chapter 3. It contained 45 light curves covering a grid of 9 spectral types (F5, F8, G0, G2, G5, G8, K0, K2 & K5) and 5 ages (0.625, 1.0, 2.0, 3.0 & 4.5 Gyr). The light curves are illustrated in Figure 6.2.

The second set was produced by A. Lanza and collaborators from Catania Observatory. They have developed a method to reproduce the Total Solar Irradiance (TSI) variations, as observed by SoHO/VIRGO, based on the differential rotation of 3




Figure 6.2: The 45 stellar micro-variability light curves simulated for the COROT blind experiment with the model described in Chapter 3.

active regions, composed of faculae and spots, whose sizes and positions on the star's surface are allowed to vary, but whose contrast factor is fixed (Lanza et al. 2003). Fits are made to 14 d portions of the TSI light curve spaced by 7 d, and linear interpolation between successive best fits is used to minimise discontinuities between fits. To produce light curves for stars other than the Sun, Lanza et al. apply a scaling factor f to the area of the active regions, and the rotation period is modified to reproduce that of the star being modelled. The factor f is deduced from an empirical scaling law established by Messina et al. (2001, 2003), relating the amplitude of

V–band variations to the active region area coverage. For stars with rotation periods longer than 12 d, for which no optical variability information is available, f was taken to be between 1.5 and 6 for F5V to K5V stars respectively. In order to account for shorter timescale variations, which cannot be explained by rotational modulation, they also inserted an additional component in their light curves, calculated by scaling the residuals from their fit to the solar light curves by a factor $3f$. This higher scaling factor was chosen arbitrarily to mimic a worst-case scenario. Ten light curves were thus generated, with spectral types G5, G0 and G8 and rotation periods 3, 10 and 20 d.

Some light curves in the resulting set (totalling 55) were then duplicated to reproduce the expected distribution of spectral types (40% F dwarfs, 40% G dwarfs, 20% K dwarfs). In order to obtain a sufficient number of non-identical stellar micro-variability light curves, C. Moutou then applied an arbitrary amplitude scaling factor between 0.5 and 2.0, and a time scaling factor between 1.0 and 1.2, as well as arbitrary time shifts. The contributions by A. Lanza and myself were treated in the same way, each light curve being assigned a stellar micro-variability component from one set or the other at random.

6.2.4 Transit signals

Twenty planetary transits were simulated by C. Moutou using the Universal Transit Modeller (UTM) of Deeg (1999), using limb-darkening coefficients calculated from ATLAS9 models and the COROT bandpasses (Barban, priv. comm.).

The choice of transit signals was aimed at sampling a variety of cases and testing detectability limits by including a large number of shallow transits in light curves of varying noise level. The planet radii used were in the range $1.6 R_{\oplus} \leq R_{\text{pl}} \leq 1.3 R_{\text{J}}$, and the periods in the range $4.8 \leq P \leq 88.4$ d. One double planet system was included. The process by which the transit signals were inserted in the light curve was not entirely random, in that care was taken to insert the deepest transits into the light curves of faint/active stars.

This part of the experiment was not designed to be realistic – 2% of the light curves contained planetary transits, which is approximately an order of magnitude larger than what is expected (Bordé et al. 2003), and the parameters of the planets do not reproduce the observed distributions. Instead, the number was chosen to allow variety while keeping the size of the dataset small for this first generation exercise, and the parameters to test the algorithms in limiting cases which were expected to be problematic.

6.2.5 Stellar mimics and variables

Ten low depth stellar eclipses were simulated. Six were grazing binaries, simulated using the algorithm of Mandel & Agol (2002) and the 'Nightfall' software of Wichman (1998). Four were diluted background binaries, simulated with UTM. In addition, one triple eclipsing system was simulated, also with UTM. Again, the parameters were chosen to explore a variety of cases rather than follow observed distributions. The background eclipsing binaries were 3 to 5 mag fainter than the principal target in the mask. As well as the V- (grazing) or U-shaped (annular) primary eclipses, the eclipsing binary light curves include secondary eclipses and, when applicable, out-of-eclipse modulation due to ellipsoidal deformation of the primary.

In addition, the light curves of stellar variables were taken from the literature and the archive of the AAVSO (American Association of Variable Star Observers) and inserted. These included a low amplitude δ -Scuti, a classical Cepheid, a β Cephei, the semi-regular variable Z Uma and the irregular Z Cam.

6.2.6 Background pollution

The relatively large size of the COROT masks imply that a number of faint background stars are expected to fall in the mask of most COROT targets. The variability of these background stars will be superimposed with that of the main target and may significantly contribute to the overall result. To simulate this effect, the contribution of one background star – computed in the same way as the principal target's, but with a different stellar variability component – was added to each light curve. The background star's magnitude followed a $2^{\Delta m}$ distribution with $0 \leq \Delta m \leq 6$, thus including stars up to 22nd magnitude (see note regarding magnitude distribution in the next Subsection).

6.2.7 The final set of 1001 light curves

Of the final 1001 light curves, 964 have no stellar or planetary transit signal (apart from the stellar micro-variability component, which all light curves have). The temperature of the simulated stars follows a uniform distribution between 4000 and 6000 K, approximating a magnitude-limited (rather than volume limited) sample. The magnitudes were drawn from a 2^m distribution, approximating an isotropic spatial distribution. This corresponds to the distribution measured from observations of the COROT fields obtained in the context of the ground-based preparatory observations program, and is in agreement with the star counts expected from the Besançon model of the Milky Way (Robin et al. 2003, 2004). All light curves have a temporal sampling rate of once per 8.5 min, except two which have 32 s sampling (this 'oversampling' mode will be




Figure 6.3: Sine-curve fitting to identify the period of the scattered light variations. Top panel: test statistic versus trial frequency. The test statistic is the reduced χ^2 of the residuals from the best fitting-sine curve at each frequency. The best-fit frequency is given by the minimum value. Bottom panel: window function amplitude spectrum.

possible for a small number of COROT light curves). The light curves that do contain planetary transits and stellar events are summarised in Tables 6.3 and 6.4.

These light curves, in units of photon counts, were supplied to the detection teams with no information about the way in which they were generated. When real COROT light curves come to be analysed, some information regarding magnitude, spectral type, luminosity class and contamination by neighbours *will* be available. However, this information, while useful at the transit candidate characterisation stage, is not fundamental at the transit detection stage. The light curves as received by the detection teams are referred to hereafter as level 0 light curves.

6.3 The transit search

The detection teams were asked to return a list of candidate transits with their parameters, as well as a list of other events identified in the course of the transit search. The results were to be sent by email to C. Moutou by May 1st 2004.

6.3.1 Method used by the IoA team (team 5)

The method proceeds along successive steps outlined in Figure 6.4 and detailed below.

6.3.1.1 Removal of residual scattered light variations

Upon receipt of the level 0 light curves, the first noticeable feature were the short period variations due to residual Earth-scattered light, which were positive or negative depending on the light curves, and the SAA gaps.

The exact period of the residual scattered light variations was determined from least-squares fitting of sine curves in the period range 0.065 to 0.075 d (see Figure 6.3). In the spirit of total ignorance of the light curve contents, this period determination was carried out individually for each light curve, though as expected the best fit period was always found to be very close to the satellite orbital period of 0.0705 d, the variations being within the errors on the fitted period. In each case, the light curve was then phase folded at the best-fit period. The phase-folded light curve was then smoothed, using a box-car filter with a width of 501 data points, and the smoothed version was subtracted from the phase-folded light curve, before remapping to the original time base and adding the original median level. The resulting light curves are referred to hereafter as level 1 light curves. An example of the scattered light removal procedure is illustrated in Figure 6.5. The effect of the removal on a section of the same light curve is also shown in the top and middle panels of Figure 6.7.

6.3.1.2 Removal of other trends and 'glitches' common to all light curves

During the scattered light removal stage, a number of 'glitches' and trends common to all light curves became apparent. This is consistent with the fact that all light curves were simulated as though they shared the same set of observation times. For example, two of the scattered light variation cycles appeared to be 'missing' approximately 78 d after the start of the level 0 light curves. The subtraction of the smoothed phase-folded light curve therefore induced two artificial 'bumps' or 'dips' at that particular location in the level 1 light curves.

In order to remove all systematic trends common to all light curves, a Principal Component Analysis (PCA) approach, based on a variant of Gram-Schmidt orthogonalisation, was initially considered. However, it involved the manipulation of 1000 by 25 056 element arrays (each light curve contains 25 056 data points), which would be computationally rather expensive. Additionally, the mean of all the light curves appears to contain most of the systematic trends and glitches which needed to be removed. The following, simplified procedure was therefore adopted:

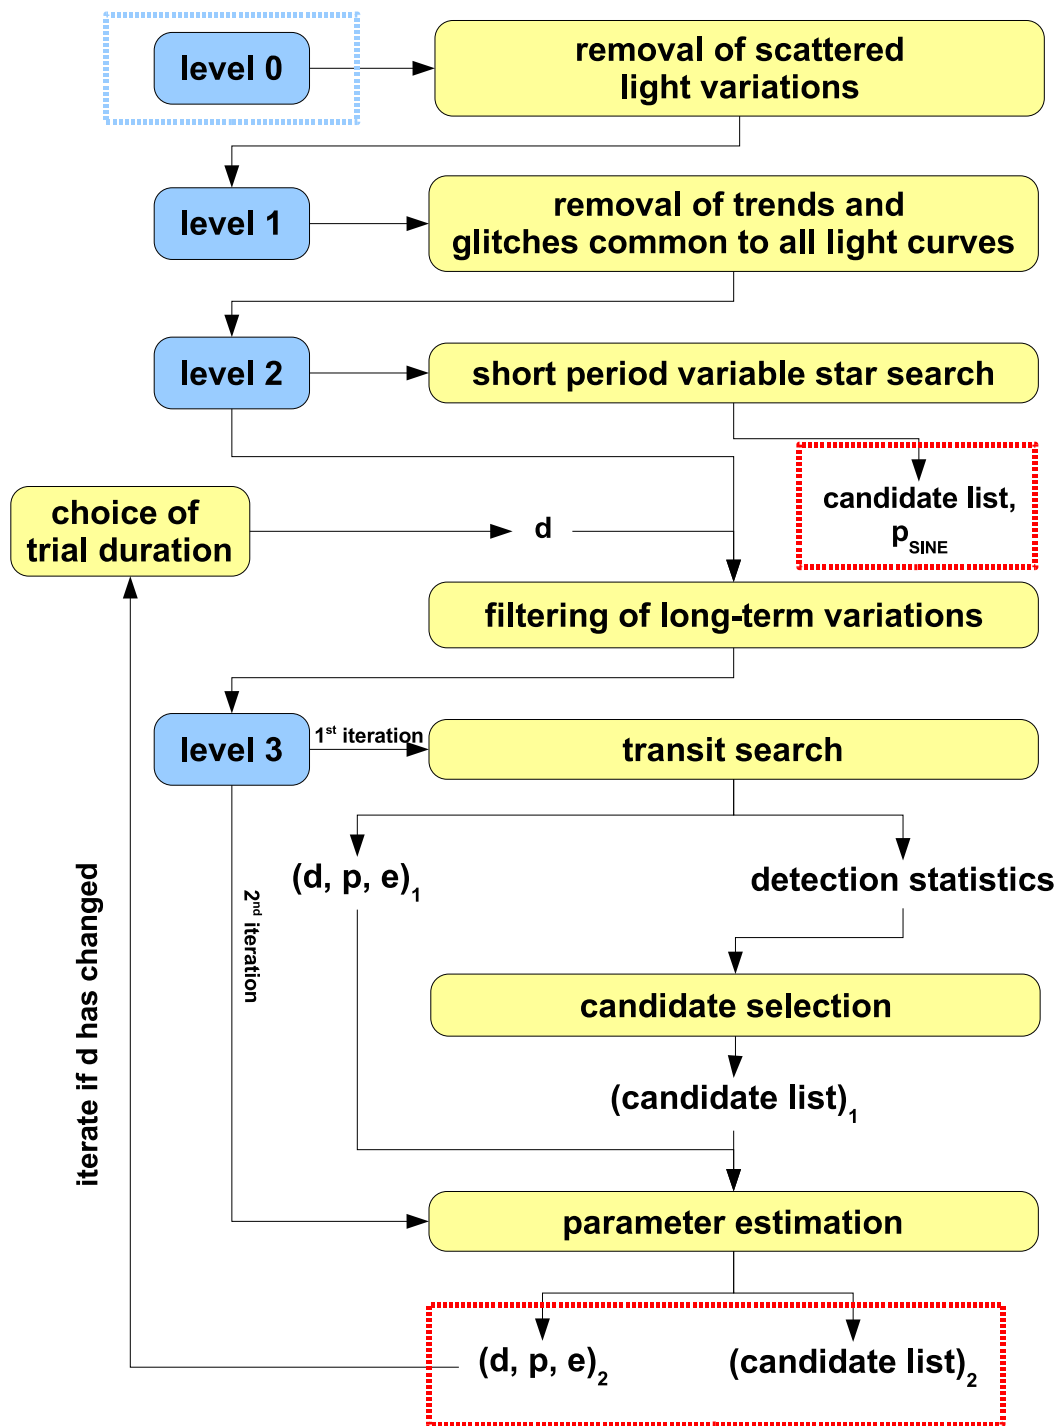


Figure 6.4: Main steps of the method employed by the IoA team, starting from the light curves as received from C. Moutou (blue dashed box) and ending with the short period variable and planetary transit candidate lists (red dashed boxes).




Figure 6.5: Removal of the scattered light variations. Top panel: Input (level 0, black) and output (level 1, red) light curves. Bottom panel: Light curve phase-folded at the best-fit period (black) and smoothed (red). The level 1 light curve (red curve in top panel) is obtained by subtracting the red from the black curve in the bottom panel and remapping to the original time base.

- The median count level and scatter σ of each level 1 light curve were computed. Note that a robust scatter estimate, $\sigma = 1.48 \times \text{MAD}$, where the MAD is the median of the absolute deviations from the median, was used throughout the present analysis. This estimate approximates the standard deviation in the case of white Gaussian noise, but is less sensitive to outliers (Hoaglin et al. 1983).
- Each level 1 light curve was median subtracted and scaled to unit scatter.
- A ‘scaled common average’ light curve was then computed by taking the median of all 1001 median subtracted, scaled light curves at each time point (NB: light curves 1000 and 1001 were simply rebinned to 8.5 min sampling and treated throughout the analysis identically to the other light curves), and scaling the result to have unit variance. An array of scatter values for each time point – the ‘scaled common error’ array – was also computed. Excerpts from the scaled common average and error arrays are shown in Figure 6.6.
- The scaled common average light curve was then ‘fitted’ to each level 1 light curve by finding the coefficient k which minimises $\sum_{i=1}^N (d_i - k m_i)^2$, where N is




Figure 6.6: Excerpts from the scaled common average light curve (top) and the scaled common error array (bottom) in the time range 76 to 88 days. The 2σ cut-off (red dashed line in bottom panel) is used to exclude common ‘glitches’ from all light curves.

the number of data points and d_i and m_i are the i^{th} points in the level 1 and scaled common average light curves respectively. If m_i is normalised such that $\sum_{i=1}^N m_i^2 = 1$ then this simplifies to $k = \sum_{i=1}^N d_i m_i$. The scaled common average light curve, multiplied by the factor k , was then subtracted from the level 1 light curve.

- The scaled common error array was also used to exclude common ‘glitches’: a 2σ cut-off was applied and time points corresponding to values of the scaled common error array above the cut-off were excluded from all light curves.
- The resulting light curves, free of common systematic trends (to first order at least) and glitches, are referred to hereafter as level 2 light curves.

Excerpts from an example of level 0, level 1 and level 2 light curve are shown in Figure 6.7, illustrating the removal of both residual scattered light variations and glitches.

6.3.1.3 Identification of short-period variable stars

At that stage, the least-squares sine-curve fitting program was run on all level 2 light curves with a range of trial periods p of 0.1 to 5 d (this range was designed to probe




Figure 6.7: Excerpts from an example of level 0 (top), level 1 (middle) and level 2 (bottom) light curves in the time range 76 to 88 days. Note the residual scattered light variations in the level 0 light curve and the ‘glitches’ at ~ 79 d in the level 1 light curve.

periods shorter than those explored through the transit search). This process took approximately 12 hours on a 2 GHz Dual-AMD Linux PC. In seven cases, the subtraction of the best-fit sine component led to a significant reduction in the reduced χ^2 of the light curve. These are summarised in Table 6.1. For these seven cases, the search was iterated to find additional periodicities with an extended period search range.

6.3.1.4 Variability filtering and transit search

The iterative non-linear filter described in Chapter 4 was applied to each light curve prior to running the box-shaped transit finder described in Chapter 2. An alternative in terms of filtering would have been to use the optimal filter (after estimating the power spectrum by least-squares fitting of sine-curves), an option also described in Chapter 4. However, the good results obtained in tests of the iterative non-linear filter and its low computational requirements made it the filter of choice, at least for the initial detection stage. The filtered light curves are referred to as level 3 light curves.

We found that pre-smoothing, using a median filter with a width of 7 data points followed by a linear filter with a width of 3 data points, before running the iterative

Table 6.1: Short-period variables identified by sine-curve fitting.

ID	p (days)	Comment
249	3.903327 44.647	a
259	0.706374 0.710492	a
384	1.781245	
553	0.461317	
599	0.937934	b
809	1.601433	c
915	1.451470	c

Comments: (a) second period found after subtracting primary periodic component, (b) eclipsing binary with orbital period p (secondary eclipses visible in phase-folded light curve), (c) eclipsing binary with orbital period $2p$ (marked difference between odd- and even-numbered eclipses).

non-linear filter to construct the continuum to be subtracted from the data¹, improved the results of the filtering, minimising the amount of transit signal removed by the filter for a given trial duration d . The filtering process is illustrated in Figure 6.8.

Three trial transit durations were used: 25, 50 and 100 time steps, corresponding to 3.6, 7.1 and 14.2 hrs respectively. The last value is above the range of durations expected for most planetary transits, but was used as a check, as previous experience had shown that false alarms due to residual stellar micro-variability tend to lead to detections at the longest trial transit duration.

Performing the transit search on the 1001 light curves at each transit duration took ~ 1 hr on a 2 GHz Dual-AMD Linux PC. For each light curve and trial duration, the following information was recorded: the best epoch e_S and detection statistic S_S for a single transit (maximum from the top panel of Figure 6.9), and the best period p , epoch e_M and detection statistic S_M for multiple transits (maxima from the middle and bottom panels respectively of Figure 6.9).

If the light curves had contained only transits and white Gaussian noise, a simple threshold in S_M (which is equal the signal-to-noise ratio of the combined transits) would have led to a given confidence level. Provided the filtering process was successful in all light curves, a simple threshold should also have been sufficient.

However, plots of S_M versus S_S for all light curves at a given transit duration (see e.g. Figure 6.10) show a large number of points with high S_S and $S_M \simeq S_S$. These correspond to light curves with a single transit-like event. Examination of these light curves showed no convincing transit candidates. It is also interesting to note that the tail increases in number and extends towards higher S_S and S_M for longer trial durations, and that most points in the tail corresponds to cases where the scatter

¹See details of the iterative non-linear filter in Chapter 4 for more detailed explanation.




Figure 6.8: Example of iterative non-linear filtering with a trial duration of 3.5 hrs for light curve ID 34. Left column: entire light curve. Right column: detail. Top row: level 2 light curve. 2nd row: pre-smoothed light curve. 3rd row: continuum. Bottom row: level 3 (filtered) light curve. This light curve contained transits with a depth of 0.07% and a period of 5.52 d. They are clearly visible in the 2nd row, and are present in the 4th, though buried in the (now close to white) noise.

in the level 3 light curve was significantly higher than expected from simple photon counting statistics. This suggests that these points are due to incompletely filtered stellar micro-variability. Variations of stellar origin are mostly on timescales longer than a transit, so that shorter trial (hence filter) durations give better results. That was expected, but the large number of light curves still significantly contaminated by residual variability even at the shortest trial duration was not. A possible explanation for this is discussed in Section 6.4.3.1. For the purposes of the present exercise, it was necessary to devise a method to discriminate between real (periodic) transit candidates and these events.

A transit signal repeated N_{tr} times in a noise-free light curve would lead to N_{tr} identical peaks in the distribution of single transit signal as a function of trial epoch, and $S_M = \sqrt{N_{tr}} \times S_S$. However, noise in the light curves and the discrete sampling in trial parameter space lead to variations in the single transit signal-to-noise ratio values at each trial epoch corresponding to a real transit, while S_S is taken to be the maximum of these values. One therefore expects S_M to be slightly smaller than $\sqrt{N_{tr}} \times S_S$. Any light curve for which $S_M \geq a + b \times S_S$ were therefore selected as potentially containing two or more transits, with $b = 1.3$ (slightly smaller than $\sqrt{2}$). The constant a , taken to be 1.0, was designed to make the threshold more stringent for less significant events. All cases below a similar threshold line with $b = 1.4$ were flagged as low confidence level candidates. Any light curves with $S_S \geq 20$ were also




Figure 6.9: Example of the transit detections statistic distributions as a function of epoch e for a single event (top), as a function of period p for a multiple event (middle) and as a function of epoch at the best period (bottom)

selected as potentially containing single but significant transits (e.g. single transit of a long-period Jupiter-like planet). This automatic selection process yielded three initial candidate lists, one for each trial duration.

The level 2 and 3 light curves were then examined by eye in the transit vicinity, and obviously spurious detections eliminated from the list. The three lists were also merged at this stage to create a single candidate list.

6.3.1.5 Basic transit parameter estimation

Once the candidate list was 'weeded out' and merged, the basic transit parameters were estimated in the following way:

- The starting point was the level 3 light curve corresponding to the trial duration closest to the true duration as estimated by eye.
- A more accurate duration was computed by estimating the 'full width at half-minimum' of the transits as follows:
 - phase-folding the light curve at the best trial period;
 - rebinning it into bins lasting approximately one fifth of the current duration




Figure 6.10: Multiple versus single transit detection statistic for trial durations of 3.5, 7.1 and 14.2 hrs (left, centre and right panels respectively). Gray symbols: measured scatter in filtered (level 3) light curves exceeds the expected scatter from photon statistics only.

estimate;

- finding the minimum in the rebinned phase-folded light curve;
 - measuring the depth of the transit from that minimum value;
 - finding the half-minimum points by moving outward from the minimum one point at a time;
 - the new duration estimate is then given by the interval between the half-minimum points,
- A new level 3 light curve was generated by applying the non-linear filter to the level 2 light curve with the newly estimated duration.
 - The duration was estimated again, also yielding a transit depth estimate and a refined epoch estimate.

A small number of very shallow transit candidates were eliminated at this last stage as convergence in the duration estimate failed to be reached.

The results, as summarised in Table 6.2, were then send to C. Moutou. The phase-folded level 3 light curves of all transit candidates (corresponding to the last iteration of the transit parameter estimation procedure) are shown an Appendix at the end of this chapter.

Table 6.2: Transit candidates and their parameters. NB: in cases where there was evidence (from visual inspection of the light curve) that the true period was a harmonic of the detected period, the suspected true period is given.

ID	N_{tr}	$\Delta F/F$ (%)	e (days)	p (days)	d (days)	Confidence	Comment
31	6	0.27	7.34	24.724	0.141	high	a
34	27	0.06	3.662	5.528	0.107	high	
85	6	0.9	0.175	26.427	0.175	high	
117	2	1.78	29.751	69.841	0.14	low	b
168	13	0.27	4.803	11.513	0.035	high	
207	2	1.5	41.45	88.486	0.279	high	c, d
249	48	28.06	2.729	3.903	0.105	high	
	3	11.69	36.957	44.657	0.21	high	e
259	52	0.3	0.327	1.414	0.14	high	a, f, g
276	2	1.54	28.564	66.985	0.142	low	h
375	6	0.04	10.024	24.905	0.175	low	i
384	–	–	–	1.781	–	high	f, j
386	9	0.66	10.073	17.117	0.035	high	
390	19	0.05	0.561	8.01	0.108	high	
406	6	0.05	10.025	24.905	0.175	low	i
460	5	0.54	2.718	32.932	0.489	high	
474	13	0.08	6.175	11.351	0.142	high	
483	4	0.03	1.606	48.643	0.176	low	i
486	31	0.08	1.893	4.83	0.07	high	a, k
533	23	1.34	2.887	6.407	0.209	high	
537	53	0.02	1.699	2.783	0.108	high	i, l
599	159	0.23	0.392	0.938	0.028	high	f, m
613	16	0.05	1.077	9.611	0.106	high	
624	22	0.05	6.709	3.238	0.142	high	
809	47	0.25	1.012	3.203	0.049	high	a, f
835	4	0.57	9.277	42.641	0.252	high	
915	51	28.51	2.84	2.903	0.105	high	a, n
917	5	0.11	25.459	30.428	0.213	high	
919	11	1.76	7.753	13.213	0.098	high	
937	18	0.35	4.932	8.459	0.105	high	
985	28	0.16	4.697	5.196	0.105	high	
1001	22	0.04	5.377	6.808	0.066	low	i

Comments: (a) eclipsing binary (secondary eclipses visible in phase-folded light curve), (b) very noisy light curve, (c) true period outside period search range, (d) detected period was half the true period, (e) triple system, second period was found by sine-fitting, (f) evidence of sinusoidal modulation, (g) transits exhibit phase shift with respect to the sinusoidal modulation, (h) light curve highly variable on short timescales, (i) detected transit very shallow, (j) no transits visible (other periodic behaviour), (k) detected period was 1.5 times the true period, (l) large number of transits improves confidence level, (m) detected period was 4 times the true period, (n) detected period was 2 times the true period.

6.3.2 Methods used by the other teams

The methods employed by the other teams for the light curve detrending and transit detection are briefly summarised below. In some cases, the information available at the time of writing is incomplete. For a fuller description see the upcoming article reporting on the exercise (Moutou et al. 2004a, in prep.) or the individual references for each method.

6.3.2.1 Team 1

A low-pass filter with a cutoff frequency between 0.059 and 0.117 d^{-1} was used to remove short-term variations (residual scattered light). This type of filter does not proceed to the edge of a dataset, and the first and last ten days of each light curve were removed before proceeding further. Local linear fits to small (0.5 to 3 d) sections of each light curve were then used to model and remove long-term stellar variations. The transit search was carried out with an algorithm described in Rauer et al. (2004), by which data points further than 3σ from the light curve average are flagged and contiguous sequences of such data points within a certain range of durations are considered as transit candidates. After removal of obviously spurious candidates, a period search was run on the epochs of the highlighted sequences. Finally, detailed visual examination was used to determine transit parameters and detect signs of stellar binaries (secondary eclipses and out-of-eclipse variations).

6.3.2.2 Team 2

In the method used by team 2, transit detection proceeds before detrending. The transit detection is carried out by comparing, at a given epoch, a short sequence of data to a short reference sequence (model transit) by plotting one versus the other and measuring the width (rms of residuals from a linear fit) of the resulting cloud. The sequence of widths at each trial epoch forms the detection curve DC, on which the detrending is applied. Detrending proceeds by applying a model of the form $DC_{i,j} = s_{i,j} + \lambda_i p_j$, where i is the detection curve number, j is the data point index, s is the signal, p is a perturbing trend common to a subset of all DCs and λ a scaling factor for each DC. Starting with an initial guess for p , the λ_i 's are found by projecting each DC onto p . A new guess for p is obtained by taking the group average of DC_i/λ_i over all i 's, and the λ 's are then recalculated. The DC with the highest scatter is taken as the first guess for p , and the resulting correction applied to those light curves which closely resemble it (high λ 's), then the most noisy of the remaining light curves is taken as the next guess, etc. . . Detrended DCs are visually examined and those showing periodic signals selected as candidates.

6.3.2.3 Team 3

Team three modelled the scattered light variations as a sum of harmonics of the fundamental orbit frequency (5 harmonics included), and the long term stellar variations as a sum of harmonics of the frequency $1/2T$, using 200 harmonics (i.e. down to timescales of 1.5 d). The 411 parameter model (5 high frequency harmonics each with sine & cosine components, 200 low frequency harmonics and a constant level) is fitted to each light curve and subtracted. This is basically identical to the least-squares fitting procedure described in Section 4.3 for power-spectrum estimation of data with gaps, but only the frequencies corresponding to the noise sources to be removed are fitted and the result is subtracted – resulting in the equivalent of a band-pass filter. The fitting process is made less computationally expensive by considering only those time points common to all light curves and using a single SVD pseudo-inverse of the fitting matrix to perform all the fits. The BLS algorithm of Kovács et al. (2002) – which is very similar to our box-shaped transit finder – was then used for transit detection. It is interesting to note (see Section 6.4.3.2), that a fit of the form $a + b/f$ where f is the trial frequency was subtracted from the distribution of their SR statistic with frequency, which is the equivalent in the BLS of the distribution of our S_M statistic with trial period. Light curves yielding a peak SR statistic greater than 7 were taken as candidates.

6.3.2.4 Team 4

Team 4 used a filter inspired from image processing techniques, the Gauging filter (Guis & Barge 2004, in prep.), for detrending. This filter treats the plot of a light curve as two 2-D images: one is black above and white below the light curve, and the other vice versa. Each image is undersampled and smoothed with a 2-D boxcar filter, the results for the two images are averaged and subtracted from the original light curve. Transit detection then proceeds with a standard matched filter, candidates being those light curves with periodic peaks in the distribution of correlation versus trial epoch, or those with one or two peaks deemed significant (according to criteria based on the entropy and variance of the peak distribution, see Guis & Barge 2004 for details).

6.4 Discussion of the results

The results of the exercise are summarised in Tables 6.3, 6.4 and 6.5. The first two of these tables list all the transit and stellar events inserted in the light curves together with relevant parameters and whether they were or weren't detected by

each group. The last table contains the false alarm and missed detection rates achieved by each group for transit and transit-like (eclipsing binary) events, and the results obtained by combining the candidates of all the groups.

Table 6.3: Characteristics of the transits that were inserted in the light curves. LD: stellar limb darkening coefficient. i : orbital inclination. a : semi-major axis. V : apparent visual magnitude. Detection flag: each sign corresponds, from left to right, to teams 1 to 5. +: the event was detected by the team in question. -: the event was missed. Contents provided by C. Moutou.

ID	R_* (R_\odot)	LD	R_{pl} (R_\odot)	p (days)	i ($^\circ$)	a (R_\odot)	V	*	Detection flag
34	0.92	0.6	0.025	5.52	91.2	12.77	13		+ + + + +
85	1.1	0.4	0.099	26.4	88.9	37.876	15	a	+ + + + +
168	0.92	0.5	0.13	11.5	87.4	20.827	15		- - + + +
207	0.92	0.5	0.11	88.4	90	79.89	16		+ + + + +
317	1.1	0.6	0.02	33.8	90.5	44.66	12		- - - - -
326	0.85	0.6	0.017	6.8	89.9	13.9	14		- - - - -
390	0.92	0.6	0.022	8.0	91	16.35	12		+ + + + +
460	1.1	0.3	0.076	32.9	89.52	23.49	15	b	+ + + + +
474	0.92	0.6	0.028	11.34	91	20.63	13		+ + + + +
533	0.92	0.7	0.095	6.4	90	7.89	16		+ + + + +
537	0.85	0.6	0.015	2.78	90.1	7.68	12		- - + - +
575	0.85	0.6	0.019	15.9	90	24.57	14		- - - - -
613	1.1	0.6	0.026	4.8	89.4	12.16	14		- + + - +
618	1.3	0.6	0.023	8.48	91	19.55	12		- - - - -
624	1.1	0.6	0.029	6.7	90.2	15.18	14		- + + + +
681	1.1	0.6	0.023	19.8	91.4	31.27	13		- - - - -
715	1.3	0.3	0.098	10.1	86.4	21.96	15		- - - - -
			0.07	63.8	89.7	75		c	- - - - -
835	1.1	0.4	0.084	42.6	89.3	52.1	15		+ + + + +
915	1.5	0.25	0.13	58.32	89.9	70	15		+ - + - -
		0.3	1.1	2.9	86	11.4		d	+ + + + +
917	0.85	0.6	0.028	30.4	90.3	37.8	13		+ + + + +
1000	1.1	0.6	0.02	33.8	90.5	44.66	14	e	- - - - -
1001	1.1	0.6	0.02	33.8	90.5	44.66	13	e	- - - - +

* (a) planet + ring, (b) planet + moon, (c) 2 planets, (d) circumbinary planet, (e) 32 s sampling rate.

6.4.1 Relative performances of the various groups

Teams 3 & 5 – the Geneva Observatory team and ourselves – detected significantly more transit events than the other groups. In fact, taking the overlap of the results of teams 3 and 5 – i.e. accepting a detection only if it was made by both teams – gives as good a result as taking the overlap of the result of all the teams (column B in Table 6.5). The transit search algorithms used by both teams were very similar: team 3 used the BLS method of Kovács et al. (2002), which is mathematically identi-

Table 6.4: Characteristics of the stellar events that were inserted in the light curves. BEB: background eclipsing binary. GrB: grazing binary. An x in the detection flag column means the event was detected but incorrectly identified as a transit. Contents provided by C. Moutou.

ID	V	Type	p (days)	$\Delta F/F$	Detection flag
31	14	BEB	24.7	0.003	+ + + + +
131	14	δ Cepheid	5.86	—	- - - + -
249	14	triple star	3.9	—	+ + + + +
259	16	GrB	1.4132	—	- + + + +
271	15	Z Cam	—	—	- - - - -
384	15	β Cephei	0.2835	0.001	- + - + +
386	15	GrB	17.1	—	x x x x x
486	15	BEB	2.4128	0.001	- - x - +
518	15	GrB	78.3	—	- - - - -
553	15	δ Scuti	0.07342	0.003	- - + + +
599	15	GrB	1.874	—	- + + x +
650	14	semi regular	—	—	- - - - -
809	15	GrB	3.2	—	- x + + +
919	16	GrB	13.2	—	+ x + x x
937	15	BEB	8.452	0.001	x x x x x
985	15	BEB	5.19	0.001	x x x x x

Table 6.5: Individual and combined performance of the detection teams. In columns A, B & C, which refer to detections made simultaneously by 1, 2 and 3 teams respectively, incorrect identifications of a stellar event count toward true detections if half of the teams that detected the event or more identified it correctly.

Team	1	2	3	4	5	A	B	C
Planetary transits		(23 inserted)						
True detections	10	11	14	11	14	16	15	13
Missed detections	13	12	9	12	9	7	8	10
False alarms	1	0	0	3	5	9	0	0
Stellar events		(17 inserted)						
True detections	4	6	8	8	9	10	9	8
Incorrect identifications	3	5	4	5	4	4	4	4
Missed detections	10	6	5	4	4	3	4	5
False alarms	0	0	0	0	0	0	0	0
All periodic transit-like events		(34 inserted)						
True detections	16	21	24	21	25	25	24	22
Missed detections	18	13	10	13	9	9	10	12
False alarms	1	0	0	3	5	9	0	0

cal to our box-shaped transit-finder, except that the light curve is phase-folded and rebinned before computing the detection statistic. This suggests that transit search algorithms based on least-squares minimisation with box-shaped transit models are a good choice. The other three teams used a variety of methods, some inspired by image processing techniques, some standard matched filter detectors. All have the common characteristic of being more complex than the methods used by teams 3 and 5. Simplicity therefore appears to be the best guarantee of robustness.

However, team 3 had no false alarms, whereas we had 5. The difference in false alarm rate between teams 3 and 5 may be due to the variability filtering and/or candidate selection methods. The fact that all our false alarms were low confidence candidates lends credibility to the latter hypothesis. The only low confidence candidate that wasn't a false alarm was light curve 1001. In that case the detected period was not the true period (i.e. the detected events were not the transits), so that it can, in a way, be considered to have been a false alarm.

6.4.2 Overall results and implications

A very reassuring result is the absence of overlap between the false alarms due to the various teams. By contrast, there is some overlap on all but one of the true detections – the only true detection made by one team only being light curve 1001, which was one of our low confidence detections. This bodes well for the exploitation of COROT light curves by teams who work in a concerted fashion, but rely on different methods, and suggests that internal exchange of candidate lists is advisable before any results are announced.

On the other hand, one might consider the fact that several of the detected eclipsing binaries were incorrectly identified as transits by most of the teams that detected them rather worrying. A total of 34 'transit-like events' – 23 planetary transits and 11 background or grazing eclipsing binaries – were inserted in the light curves. Of these (counting only detections made by at least two teams simultaneously), 15 were correctly identified as planetary transits, 9 were correctly identified as stellar eclipses, 12 were not detected, and 4 were detected but incorrectly identified as planetary transits by more than half of the teams that detected them.

The fraction of light curves containing stellar binary eclipses was more or less realistic, but the fraction of light curves containing planetary transits was approximately ten times what is expected (from extrapolation of the discovery rates of present RV surveys). If the trends from the present exercise are extrapolated to a more realistic sample, we would expect, for (say) 10 000 light curves, ~ 20 planetary transits and ~ 130 low-depth stellar eclipses to be detectable; leading to ~ 150 transit candidates of which ~ 130 would be stellar but only ~ 90 identified as such,




Figure 6.11: Depth $\Delta F/F$ versus number N_{tr} for all inserted planetary transits. Plus signs: non-detected events. Diamonds: events detected by all teams. Circles: events detected by 1 to 4 teams. Dashed line: estimated COROT detectability limit, $\Delta F/F \simeq -8.5 \cdot 10^{-6} N_{tr} + 6.5 \cdot 10^{-4}$. The superimposed plus sign and circle corresponds to three events: two undetected (IDs 317 and 1000) and one detected by 1 team only (ID 1001). The plus sign above the line corresponds to a very variable light curve (ID 715). Plot provided by C. Moutou.

leaving ~ 60 planetary transit candidates: a false alarm rate of over 65%! Such a high false alarm rate would not be acceptable in view of the time-consuming and difficult follow-up observations that are needed to weed them out. However, several points of COROT's observation strategy which are specifically aimed at reducing this type of false alarm were not included in the present exercise. It would also have been possible to correctly identify some of the eclipsing binaries from a more detailed examination of their light curves, but most teams chose to focus on the detection process in the limited time available. The false alarms rate due to background or grazing eclipsing binaries would therefore not necessary be as high as this exercise suggests.

C. Moutou used the results of the exercise to estimate detectability limits for COROT, by plotting the transit depth $\Delta F/F$ versus the number of transits N_{tr} (Figure 6.11) and drawing an empirical line, corresponding to a linear relation between p and N_{tr} , that separates events detected by more than one team from those that went undetected or were detected by one team only. The minimum detectable radii are summarised in Table 6.6 as a function of period for different stellar radii (i.e. spectral types).

Table 6.6: Estimated minimum detectable planet radius as a function of period p or number of transits N_{tr} . R_{pl} is expressed in units of the stellar radius R_* , or in units of Earth radii R_{\oplus} for different stellar radii, corresponding to a given spectral type.

p (days)	N_{tr}	Minimum R_{pl}				
		fractional (R_*)	F0V star (R_{\oplus})	G0V star (R_{\oplus})	G2V star (R_{\oplus})	K0V star (R_{\oplus})
75	2	0.0252	4.12	3.01	2.83	2.33
50	3	0.0250	4.08	3.00	2.81	2.31
30	5	0.0246	4.03	2.95	2.76	2.28
15	10	0.0238	3.89	2.85	2.67	2.20
10	15	0.0229	3.74	2.74	2.57	2.12
3	50	0.0150	2.45	1.80	1.68	1.39

These limits are slightly more pessimistic than those suggested by the results of the simulations described in Chapter 5 (Section 5.3.3). The latter implied a detection limit between 1.5 and $2.0 R_{\oplus}$ for a period of 30d and a G2V parent star, while the corresponding limiting radius, as given by the line shown on Figure 6.11, is $2.76 R_{\oplus}$. This difference can probably be explained by the more complete treatment of the instrumental noise in the present exercise. The method adopted by A. Lanza to simulate short-timescale stellar micro-variability is also more pessimistic than ours, as they effectively apply a scaling factor of 3 while we apply none.

6.4.3 Lessons learnt

6.4.3.1 Simulated light curves

A possible explanation for the residual micro-variability which, in a number of light curves, caused our algorithm to produce high single transit detection statistics (and multiple transit statistics barely higher than the single transit values) can be arrived at by noting that a number of the level 2 light curves contained remarkably similar transit-like variations (i.e. relatively sharp dips lasting a few hours). As shown in Figure 6.12, the resemblance between different light curves (during but also around the dip), together with the lack of repeat events in a given light curve, suggest that these are artifacts of the simulated stellar micro-variability components rather than inserted transit events. The dips are slightly distorted and shifted from one-another – presumably due to the scaling and shifting applied to the stellar micro-variability components to generate a sufficient number of non-identical light curves – but the similarities are striking. It is our opinion that this is due to a problem with the use by




Figure 6.12: Two examples (one per column) of recurrent transit-like features in similar looking light curves. All panels have the same time range. Arrows of the same colour on different panels point to the same transit like dips, which occur at different times due to the shifting and scaling applied to the stellar micro-variability light curves. Almost all the highly variable light curves in the sample closely resemble one of these two examples. Note that ID 117 is one of our false alarms.

A. Lanza et al. of scaled residuals to simulate the short-term variability component of their simulated light curves. This procedure is likely to amplify any artifacts in the residuals resulting from imperfections of the rotational modulation fit to the solar irradiance variations, and results in transit-like dips which we believe are not realistic. As these dips last only a few hours, they were not removed by the filtering process, and therefore led to a number of spurious detections.

6.4.3.2 Performance of our filtering and transit detection method

Our team obtained the largest number of correct detections. We also reported 5 ‘false alarms’, although all of these were low confidence detections. The team whose results in terms of sensitivity and robustness can be considered the best is probably team 3, and our results are identical to theirs if the low confidence detections are discounted. It is interesting to discuss the similarities and differences in the methods used by the two teams.

The least squares fitting method used by team 3 is equivalent to a bandpass filter, where the low-frequency continuum and the high-frequency noise are reconstructed by fitting sine curves up to a limiting frequency, then subtracted from the original light curve. In theory, there is no reason to expect this method to be more

capable of discriminating between stellar and transit signals than the iterative non-linear filter we used. The transit search method they used was very similar to ours, so that one wouldn't expect significant differences to arise from that step in the detection process. On the other hand, their candidate selection was based on a single threshold in their SR detection statistic.

Such an approach was not adopted in our case because, as discussed previously, a single threshold applied to the maximum of the distribution of our S_M detection statistic as a function of period gave unsatisfactory results. However, comparison of the two methods brought to our attention the presence of a systematic trend in the distribution of the S_M statistic with period, which is clearly visible in the middle panel of Figure 6.9. This trend is due to the change in the number of independent models tested with period. Among other parameters, the number of trial epochs and the number of in-transit points change with trial period. These effects are less pronounced in the BLS method (though a correction for this effect was applied by Team 3), in which the light curve is phase-folded at the trial period and rebinned into a fixed number of bins. However, that approach gives even statistical weight to each bin, rather than to each data point, which is not optimal, particularly in the presence of irregular sampling or gaps. The number of period cycles in the full light curve also changes with trial period, leading (for strong periodic signals) to discrete steps in the distribution of multiple transit detection statistic with trial period at harmonics of the true period.

It is difficult to derive a theoretical correction for the overall trend observed, because the value plotted for each trial period is itself the result of finding the maximum in the distribution of detection statistics over all trial epochs. Theoretically computing how the trend induced by the varying number of independent tests propagates through this maximum finding process would be prohibitively complex. We therefore experimented with an empirically derived weighting scheme, dividing the value of the detection statistic obtained for each trial period by a factor which scales with the 0.1th power of the trial period (normalised to the middle of the trial period range). This leads to a significant reduction of the trend, and was therefore adopted. This modification alone removes some of the false alarms.

It is also possible to improve our candidate selection by automatically determining the best trial duration, taking into account the ratio R of the scatter measured in the level 3 light curve to that expected from photon statistics (based on the median photon count), and by selecting candidates not only on the basis of the multiple and single transit detection statistics themselves, but also on the significance of the maximum in the distribution of S_M versus trial period. After a process of optimisation by trial and error, where various forms of threshold were experimented




Figure 6.13: Plots of S_M versus S_S (left) and S_M versus $\text{SNR}(S_M)$ (right) after slight modifications to the transit search algorithm and automatic duration selection. Grey: light curves with $R > 1$. Red: Light curves actually containing planetary transits or transit-like stellar events. Solid line: main detection threshold. Dashed lines: upper (high confidence) and lower (low confidence) thresholds.

with, the following procedure seems to give the best results:

- A single trial transit duration (0.18 d) is used in the filtering process, but three trial durations (0.18, 0.27 and 0.35 d) are used for the transit search. The decision to take this approach was based on the fact that almost all the candidates which were selected on the basis of the two higher trial durations and not the lower in the initial search were either removed at the individual light curve examination stage or were false alarms.
- For each light curve, the optimum trial duration is chosen in the following way:
 - If $R > 1$, which indicates residual stellar variability in the filtered light curve, the best trial duration is taken to be the shortest, which is the one that was used for the filtering, as using a longer filter duration would only lead to more residual stellar variability and hence a higher probability of false detection.
 - If $R < 1$, the best trial duration is taken to be that which leads to the highest ratio of multiple to single transit detection statistic S_M/S_S .
- The candidate selection is then based on the two plots shown in Figure 6.13: S_M versus S_S , but also S_M versus the 'signal-to-noise ratio' of the multiple transit detection statistic, $\text{SNR}(S_M)$, that is the ratio of $S_M - \langle S_M \rangle$ to $\sigma(S_M)$, where $\langle S_M \rangle$ is the median of the distribution of S_M with trial period and $\sigma(S_M)$ is the scatter of

this distribution.

- Rather than choosing selection thresholds on the basis of any theoretical consideration, they were chosen on the basis of the distribution of the points on the two plots themselves:
 - The plots were generated for all light curves in a ‘blind’ fashion, that is without showing which points corresponded to light curves actually containing transits, highlighting only those points for which $R > 1$ in grey.
 - Threshold lines were placed where a clear separation between the bulk of the points and the outliers was visible (solid lines on Figure 6.13). In the S_M versus S_S plot, this results in a line defined by: $(S_M > 7 \text{ AND } S_M > 1.65 \times S_S)$ OR $S_M > 20$. In the S_M versus $\text{SNR}(S_M)$ plot, the line is defined by: $S_M > 20$ OR $\text{SNR}(S_M) > 11$.
 - Slightly higher and lower threshold lines, designed to respectively exclude and include a few more points in each case, were also chosen (dashed lines on Figure 6.13).
 - Selection then proceeds according to a points system: one point for a point above the lower dashed line, two for a point above the solid line, and three for a point above the upper dashed line in each plot. This results in points scores between 0 and 6, and candidates are light curves with scores of 3 and above. This procedure enables one to probe slightly into the bulk of the points in one of the two plots provided the point corresponding to the same light curve is well separated from the bulk in the other plot.

The above procedure is almost automatic – the user must only choose the position of the threshold lines, but only on the basis of the distribution of the points on the two scatter plots rather than by examining light curves. It leads to zero false alarms and 10 missed detections (counting both planetary transits and transit-like stellar events). The missed detections are the same as before plus ID 460 which was previously detected but is now missed – it is a long duration event.

A number of light curves with scores of 1 or 2 have best durations other than the filter duration used. In those cases, we tried repeating the transit search using the best duration as the filter duration, then applying the same selection procedure. This yields two more detections – IDs 460 and 317 – but also two false alarms – IDs 309 and 995.

Modifications in the transit search and candidate selection procedure therefore allow us to obtain 8 missed detections and two false alarms without having to examine any light curves individually. We cannot objectively say whether the two false alarms would have been excluded at the light curve examination stage as we

already know the content of the light curves and could not reproduce the process in an unbiased way.

6.4.3.3 Parameter estimation and rejection of stellar events

The method used for parameter estimation was very basic. It was also made difficult by the need to make a guess of the transit duration to filter the light curve before the parameters could be measured. In future, methods to filter the light curve without a prior guess of the duration (but given knowledge of the period) will be investigated. The transit parameters will be measured in a more optimal way, e.g. with a matched filter applied to the phase folded-light curve, using the same type of transit modelling software to generate model transit light curves as was used to generate the signals inserted in the light curves for the present exercise.

Can careful consideration of the shape of a suspected planetary transit help reduce the contamination by stellar events? We consider three categories of stellar eclipses that can mimic planetary transits: grazing, high mass ratio, and blended (the light of a third star aligned with the eclipsing system dilutes the eclipses). The first and some of the second may be identified by relating the transit observables (depth $\Delta F/F$, total duration d_1 , duration of totality d_2 , period p) to the geometry of the system (as illustrated on Figure 1.1).

It is immediately apparent, for example, that the lack of a flat-bottom to the eclipse ($d_2 = 0$) is a tell-tale sign of a grazing event. This is the case for light curve IDs 249, 259, 599, 809 and 915 (which were identified as stellar events due to the presence of secondary eclipses or out-of-eclipse variations), but also of IDs 386 and 919 (which we reported as candidate planetary transits when they were in fact grazing eclipsing binaries). Of course, a planetary transit can also be grazing, but grazing planetary transits are even shallower and rarer than non-grazing ones, which makes it all the more unlikely that one would have been observed, let alone detected.

High mass ratio events can be identified by computing the star's density from the light curve, as discussed by Seager & Mallén-Ornelas (2003). In Chapter 1 we derived the dependence of the transit depth (Equation 1.6), the full transit duration (Equation 1.11) and the duration of totality (Equation 1.12) on the star radius R_* , the planet radius R_p , the period p , the orbital distance a and the inclination i of the system, assuming that the orbit is circular, that the planet is dark, that there is only one star – the eclipsed object – in the system, and that there is a flat bottom to the eclipse – ie the impact parameter $b \leq R_*$. In this framework, Seager & Mallén-Ornelas (2003) show that, these equations can be used to derive R_p/R_* , b/R_* and a/R_* from the light curve observables ($\Delta F/F$, d_1 , d_2 and p) alone, and, using them in combination with Kepler's third law (Equation 1.4), the star density ρ_*/ρ_\odot (under

the additional assumption that $M_p \ll M_*$). In theory, we should thus be able to identify eclipsing binaries where the primary is a giant star, without a spectrum or multi-colour information (provided the errors on the observables are low enough, but the high time sampling and photometric precision of COROT should allow this). This method cannot, however, be used to exclude cases where the primary is a main sequence star and the secondary a brown dwarf, as both the density of the eclipsed object and the radius of the transiting object will be consistent with a planet orbiting a main-sequence star.

The unique solution to the transit equations discussed above breaks down in the case of a blend, i.e. if light emitted by third object, spatially coincident or fortuitously aligned, dilutes the observed eclipses. In such cases, any secondary eclipses may be diluted to the point of being undetectable. Other observations than the single bandpass light curve alone are thus needed to identify these events as stellar rather than planetary. This is the case for the remaining two light curves which we incorrectly identified as candidate planetary transits, IDs 937 and 985.

6.5 Future prospects for COROT blind exercises

The present exercise was considered extremely useful by all participants, and a second generation exercise is foreseen for early 2005. Improvements will hopefully include the inclusion of colour information and improvements in the simulation of stellar micro-variability, based among other things on the soon to be released MOST data. It would also be desirable to simulate an entire COROT field, with stellar distributions and crowding derived from galactic models and the results of preparatory ground-based observations. It is foreseen that this second exercise will have a higher profile, groups outside the COROT EWG (e.g. the *Kepler* team) being invited to participate. More emphasis will also be devoted to the use of oversampling and to the determination of transit parameters after detection.

In the mean time the light curves from the present exercise will be made available to the entire transit search community, while the teams which participated will continue to attempt to improve their algorithms now that the content of the light curves is known.

From our point of view, this exercise also highlighted areas where future efforts should be concentrated. There is room for improvement in the variability filtering and transit candidate selection procedures. The pre-processing steps (in particular the removal of systematic trends) and variable star search tools, of which only the very basic lines were laid out here, will also play an important part in maximising the scientific return from COROT and other transit search data, and they need to be

developed further.

Appendix: Phase-folded light curve plots

Figures 6.14 to 6.21 contain plots of the normalised phase folded light curves for all the transit candidates reported by the IoA team. There is one row per transit candidate. In each case the left panel shows the full phase-folded light curve, the right panel the phase folded light curve around the transit. The blue line in the right panel is the binned phase-folded light curve from which the transit parameters are computed. The red line in both panel shows the box-shaped transit model corresponding to the reported parameters. The light curve ID is indicated in the top of the left panel. The period and input epoch (in days), as well as the filtering duration (in number of time steps and in days), are listed at the bottom of the left panel. The measured duration (in number of time steps and in days), epoch (in days) and transit depth (in %) are listed at the bottom of the right panel.

In the cases where sinusoidal variations corresponding to the period of the transits were visible (IDs 259, 384, 553, 599 & 809), the phase folded level 2 light curve is shown in a single panel below the two panels showing the level 3 light curve. The two rows for light curve 249 correspond to the two sets of eclipses with different periods (star 249 turns out to be a triple system). Light curves 384 and 553 showed only sinusoidal variations and no transits, so that only the phase folded level 2 light curve is shown. The plots for star 259 clearly show the shift in phase of the transit relative to the sinusoidal variations. This is probably due to a small error in the detected period.

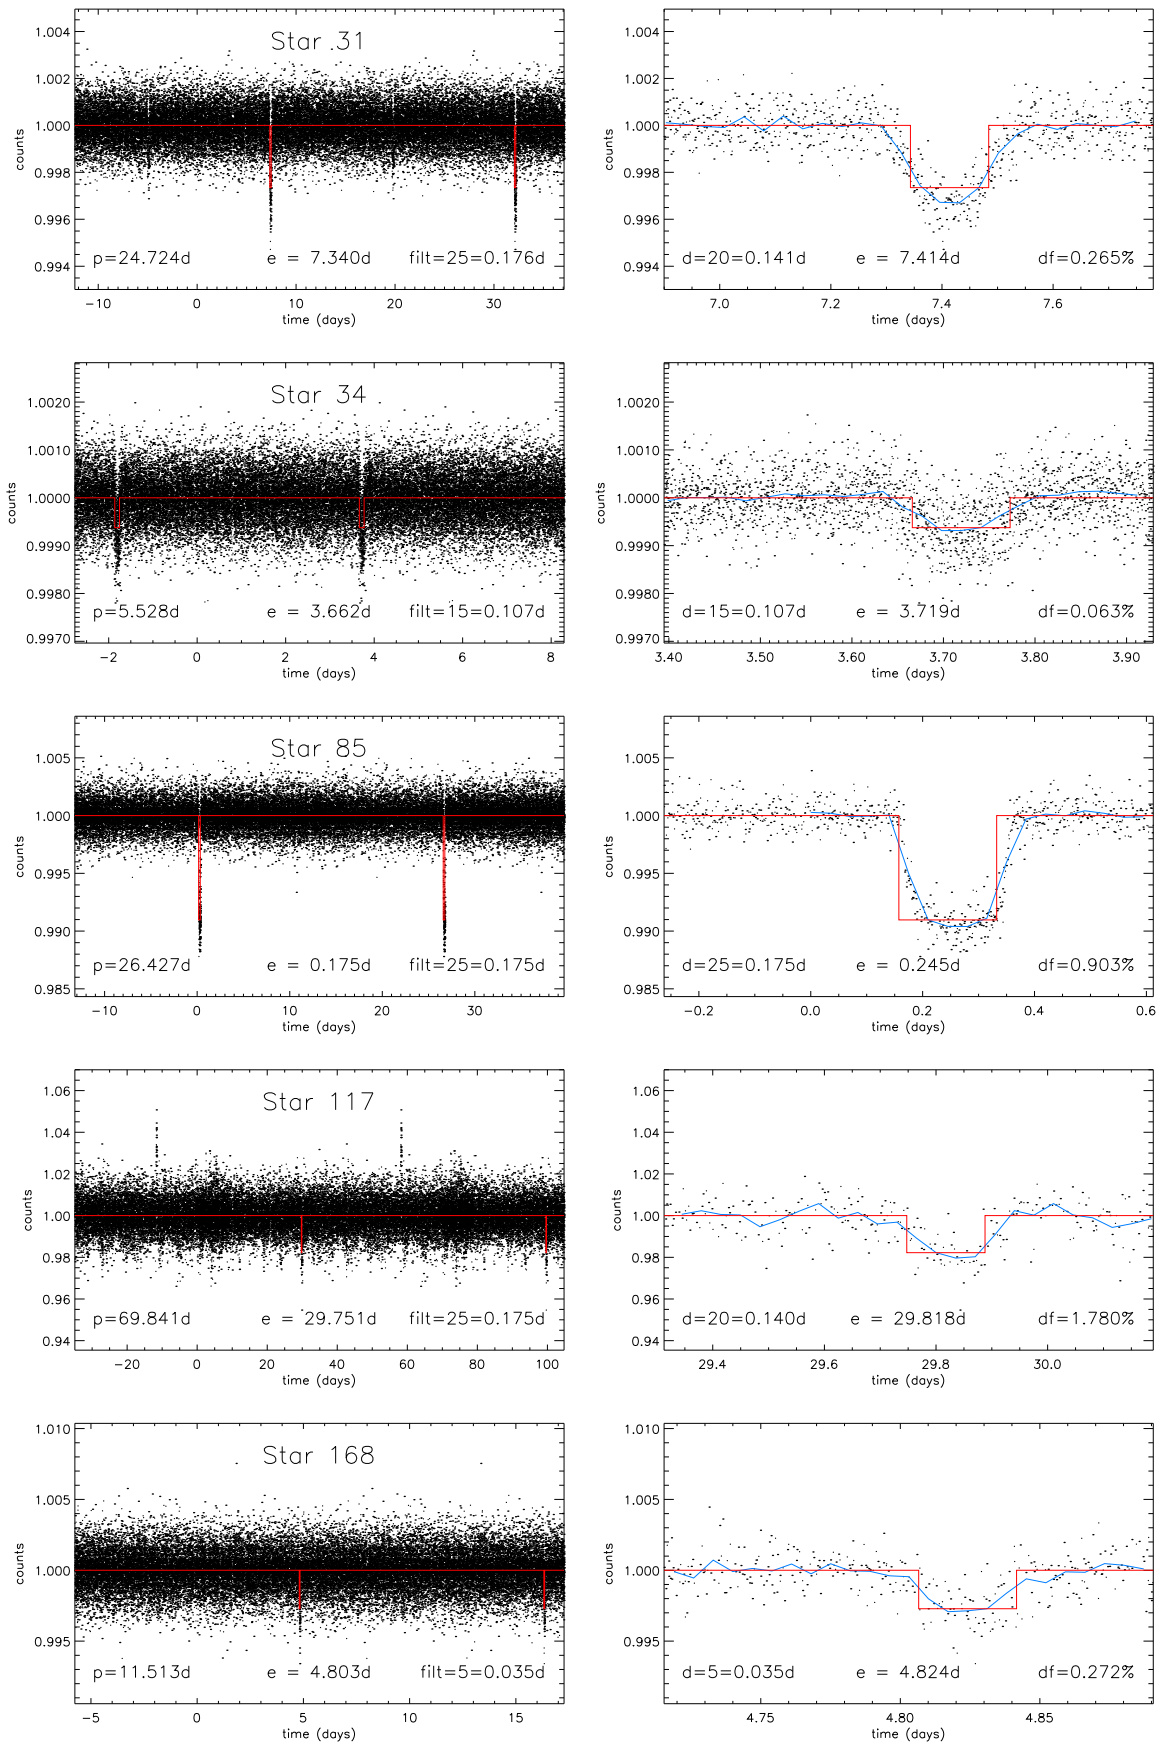


Figure 6.14: Light curves of transit / variable star candidates – page 1.

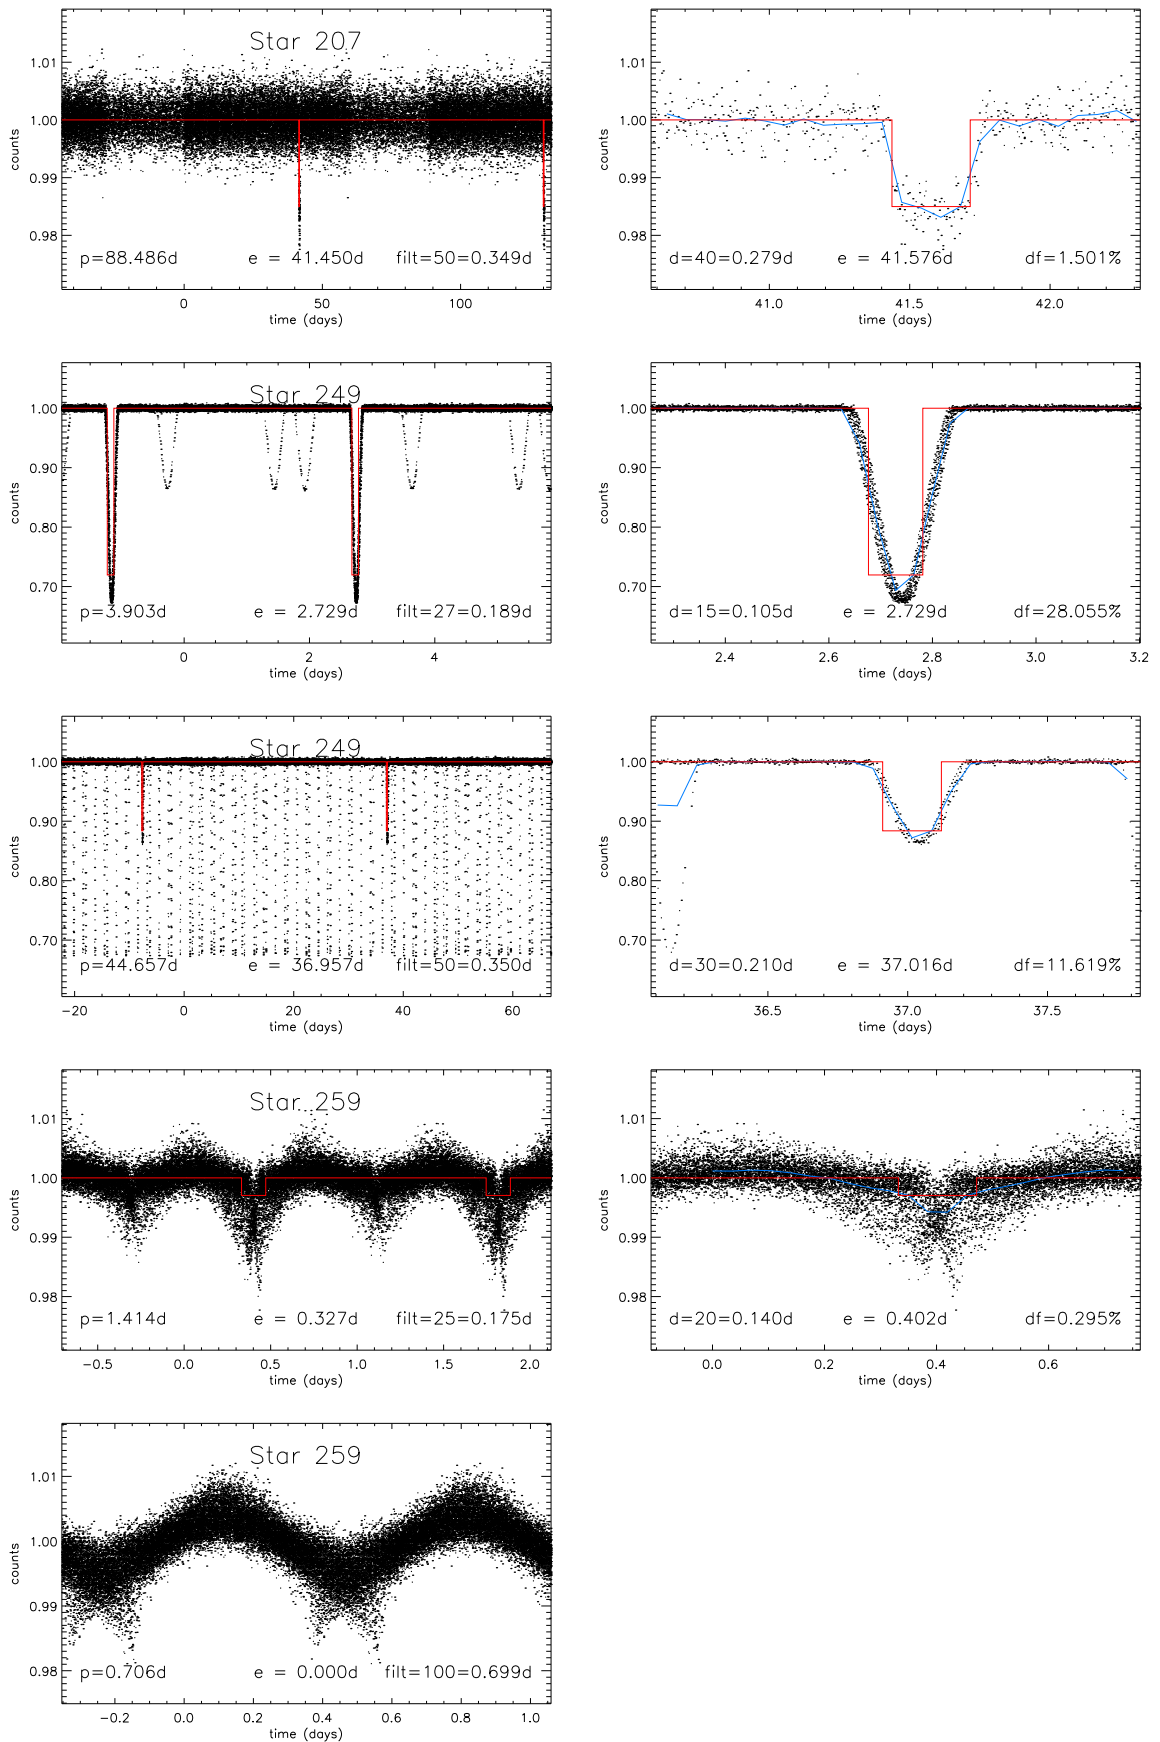


Figure 6.15: Light curves of transit / variable star candidates – page 2.

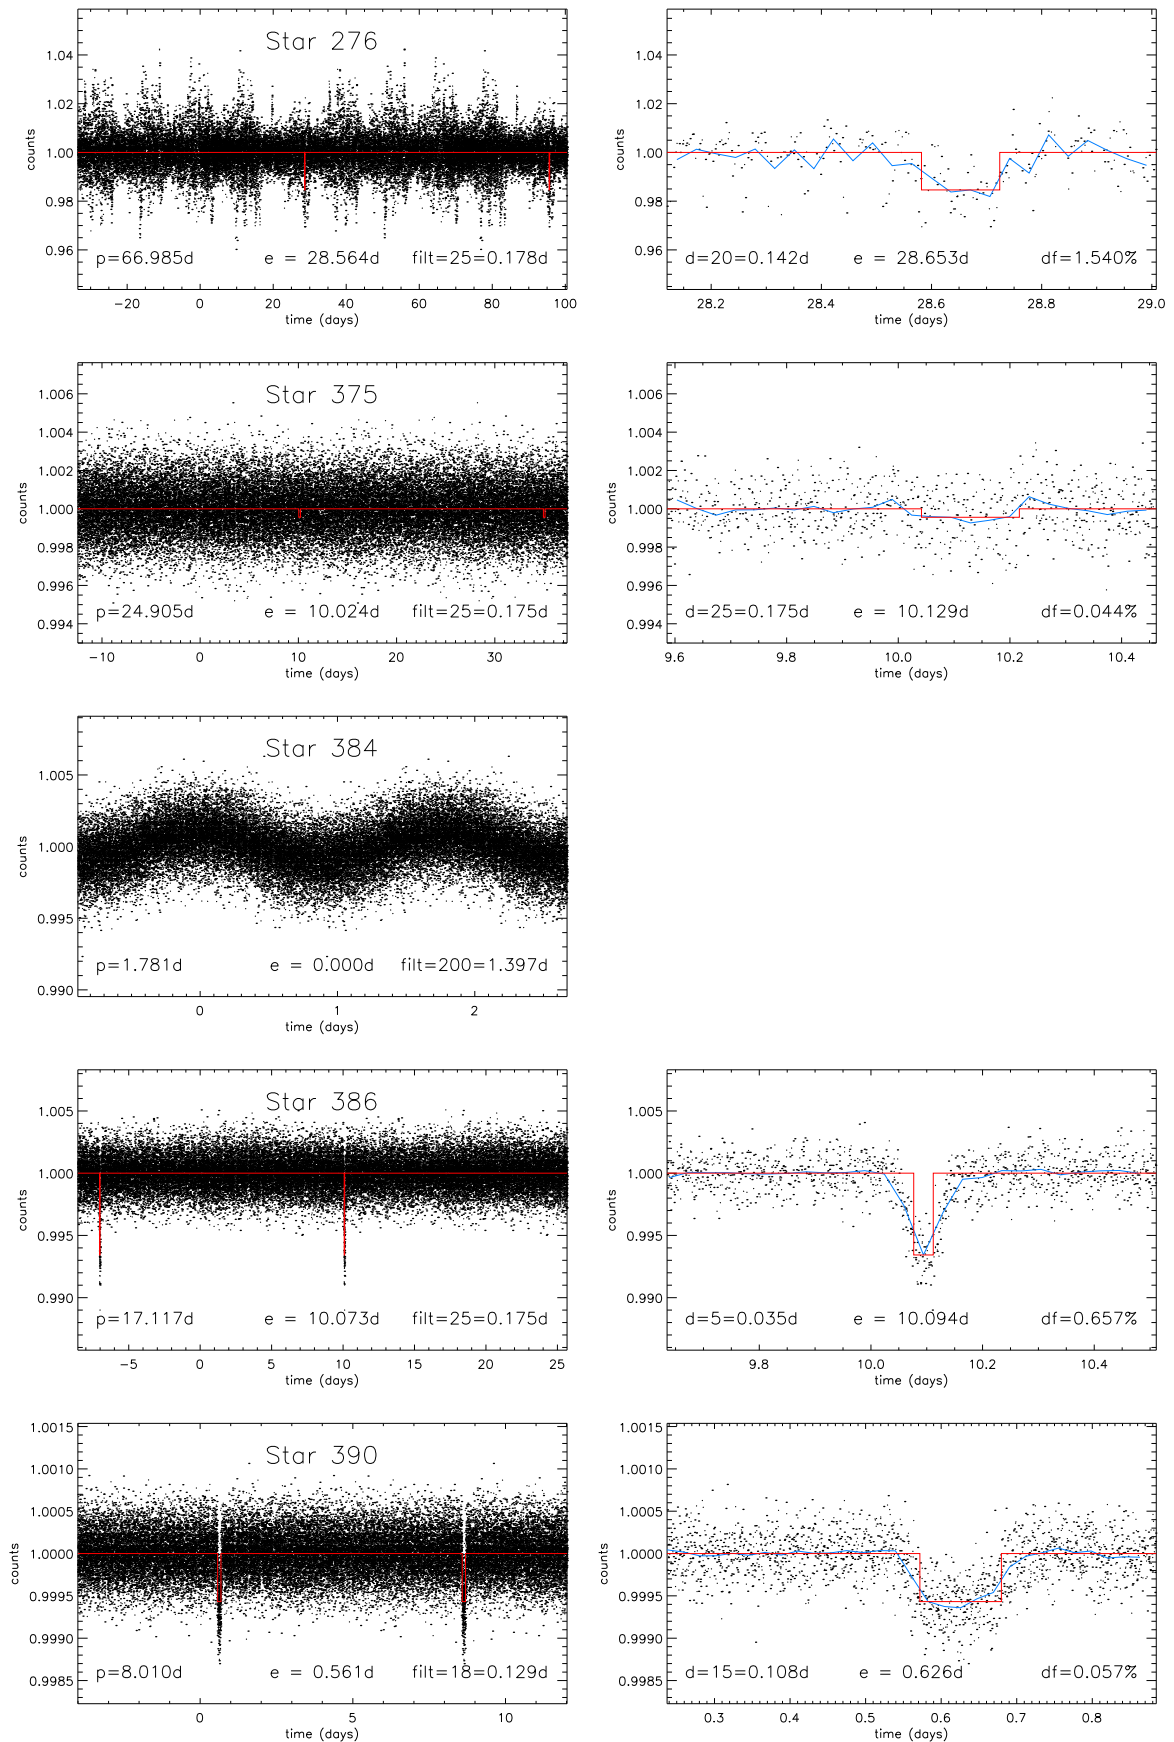


Figure 6.16: Light curves of transit / variable star candidates – page 3.

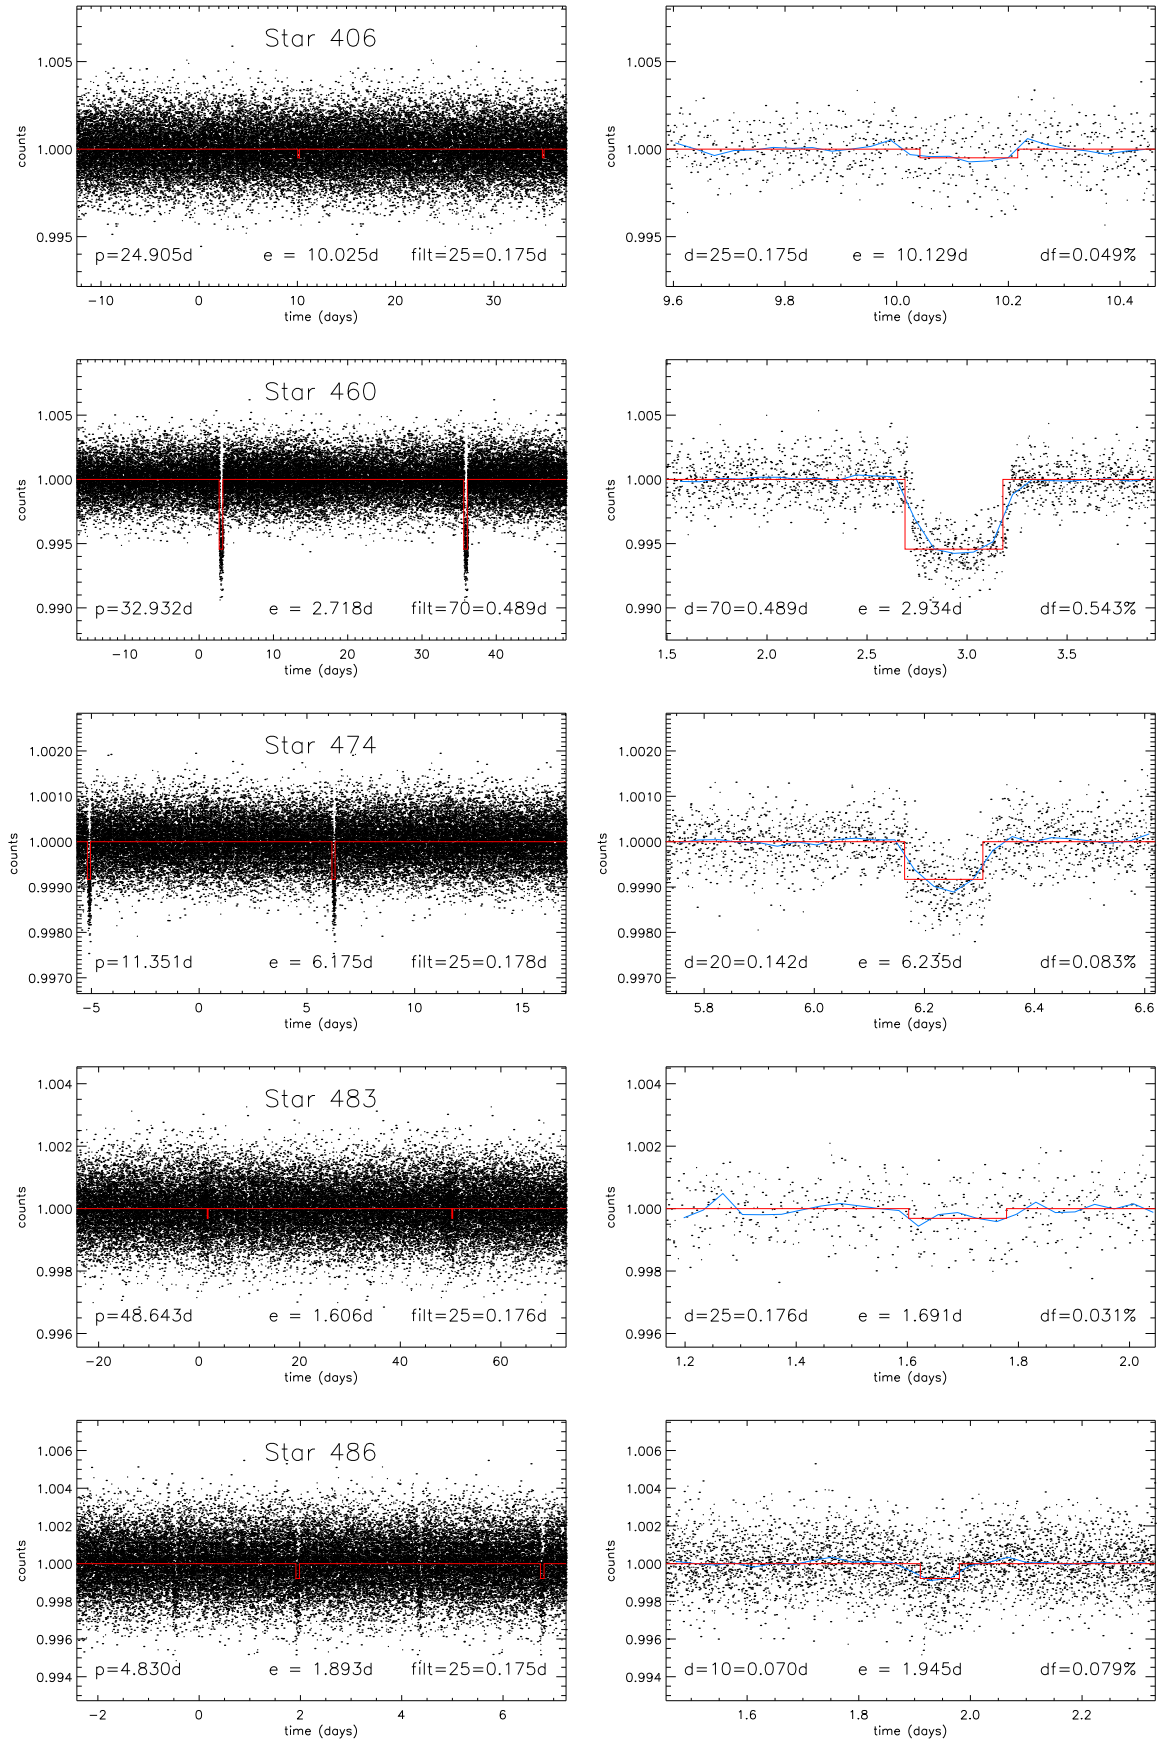


Figure 6.17: Light curves of transit / variable star candidates – page 4.

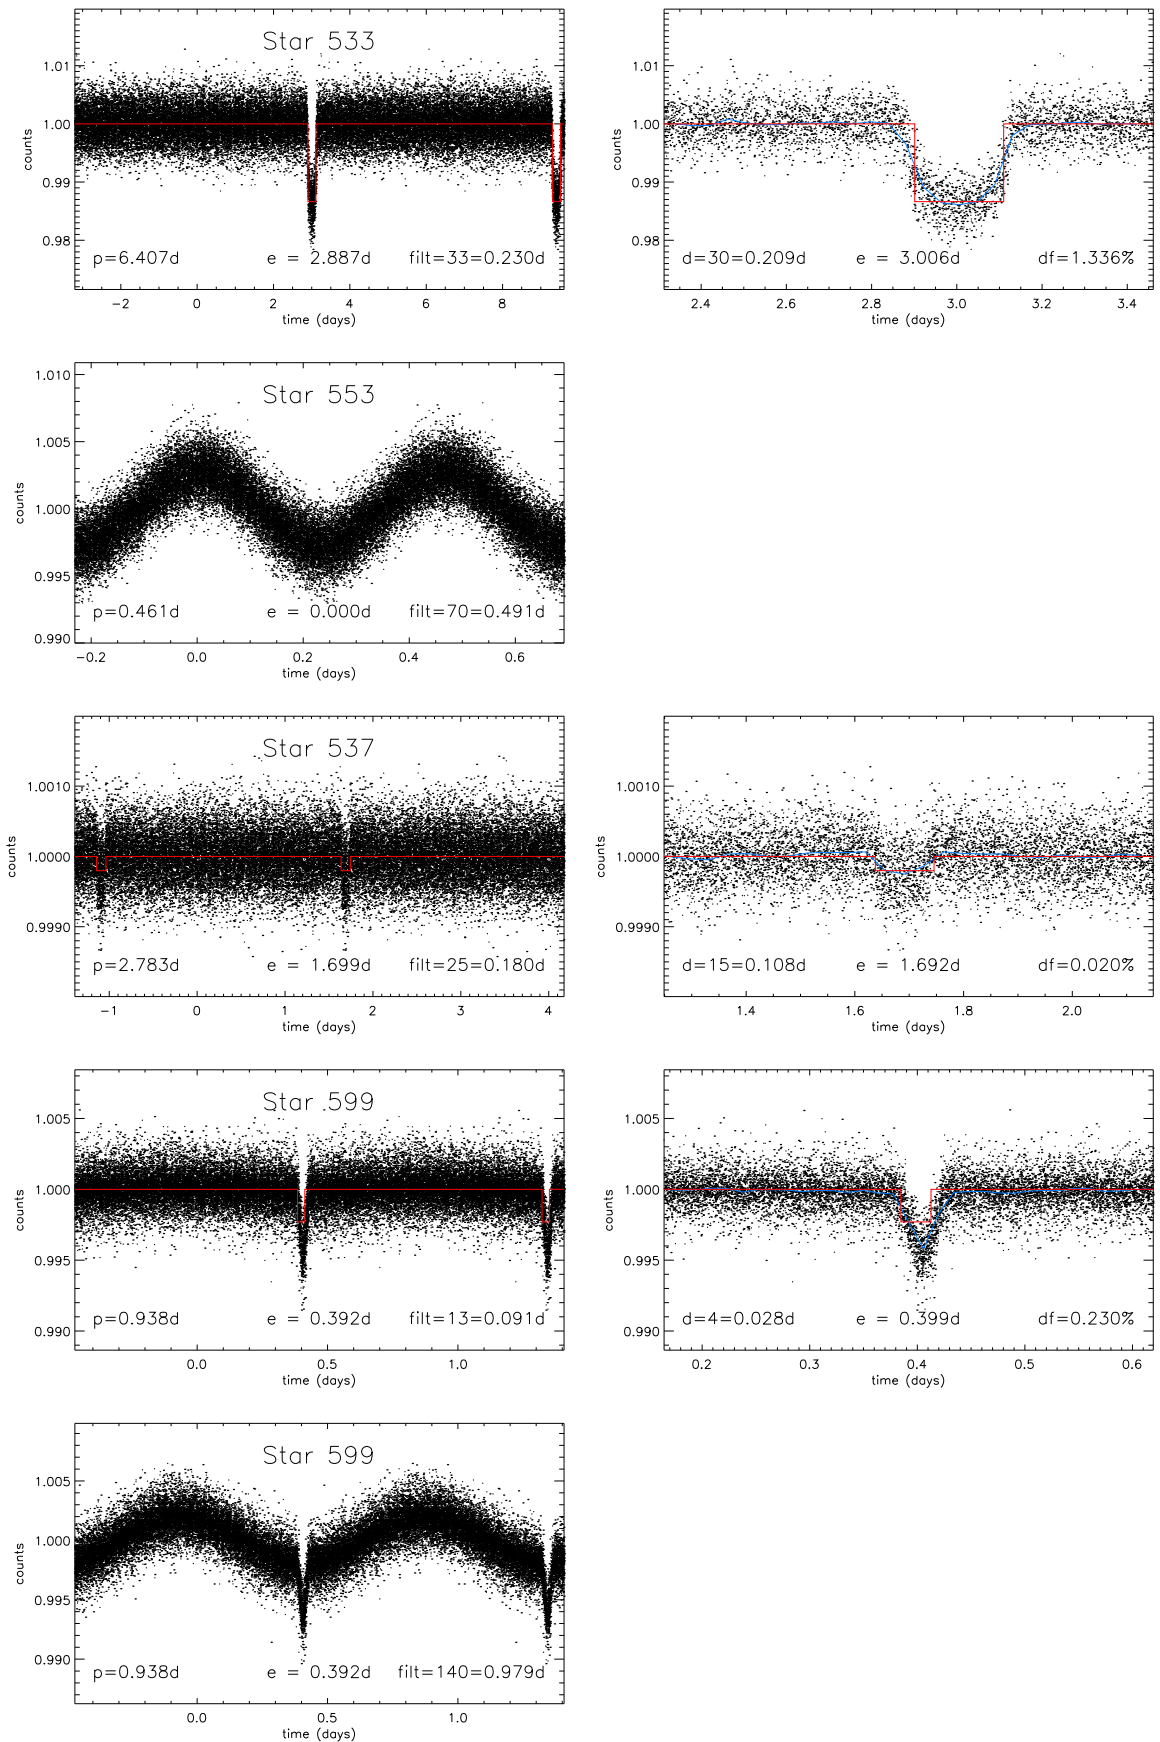


Figure 6.18: Light curves of transit / variable star candidates – page 5.

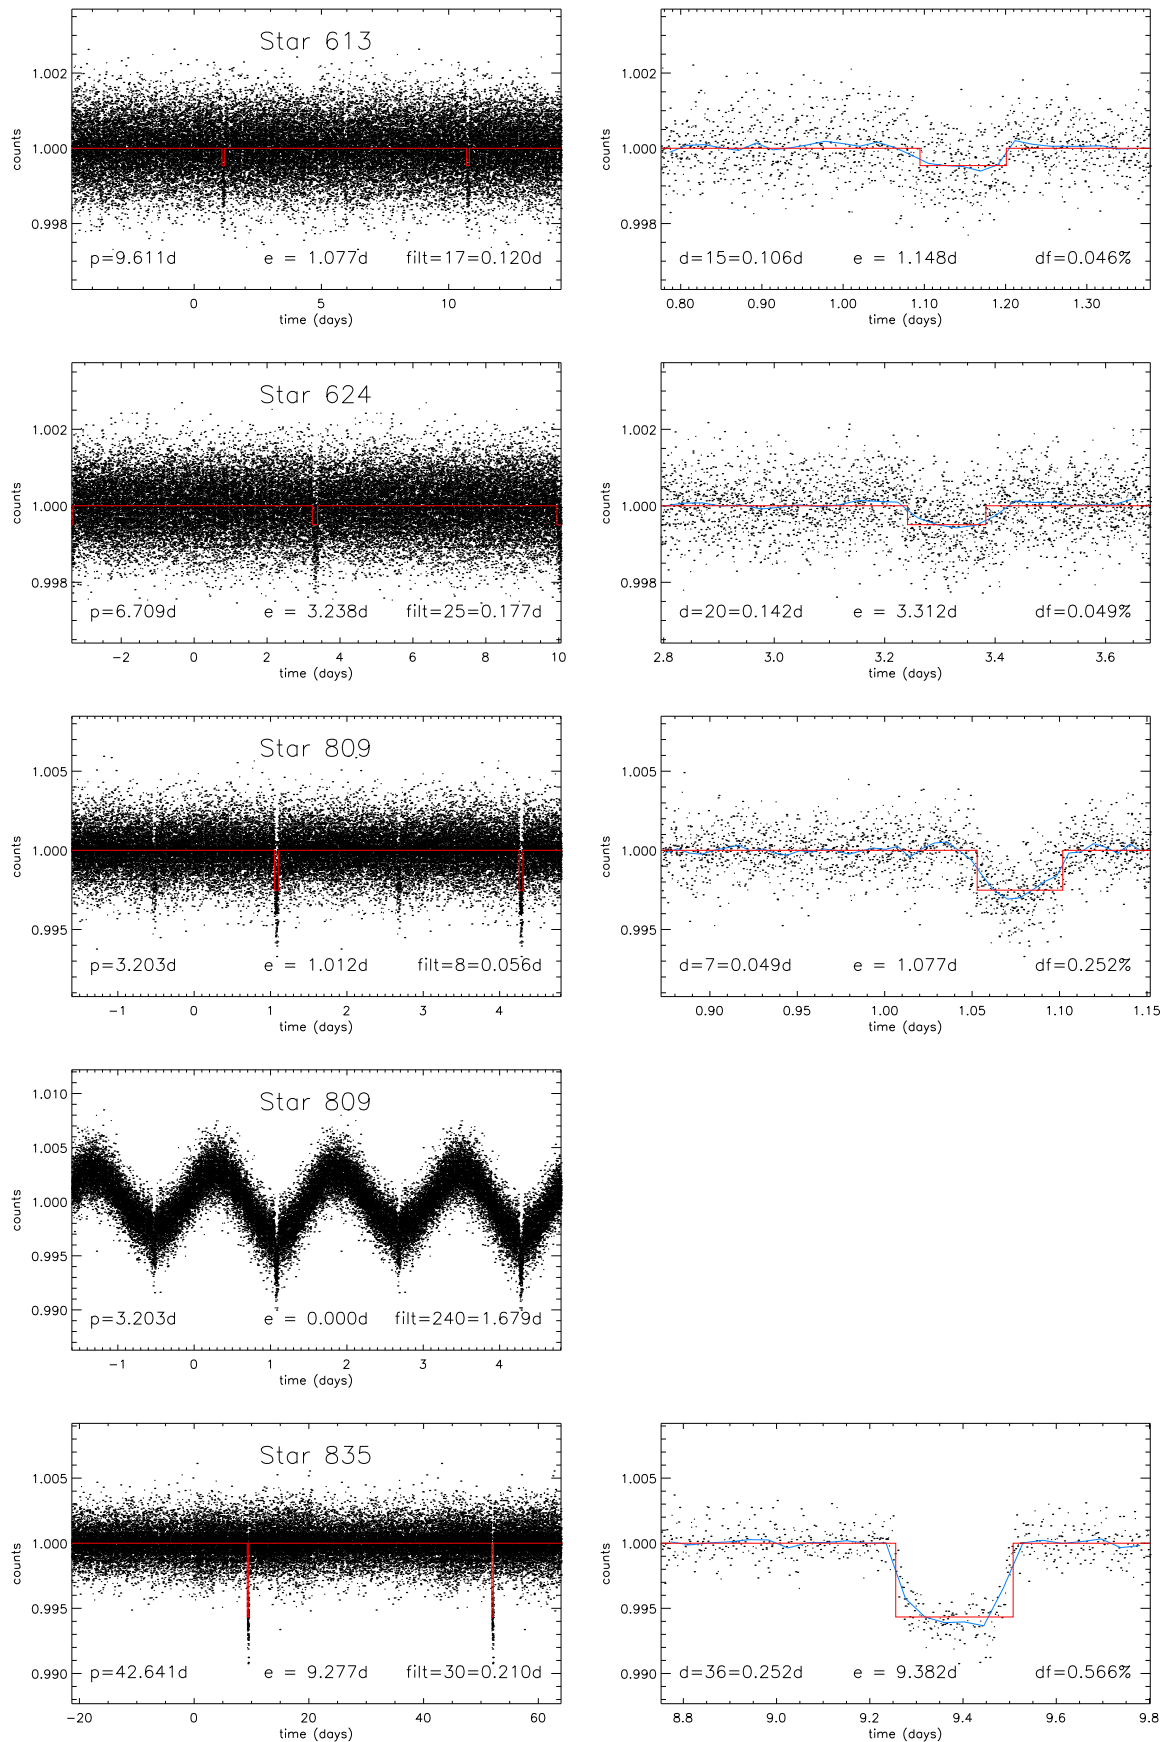


Figure 6.19: Light curves of transit / variable star candidates – page 6.

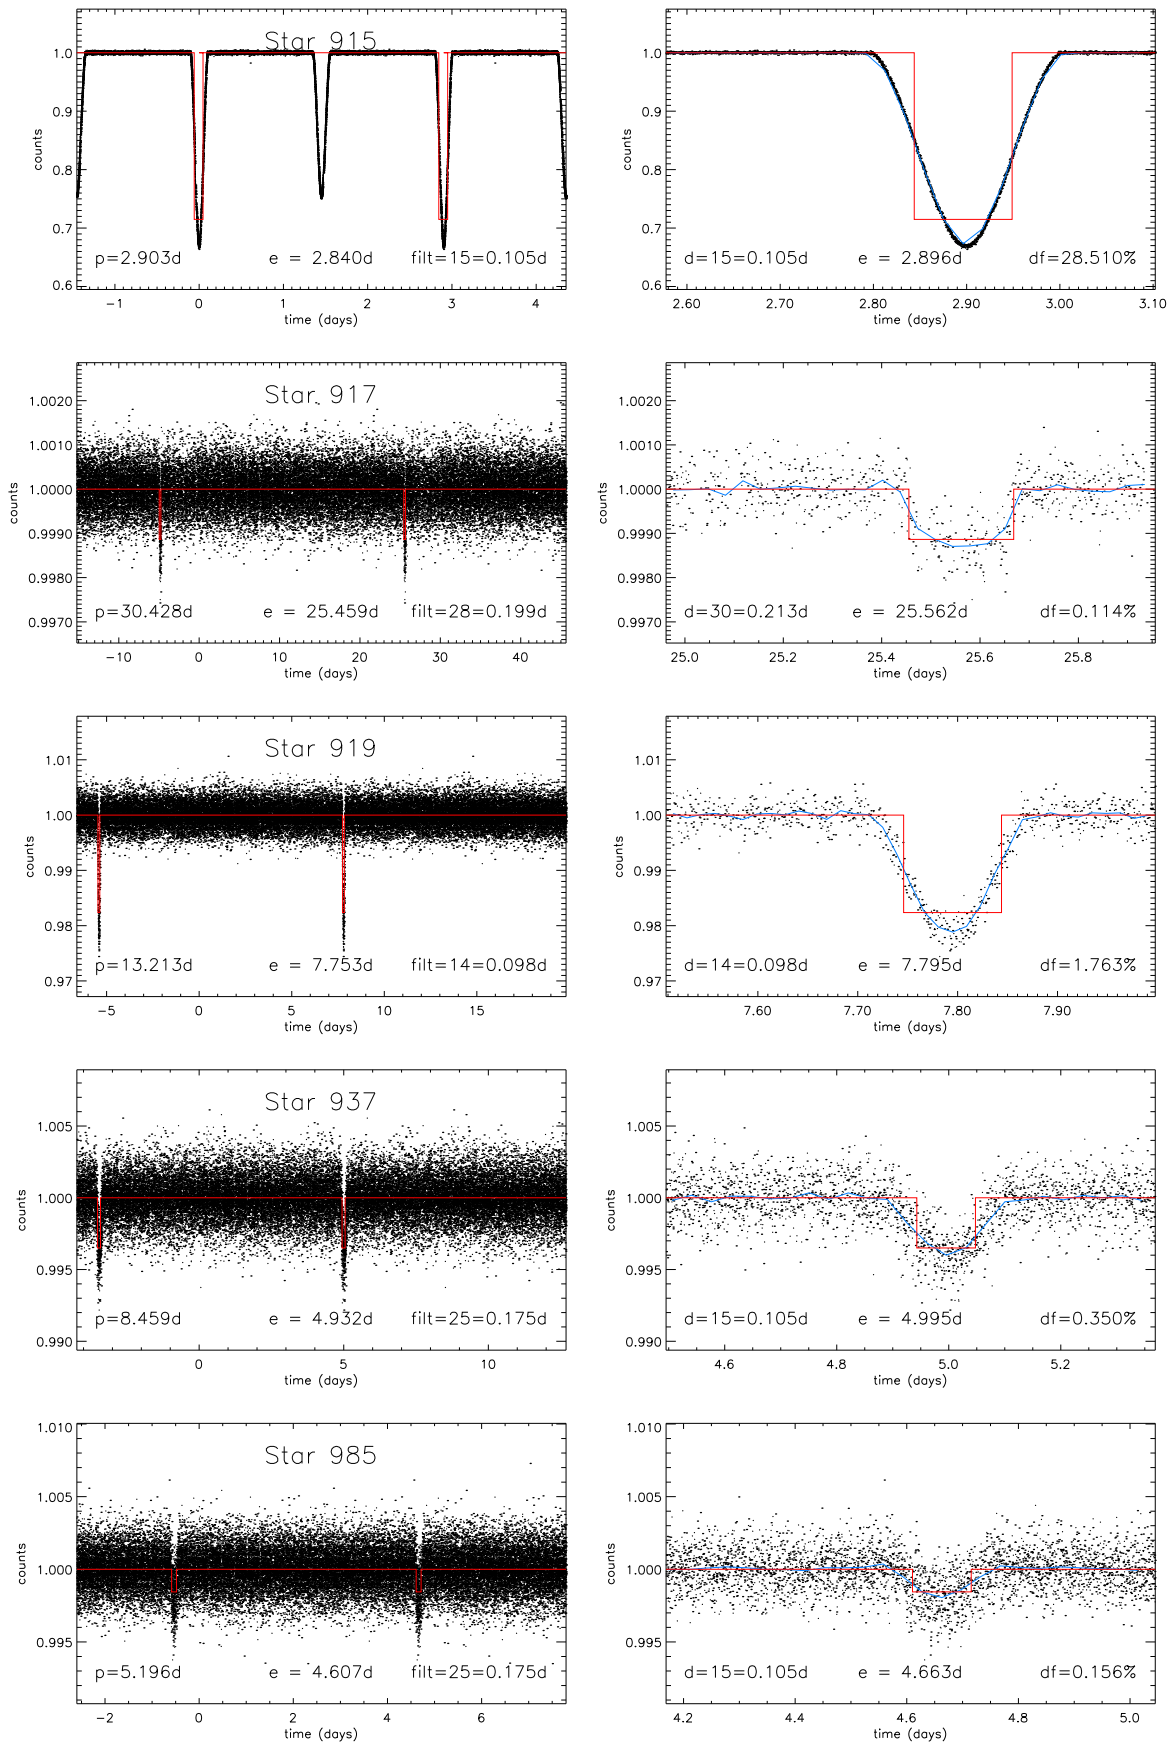


Figure 6.20: Light curves of transit / variable star candidates – page 7.




Figure 6.21: Light curves of transit / variable star candidates – page 8.

Chapter 7

Application to ground-based data

While the tools developed in Chapters 2 & 4 were primarily intended for application to data from space-based transit searches, their application to existing ground-based data provides a fertile testing ground, as well as the potential for scientific return on a shorter timescale. Applications to real data are also bound to highlight problems which did not manifest themselves in the tests on simulated data described so far. In the present Chapter, a brief overview of early experiments with real data is presented. This is not intended as a report on a fully-fledged transit search, but rather as an example illustrating the learning process and the kinds of problems that were encountered.

7.1 The data

The dataset used is part of the University of New South Wales (UNSW) transit search project¹, which uses the Automated Patrol Telescope (APT) at Siding Springs Observatory (SSO). Mike Irwin and myself have been involved in a collaboration with the UNSW transit search team since 2002, our contribution focusing mainly on improving the photometric accuracy of the light curves and on searching for variable stars and transits.

7.1.1 Telescope and instrument

The APT's design is similar to that of a Schmidt camera, but uses a 3-element lens to achieve a wide, corrected field-of-view (FOV). It is entirely computer-controlled and can be operated remotely or – within limits set by the reliability of the weather monitoring systems – automatically. It has a 0.5m aperture and f/1 optics yielding

¹See <http://www.phys.unsw.edu.au/astro/research/planetsearch.html>.



Figure 7.1: The Automated Patrol Telescope. Figure 7.2: A sample APT image of NGC 6633.

a 5-degree flat-field. The current CCD covers a 2×3 degree FOV with 800×1200 pixels, each 9.4 arcsec on a side. Intra-pixel sensitivity changes of the CCD become significant for such under-sampled images, where the light of a star covers only a few pixels. The UNSW team therefore developed an observing technique to overcome this problem, whereby the telescope is systematically moved in a 1×1 or 2×2 pixel 'raster-scan' pattern during each exposure. This has the effect of broadening the PSF slightly, but any intra-pixel sensitivity induced variations are averaged out.

At the time the data used here were taken, the raster-scanning technique was being experimented with and different scan widths were tried on different nights. We used only data taken with a 1×1 pixel scan, which we found to give the best results. The resulting dataset spans 6 nights, with the 4th night (when no scanning was used) missing. Exposure times were 150s, which leads to stars brighter than $V \sim 9$ being saturated.

The telescope is equipped with broadband B , V , R & I filters, but filter changes can only be carried out manually, so that all transit search observations in a given run are performed using a single filter, usually V (as is the case for the data used here).

7.1.2 Target

The data used in this Chapter are from a field centred on the open cluster NGC 6633.

NGC 6633 ($\alpha = 18^{\text{h}}27^{\text{m}}$, $\delta = +06^{\circ}34'$; RA = $11^{\text{h}}07^{\text{m}}$; $l = 36.1^{\circ}$, $b = 8.3^{\circ}$) is a young open cluster of solar metallicity. Its age is $\log t = 8.63$, its reddening $E(B - V) = 0.18$ and its distance modulus $(V - M_V)_0 = 7.77$ (all cluster information taken from Martín et al. 2004). Its youth, moderate reddening and relative proximity, together with the fact that it falls within the COROT visibility zone towards the galactic centre, make it a particularly interesting target for temporal variability studies. Consequently, this cluster was the focus of a multi-site monitoring campaign in August 2002 (Martín et al. 2004). Prior to that campaign, two γ Doradus variable candidates were known in NGC 6633 from monitoring in the Strömgren system (Martín & Rodríguez 2002). Martín et al. (2004) reported on the analysis of data from the multi-site campaign for 192 objects with $V < 12$. Of the 60 stars which were catalogued as belonging to the cluster, they found 12% were variable, and a further 28% were suspected of variability. Variable types included δ Scuti stars, γ Doradus stars, Ap-type variables, eclipsing binaries and long-period variables. The APT data were taken as part of this campaign, but were in the process of being reduced at the time the Martín et al. (2004) paper was written, and were thus not included in that analysis.

Note that, due to the fact that NGC 6633 is situated in the galactic plane, the field is crowded and most of the stars in the magnitude range of interest may not be members of the cluster.

7.1.3 Data reduction and photometry

The data were reduced using a purpose-made pipeline developed by the UNSW planet search collaboration on the basis of the CASU² pipeline for the Wide Field Camera (WFC) Survey on the Isaac Newton Telescope (INT). Images were de-biased and trimmed, bad pixels were detected and interpolated over, and a flat field correction was applied using a single master flat-field image generated from twilight flats from several nights.

The photometry and production of light curves were then carried out in Cambridge. A master image was produced by stacking a large number of individual frames, and a deep master catalogue was then generated by running the *imcore* source extraction software on the master image. The source extraction software was also run on the individual frames, allowing a very precise astrometric transformation between each frame and the master image to be computed. Standard aperture photometry (with a core radius of 2 pixels, approximately matching the FWHM of the

²Cambridge Astronomical Survey Unit.




Figure 7.3: RMS precision versus V-band magnitude for 1 night of observations in the field of NGC 6633 (exposure time 150s, 107 observations). The red dashed line represents the expected rms from object and sky photon noise.

images) was then carried out on each frame using the positions derived from the master catalogue. Finally, a correction equal to the average departure from the light curve median over each frame was applied (this is equivalent to a very crude relative zero point, including atmospheric extinction correction). Frames for which the global rms exceeded 10 mmag were discarded as indicative of non-photometric conditions.

The resulting rms precision is ~ 2 mmag for unsaturated stars brighter than $V \sim 10.5$ within a single night (see Figure 7.3). The backbone of points at the bright end is shifted slightly upward from the expected value from photon noise statistics, showing that other low-level noise sources remain.

7.1.4 Post-processing

The light curves contain residual systematic trends, which result in a degradation of the precision when several nights are joined together (see Figure 7.4, left panel). Examination of the light curves reveals that most display parabolic trends within each night (as one would expect from atmospheric extinction and airmass dependent seeing), as well as night-to-night shifts. These systematics are very damaging for transit searches, as downwards trends at the edges of nights mimic transit ingress or egress events, while a long term trend or night-to-night shift can lead to a spuriously high transit detection statistic because the duration of a night of observations (≤ 6 hours in this case) falls within the range of reasonable transit durations.

A more accurate extinction correction, by which nightly extinction coefficients are derived from iterative parabolic fits to the light curves, was initially experimented




Figure 7.4: RMS precision versus V-band magnitude for the 5 nights of observations in the field of Open Cluster NGC 6633 (exposure time 150s, 417 observations) without (left) and with (right) systematics correction.

with, but in the present case was not found to significantly improve the results.

The standard practice to obtain a “systematics free” light curve for a given object is to identify trends common to a number of nearby comparison stars, which are assumed to be constant apart from noise and the systematics of interest, and to subtract those trends from the light curve of the star of interest (see for example Gilmore 1979). In the age of small field-of-view detectors with non-linear response, care had to be taken to account for the non-linearities and to ensure that the comparison stars really were constant, or to account for the possibility that they might not be. With wide field-of-view CCDs, many comparison stars are available. One can rely on averaging over many comparison stars to iron out the effects of intrinsic variability of any given comparison star. This is, in effect, implemented through the global zero-point correction described at the end of Section 7.1.3.

However, one must also take into account the possibility that the systematics are position dependent, as differential airmass and PSF variations across the field-of-view may become important. In the context of an INT/WFC monitoring campaign in the field of open cluster NGC 6819, Street (2002) implemented a position-dependent zero-point correction. The zero point for each frame (i.e. the median, for a given image i , of the individual deviations $\delta_{ij} = \overline{m_j} - m_{ij}$ from the light curve average or median $\overline{m_j}$ for each star j) is computed not for the entire image but for a square region of the image (the optimal size of which is determined through trial and error).

In the present work, we implemented a similar procedure, but rather than subtracting the zero-point for each region we proceeded, once the zero-points were

computed for each region of an image, to fit a 2-D polynomial to those zero-points as a function of x and y position. The correction applied for each star in each image was then derived from this fit, rather than from the zero-points themselves. It is therefore a smooth function derived from the entire grid of regions rather than a discontinuous set of values. Some care must be taken in the calculation of the zero-points for each region. The size of the regions was set to 100×100 pixels, after some experimentation with different values. This is low compared to the value of 500×500 used by Street (2002), but the APT pixels are much larger than the INT/WFC pixels. The smooth nature of the correction also makes it less sensitive to the size of region used. Only objects classified by the source extraction software as stellar (i.e. circular) and falling within the magnitude range $9 \leq V \leq 16$ are included in the zero-point calculation. Each star is weighted according to the inverse of the overall variance of its light curve. Cases where $|\delta m_{ij}| > 0.25$ are also excluded, as they correspond either to true variable stars or to outlying points in an otherwise constant light curve. A 2nd order polynomial was used. The fits are not improved by increasing the order of the polynomial.

This method was found to yield a small improvement in the overall rms (see Figure 7.4, right panel), reducing the spread about the backbone of points for a given magnitude as well as the number of stars with rms values significantly above the expected value for their magnitude. It made a significant difference to the number of spurious transit candidates with high detection statistics.

While the steps described here should remove any position dependent zero-point effects, it seems that not all the systematic trends are fully removed. Certain groups of light curves, unrelated in position or magnitude, still appear to show common trends. Stars of potentially very different temperature and colour are included in the computation of the correction applied to any given light curves, and this may affect the results, but no colour information is available to – for example – select colour-dependent comparison stars. Other, unforeseen parameters may also affect the systematics. Street (2002), who also observed such residual systematics after applying her position-dependent zero-point correction, implemented a method to identify those light curves whose residuals were most similar to each light curve in turn, and to derive and apply a correction based on that set of light curves. A similar approach has recently been published by Kovács et al. (2004), while a PCA-based approach for the same purpose has been proposed by Tamuz et al. (2005). These methods, or variants thereof, will be investigated in the future to see if they can further improve the quality of the light curves.

After the systematics correction, data points deviating from the light curve median by 5σ or more (σ being the MAD-estimated scatter rather than the standard

deviation) were removed from each light curve. To avoid removing real features, this clipping stage was suppressed in light curves with more than 5 such outliers (which are generally eclipsing binaries or flare stars).

7.2 Transit search

Before entering the transit search itself, an aside on variable star searching is appropriate. A search for variable stars, using the same reduced chi-squared criterion and sine-fitting technique as mentioned in Section 6.3.1.3 for the COROT blind exercise, was carried out on the 5 nights of APT data, though without the post-processing steps described in the previous Section, which had not yet been developed. This variable star search, some of the results of which were published in Hidas et al. (2003), was intended primarily to demonstrate the precision achieved by the APT observations and the potential of the survey, and concerned only the thousand brightest stars. We identified approximately 100 variables, including all those reported by Martín et al. (2004) in the magnitude range covered by the APT data, but also tens of new variables of various types, from low-amplitude irregular variables to eclipsing binaries, multi-periodic pulsating stars and RR Lyrae stars. Variable star identification and classification is an integral part of the scientific exploitation of transit search data, and efforts in this direction are ongoing, but we concentrate in the present Chapter on the transit search itself.

The transit search algorithm described in Section 2.2, modified to search over a range of trial durations as well as periods and epochs, was applied to all light curves. Trial durations ranged from 0.04 to 0.25 days (1 to 6 hours) and 0.8 to 4.0 days, with increments of 0.01 day for both. The trial epoch increment was set to the average interval between consecutive observations (excluding daily gaps), which was 0.0026 days (3.75 min). The parameters of the best transit candidate (corresponding to the highest multiple transit detection statistic S) were recorded for each light curve.

In order to establish an appropriate detection threshold, the cumulative distribution of the multiple transit detection statistic S obtained for each light curve was plotted. The result is shown in Figure 7.5 (black line). It is immediately apparent that the vast majority of the light curves contain at least one event with $S \geq 6$, the threshold derived from Monte Carlo simulations to be optimal for white noise only (see Section 2.2). While there is a distinct tail above $S > 30$, these cases correspond to very deep eclipses ($> 10\%$) which are unlikely to be of planetary origin, and restricting oneself to those would defeat the point of the transit search. There is no clear feature in the distribution to indicate an appropriate cutoff lower down, and many of the cases in the range $10 < S < 30$ correspond to spurious detections due to




Figure 7.5: Cumulative distribution of the multiple transit detection statistics for all stars with $9 < V < 16$. Black line: no filtering; gray line: after filtering of the long term trends.

night-to-night shifts or long-term trends in the light curves.

Given the lack of a natural way to determine an optimal threshold, one can select one that will give the desired number of candidates. Artificial though it may be, this provides a means of making a rapid foray into the data, and the threshold can be relaxed later if wished. But first, it is desirable to try to weed out at least some of the spurious detections automatically. To this end, two approaches were investigated.

The first was to repeat the transit search, preceded by a long-term variability filtering stage. The filter used was a simplified version of the least-squares filter described in Chapter 4, where rather than fitting sine curves up to relatively high frequencies and applying a matched or Wiener filter, one simply fits and subtracts sine-curves up to a given frequency. In effect a highpass filter, this is almost identical to the approach used by team 3 in the COROT blind exercise, though here sine curves of increasing frequency were fitted and subtracted one by one rather than simultaneously. The fitted frequencies were harmonics of the fundamental frequency $1/2T$ (T being the total time span of the light curves), and the highest fitted frequency was chosen to correspond to five times the longest trial duration (i.e. 1.25 day, which ensures the removal of any long term trends and most night-to-night shifts). Note that the dataset is very short, hence only a total of 8 frequencies are fitted in this way.

The resulting detection statistic distribution is shown as the gray line on Figure 7.5. The tail of high S events has disappeared (as the filter removes some signal from very deep eclipses), and the overall distribution has been shifted to lower S . Examination of the 100 light curves with the highest S (which brings us, in this case, down to $S = 10.2$) reveals that long term trends or night-to-night shifts are no longer a

problem, but many spurious detections arise at periods close to 1.0, 1.5 and 2.0 days. These are due to downward trends at the edges of nights, which are still present in many light curves despite the post-processing steps described in Section 7.1.4. We find $\sim 30\%$ of the light curves show transit or eclipse-like dips, and $\sim 10\%$ show other forms of variability (RR Lyrae-like, short-period pulsations, ...). The remaining $\sim 60\%$ either appear to be spurious detections due to night-edge effects, or have insufficient phase coverage to establish confidence in the detections. Excluding periods within 0.02 day of 1.0, 1.5 or 2.0 days, removes two of the shallower eclipse-like dips and one variable while reducing the number of spurious detections to ~ 10 .

Another approach is to examine the parameters of the high S detections in the case where no filter was used, to see if a criterion based on these parameters could separate the true transit-like events from the detections induced by long-term trends or night-to-night shifts. Of the cases with $S > 20$, it is notable that all but one of the spurious detections have durations larger than 0.23 day, while all the transit-like events have durations below 0.17 day. The light curves for the 100 highest S cases with $d < 0.17$ day were thus examined by eye. We found $\sim 40\%$ had transit-like eclipses, 3 light curves showing signs of other variability, 2 due to night-to-night shifts, and the remainder could either be ascribed to night-edge effects, or the phase-coverage of the transit candidate was too poor to assess it.

It is interesting to note that, whether using filtering or not, a number of shallow transit candidates fall at the end of the first night. The fact that many stars simultaneously show a dip at the same time suggests that it is a systematic effect. Such ‘transits’ are also only partially covered, the egress falling after the end of the night. Shallow transits at the end of a night were thus only accepted if more than one transit was seen and the phase coverage of the transit was sufficient.

A final list of 40 transit- or eclipse-like events was compiled by merging the candidates from the two approaches described above that appeared convincing upon visual inspection of their light curves. The resulting list is shown in Table 7.1. Most candidates were detected through both approaches, but 4 were detected with the first approach (filtering) only and 12 with the second approach (no filter) only. Plots of the phase folded light curves for these candidates, as well as examples of some typical spurious detections, are shown in an Appendix at the end of this Chapter.

7.3 Discussion

The transit search tools developed in Chapters 2 & 4 were applied to a typical ground based transit search dataset, and allowed the detection of several tens of transit-like events, 7 of which have depths ≤ 50 mmag. The light curve analysis was carried out

Table 7.1: Transit or eclipse candidates with their parameters.

ID (1)	V (2)	σ (3)	d (4)	p (5)	e (6)	S_M (7)	dV (8)	dF/F (9)	N_{tr} (10)	approach (11)	approach (12)
164	9.963	0.00268	0.123	0.879	0.883	397.880	151	13.0	1	1	b
355	10.875	0.00548	0.123	1.778	1.095	77.712	55	5.0	2	1	b
742	11.675	0.00712	0.071	1.089	1.589	19.670	21	1.9	2	2	2
1007	11.999	0.00877	0.091	0.995	1.076	21.171	18	1.7	5	1	b
1162	12.119	0.01507	0.060	1.024	1.810	88.278	198	16.7	3	1	b
1397	12.295	0.01437	0.144	0.858	1.549	151.343	259	21.3	2	1	b
1640	12.461	0.06844	0.060	0.966	1.131	16.621	126	11.0	4	1	b
1752	12.524	0.01669	0.060	1.106	0.844	18.052	40	3.7	3	1	b
1816	12.568	0.01746	0.154	1.118	1.757	184.181	410	31.5	1	1	2
1965	12.646	0.01015	0.112	1.109	1.773	27.842	41	3.7	1	1	b
2041	12.683	0.01174	0.070	0.879	1.643	28.836	64	5.8	1	1	b
2042	12.682	0.01960	0.164	0.811	0.938	12.821	22	2.2	3	1	2
2272	12.777	0.01104	0.123	0.801	1.410	85.210	131	11.4	1	1	b
2547	12.895	0.06896	0.112	1.309	0.886	15.576	139	12.1	2	1	b
2787	12.983	0.02092	0.133	0.879	0.881	63.703	183	15.6	1	1	b
2829	13.001	0.02622	0.091	0.971	1.089	13.888	30	2.7	5	1	b
2856	13.008	0.04935	0.081	1.008	0.857	14.577	69	6.2	4	1	2
2893	13.030	0.01751	0.050	1.233	0.844	25.226	90	8.0	2	1	b
3498	13.217	0.03943	0.133	1.109	1.778	21.229	112	9.8	1	1	2
3557	13.244	0.09491	0.091	2.061	0.920	11.738	116	10.1	3	2	1
4046	13.382	0.01630	0.112	1.096	1.799	12.242	28	2.6	1	1	2
4095	13.392	0.11796	0.070	1.895	1.058	11.380	164	14.1	3	2	1
4254	13.435	0.01525	0.102	1.212	1.984	13.682	32	2.9	1	2	1
4290	13.440	0.03046	0.133	1.217	0.969	46.500	160	13.8	2	1	b
4389	13.470	0.02209	0.155	0.834	1.120	24.535	55	5.0	3	1	2
4585	13.521	0.04540	0.174	1.075	1.637	33.823	138	12.0	3	2	1
5198	13.658	0.02101	0.112	0.869	0.961	33.480	101	8.9	1	1	b
5355	13.694	0.07742	0.070	1.930	1.016	14.462	131	11.4	3	1	b
5638	13.751	0.03076	0.060	0.820	0.961	16.576	71	6.3	3	1	b
5782	13.781	0.04338	0.124	1.080	0.957	33.349	156	13.4	3	1	b
6340	13.911	0.08166	0.112	0.826	1.246	20.595	200	16.8	2	1	2
6897	14.021	0.10043	0.102	0.997	1.063	21.349	199	16.8	5	1	2
7249	14.095	0.03491	0.081	1.136	0.925	13.715	62	5.6	2	1	2
7295	14.101	0.05452	0.071	1.200	0.878	22.212	182	15.5	2	1	b
7752	14.201	0.05209	0.092	1.049	1.781	13.291	73	6.7	3	1	2
8208	14.291	0.09715	0.081	1.372	1.906	33.079	476	35.5	2	1	b
8209	14.290	0.04938	0.060	1.464	1.951	24.600	203	17.1	2	1	b
9963	14.632	0.07760	0.071	0.988	0.875	21.218	241	20.0	3	1	b
11231	14.918	0.22918	0.050	0.978	1.001	18.914	464	34.8	5	1	b
11566	14.997	0.10667	0.060	1.824	1.063	12.359	196	16.5	2	1	2

Columns: (1) star ID; (2) median magnitude; (3) fractional light curve scatter; (4) transit duration (days); (5) transit period (days); (6) time of start of first transit (days, first data point is at $t = 0.88$ day); (7) detections statistic; (8) transit depth (mmag); (9) fractional transit depth in flux units (%); (10) number of transits actually sampled; (11) approach the parameters were obtained with; (12) approach the event was detected with (b means both).

in just over one working week, i.e. the time span of the observations. This bodes well for the prospects for semi-automated transit searches in such datasets, but a number of issues remain to be addressed.

The UNSW team have independently carried out a transit search in this same dataset, though using all eight nights collected in 2002 and not just the nights with 1×1 rastering (Hidas et al. 2005). Although they did employ a systematics removal technique similar to the 2-D polynomial fitting described in Section 7.1.4, they still found too many false alarms were induced by nightly trends in the data to be able to use automated transit search algorithms, and therefore opted for a visual inspection of each light curve, limiting themselves to the 2–3000 brightest stars. They thus identified 4 transit candidates with depths below 100 mmag, in light curves 306, 742, 2041 and 1965, as well as two lower confidence candidates, in light curves 795 & 252. 742, 2041 and 1965 are amongst our candidates, but 306 was not included because it only gave rise to a detection statistic of $S \sim 8$, below the cut off in both candidate selection approaches we investigated. The UNSW went on to follow-up their 4 candidates, first checking the transit repeated in data from subsequent seasons, then examining Digitized Sky Survey plates to establish the degree of blending in the APT apertures, and finally obtaining higher resolution time-resolved photometry and low resolution spectroscopy. They found the brightest star in the apertures corresponding to light curves 306, 2041 and 1965 to show large amplitude radial velocity variations typical of stellar mass eclipsing binary objects. In the case of light curve 742, the transit occurs across the second brightest star in the APT aperture, and spectroscopic observations showing a double lined profile whose separation varies in phase with the transits suggest a system composed of two nearly identical K7 stars in a grazing eclipsing binary configuration.

Whether my implementation of the post-processing steps was more successful, or whether I simply accepted a need to weed through a larger number of false alarms, I was able to apply a mostly automated transit search tool to these data. This allowed me to search all the light curves and not only those of the brightest 2–3000 stars (though of course the chances of detecting a planetary transit are much lower in the light curves of fainter stars). I identified, in addition to 3 of the candidates found by Hidas et al. (2005), a number of shallow transit candidates (light curves 1007, 1752, 2042, 2829, 4046 and 4254). Light curve 1007 has relatively poor phase coverage and may have been rejected by the UNSW team on that basis, though it has been included here because there do seem to be a few points belonging to the transit egress. It will be very interesting to see whether the transits repeat at the predicted times in the 2003 and 2004 data. Light curve 1752 shows clear signs of out-of-eclipse variations, and the eclipses are thus unlikely to be due to a

planet, but the light curve shape is intriguing. Although the period identified for light curve 2042 is clearly incorrect, one can see from the phase-folded light curve that it is an eclipsing binary with visible secondary eclipses. Light curve 2829 is an interesting case with a clean transit which may be flat bottomed, and is certainly worthy of follow-up. The reliability of cases 4046 and 4254, which are much fainter and have larger rms, is harder to establish, as only one transit was observed in each case. Should additional transits be found in the 2003 and 2004 data, these would also make interesting candidates.

Hidas et al. (2005) also found 9 deep eclipsing binary (W Uma and Algol) systems and 31 other variable stars. We identified 6 of the EBs (light curves 1162, 1640, 2272, 2547, 2787 & 2856). The other 3 were all W Uma type systems which are almost continually in eclipse. For those, the assumptions (on which the transit search algorithm is based) that the mean out-of-transit level is equal to the overall light curve mean, and that the time in-transit is short compared to the total light curve duration, break down. In practice, the eclipse minima are detected but the detection statistic is below the threshold used here ($S \sim 10$ in all three cases), because the light curve scatter is high. Failure to detect these systems is not considered as a failure of the algorithm, as it was not designed specifically to detect such signals.

However, the 3 eclipsing binaries and the one transit candidate of Hidas et al. (2005) that were not included were only slightly below our detection threshold. This highlights the fact that, given the current state of the post-processing and transit search process, examining the 100 light curves with the highest S is likely to miss many interesting candidates. Ideally, one would wish to examine all light curves with $S > 6$, but in the present case that would include 1000 cases, the vast majority of which are likely to be false alarms. Additional work must therefore concentrate on reducing the false alarm rate before a relatively complete automated transit search can be carried out.

Three principal sources of false alarms³ were encountered: long term trends, night-to-night shifts and night-edge effects (generally at the end, rather than the beginning, of each night). The first two are effectively removed by a simple high-pass filter implemented via sine-curve fitting, but at the cost of removing part of real transit-like signals and thus a loss of sensitivity. The third suggests that the systematics correction applied can be improved. Diagnostic tests to understand the origin of the uncorrected systematics are underway. There appears to be a strong gradient in the summed light curve residuals as a function of y -position at one edge of the

³By false alarms, we mean here detections of events that turn out not to be transit-like, as opposed to detections of transit-like events that turn out not to be planetary (e.g. eclipsing binaries). Short period variable stars, which sometimes give rise to high transit detection statistics, are also not considered as "false alarms" here because they are indeed variable in a way which *should* trigger the algorithm.

CCD (this could be due, for example, to detector artifacts or vignetting), which is not completely removed by the systematics correction applied. A number of potential improvements to the systematics correction are under investigation, but the most effective may be to exclude stars falling in the problematic part of the CCD. Alternative systematics correction methods from the literature are also under review.

It is important to note that, while some degree of visual examination of candidates' light curves will always be necessary, the goal of any automated transit search system is to minimise this need. The amount of visual inspection carried out in the present Chapter – where $\sim 60\%$ of the light curves examined turned out to be spurious candidates – is certainly suboptimal.

The dataset investigated here was very short, and some of the problems encountered in this analysis can be expected to be less important for longer datasets. In particular, sensitivity to spurious periods close to one day or aliases thereof should decrease. NGC 6633 was observed during several runs in 2003 and 2004 with the APT, and the post-processing of these data is underway. If any of the transit candidates which are listed here remain after the analysis of the data from the 2004 season, follow-up observations will be undertaken, starting with higher resolution multi-band photometry to identify possible blends and provide a preliminary assessment of cluster membership.

Appendix: Phase-folded light curve plots

Figures 7.6 to 7.13 contain plots of the normalised phase folded light curves for all the transit or eclipse candidates. There is one row per transit candidate. In each case the left panel shows the full phase-folded light curve, the right panel the phase folded light curve around the transit. The red line in both panels shows the box-shaped transit model corresponding to the reported parameters. The light curve ID is indicated in the top of the left panel. The transit duration, period and epoch (in days) are listed at the bottom of the left panel, the transit depth (in mmag and in %) at the bottom of the right panel. Note that the unfiltered light curves are plotted in all cases, even when the detection was made after filtering, as the filter sometimes distorts the transit shape. In several cases (e.g. ID 2547) the detected period is clearly incorrect, but there is a transit-like event nonetheless. In such cases, the phase-folded light curve is plotted with the detected period, as the plots are automatically produced by the program implementing the detection algorithm.

Figures 7.14 to 7.16 contain similar plots for some representative spurious detection cases (long term variations, night-to-night shift and night-edge effects respectively). Those figures contain an additional panel with the unfolded light curve, which

helps highlight the spurious nature of those candidates.

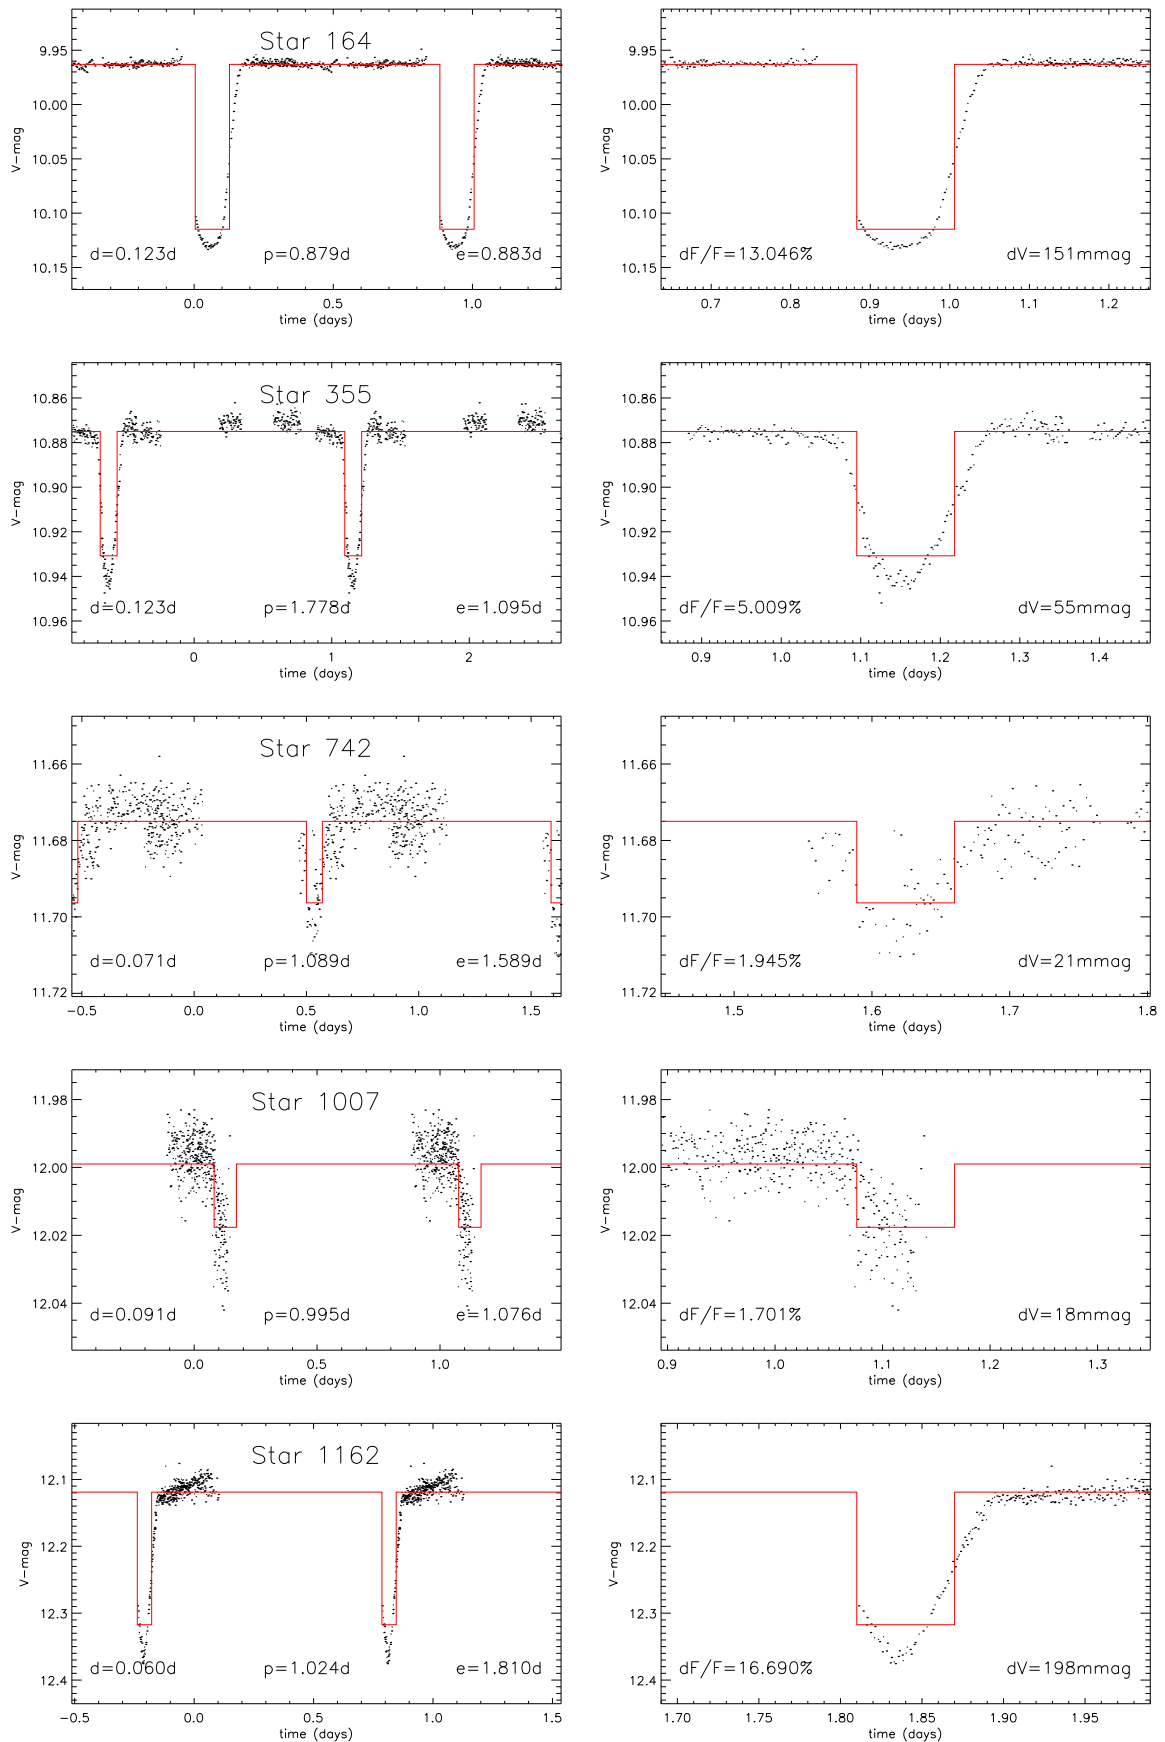


Figure 7.6: Light curves of transit / eclipse candidates – page 1.

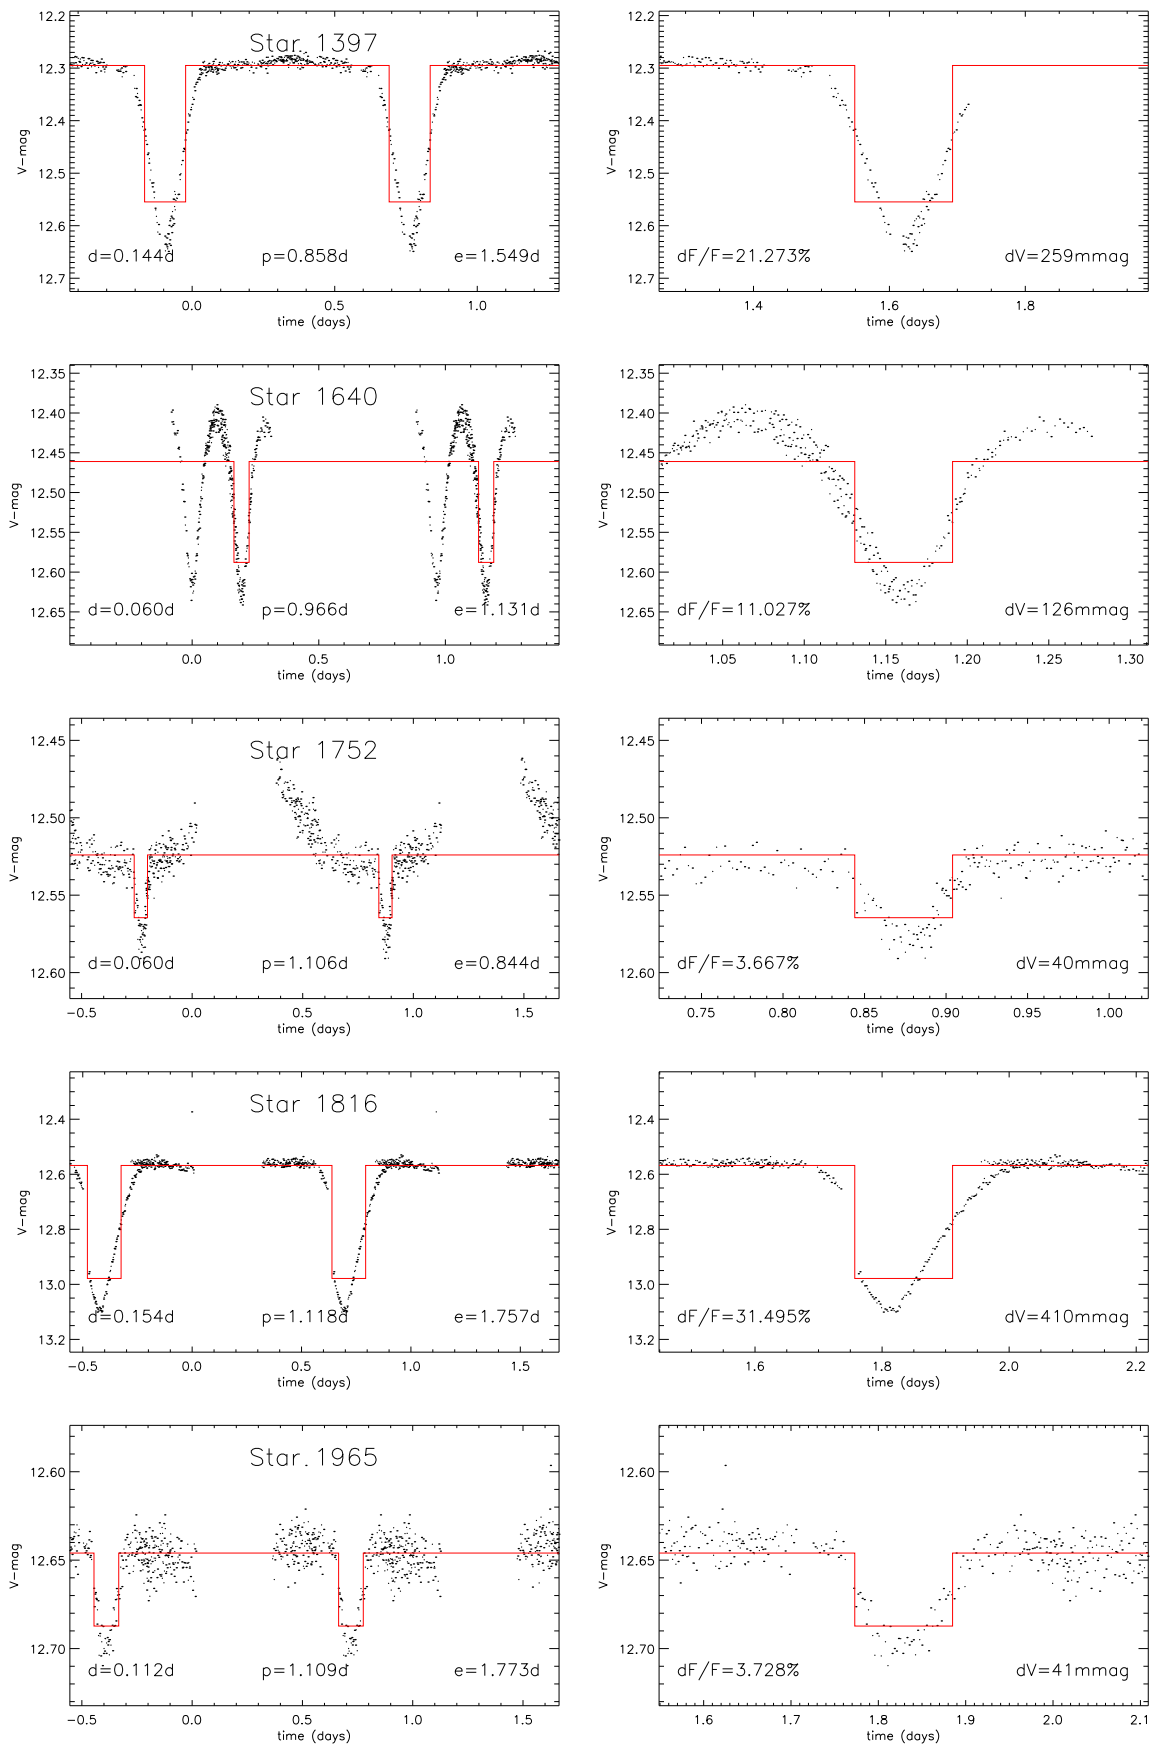


Figure 7.7: Light curves of transit / eclipse candidates – page 2.

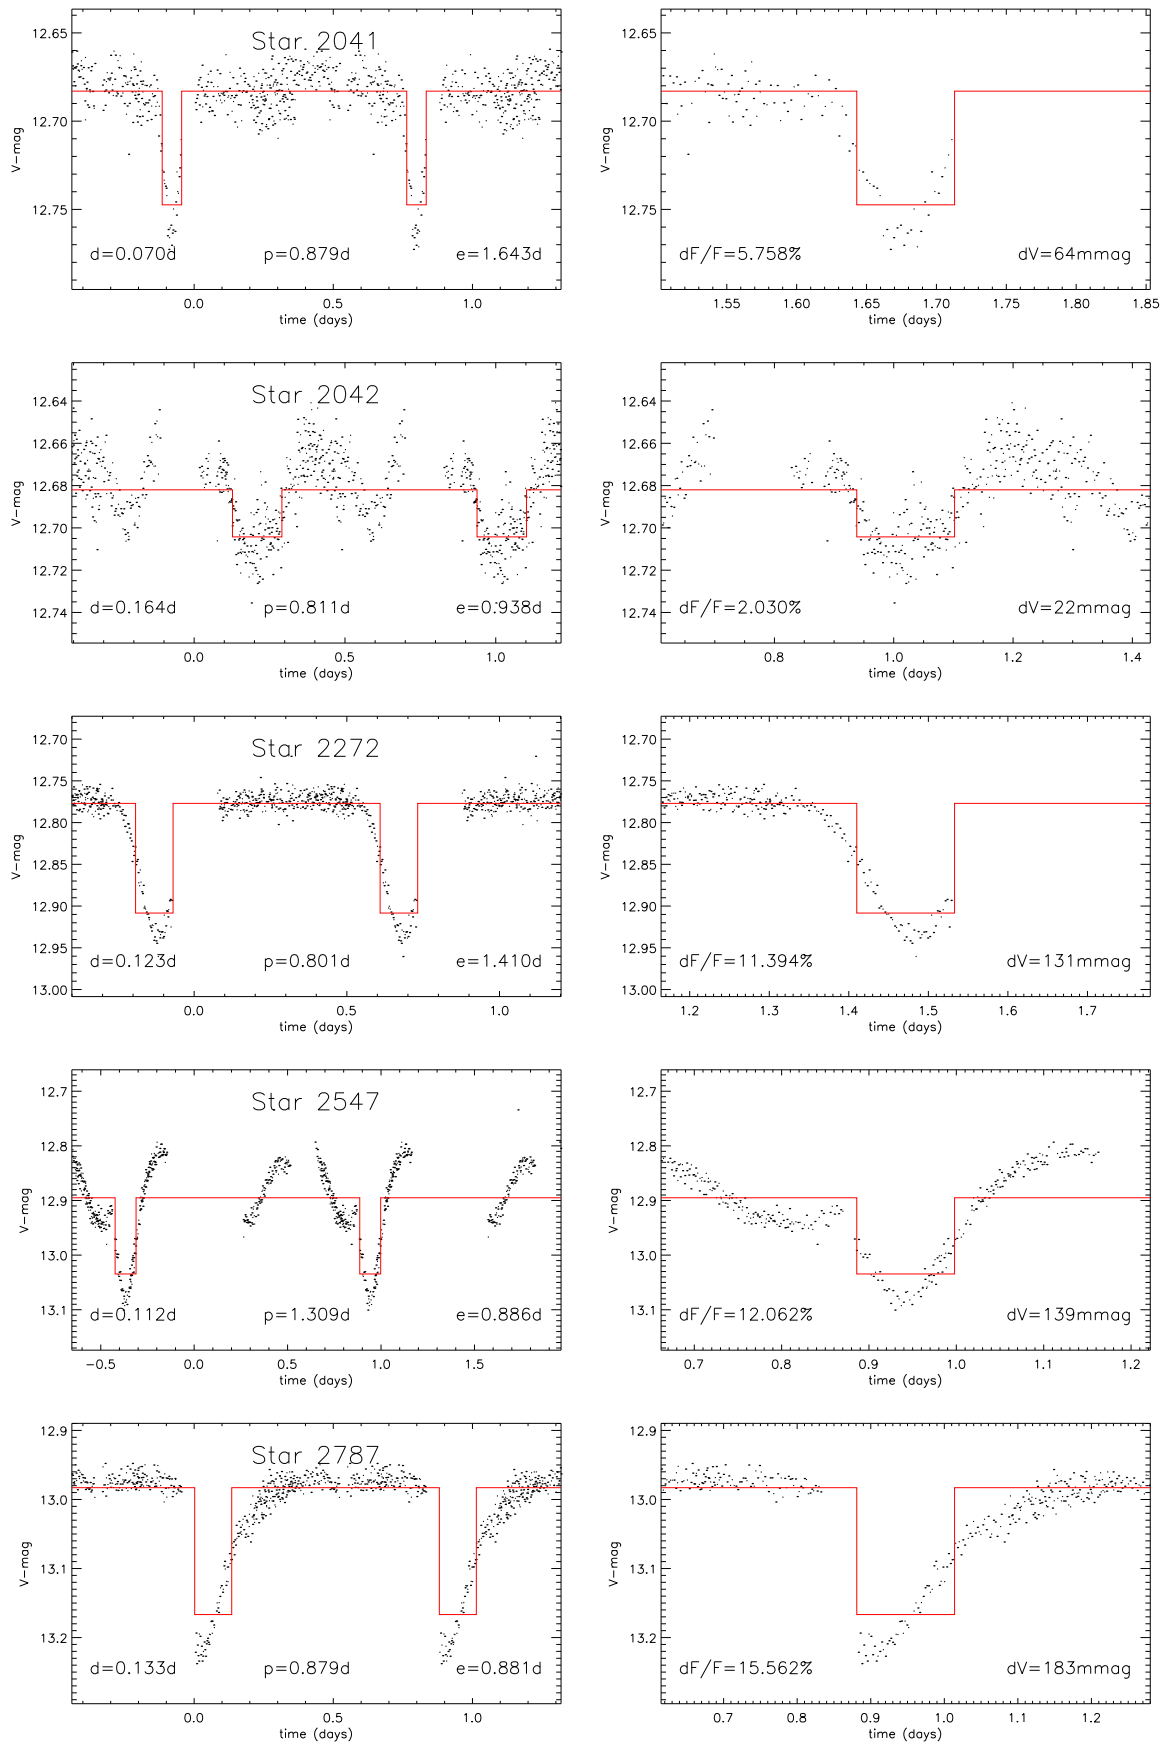


Figure 7.8: Light curves of transit / eclipse candidates – page 3.

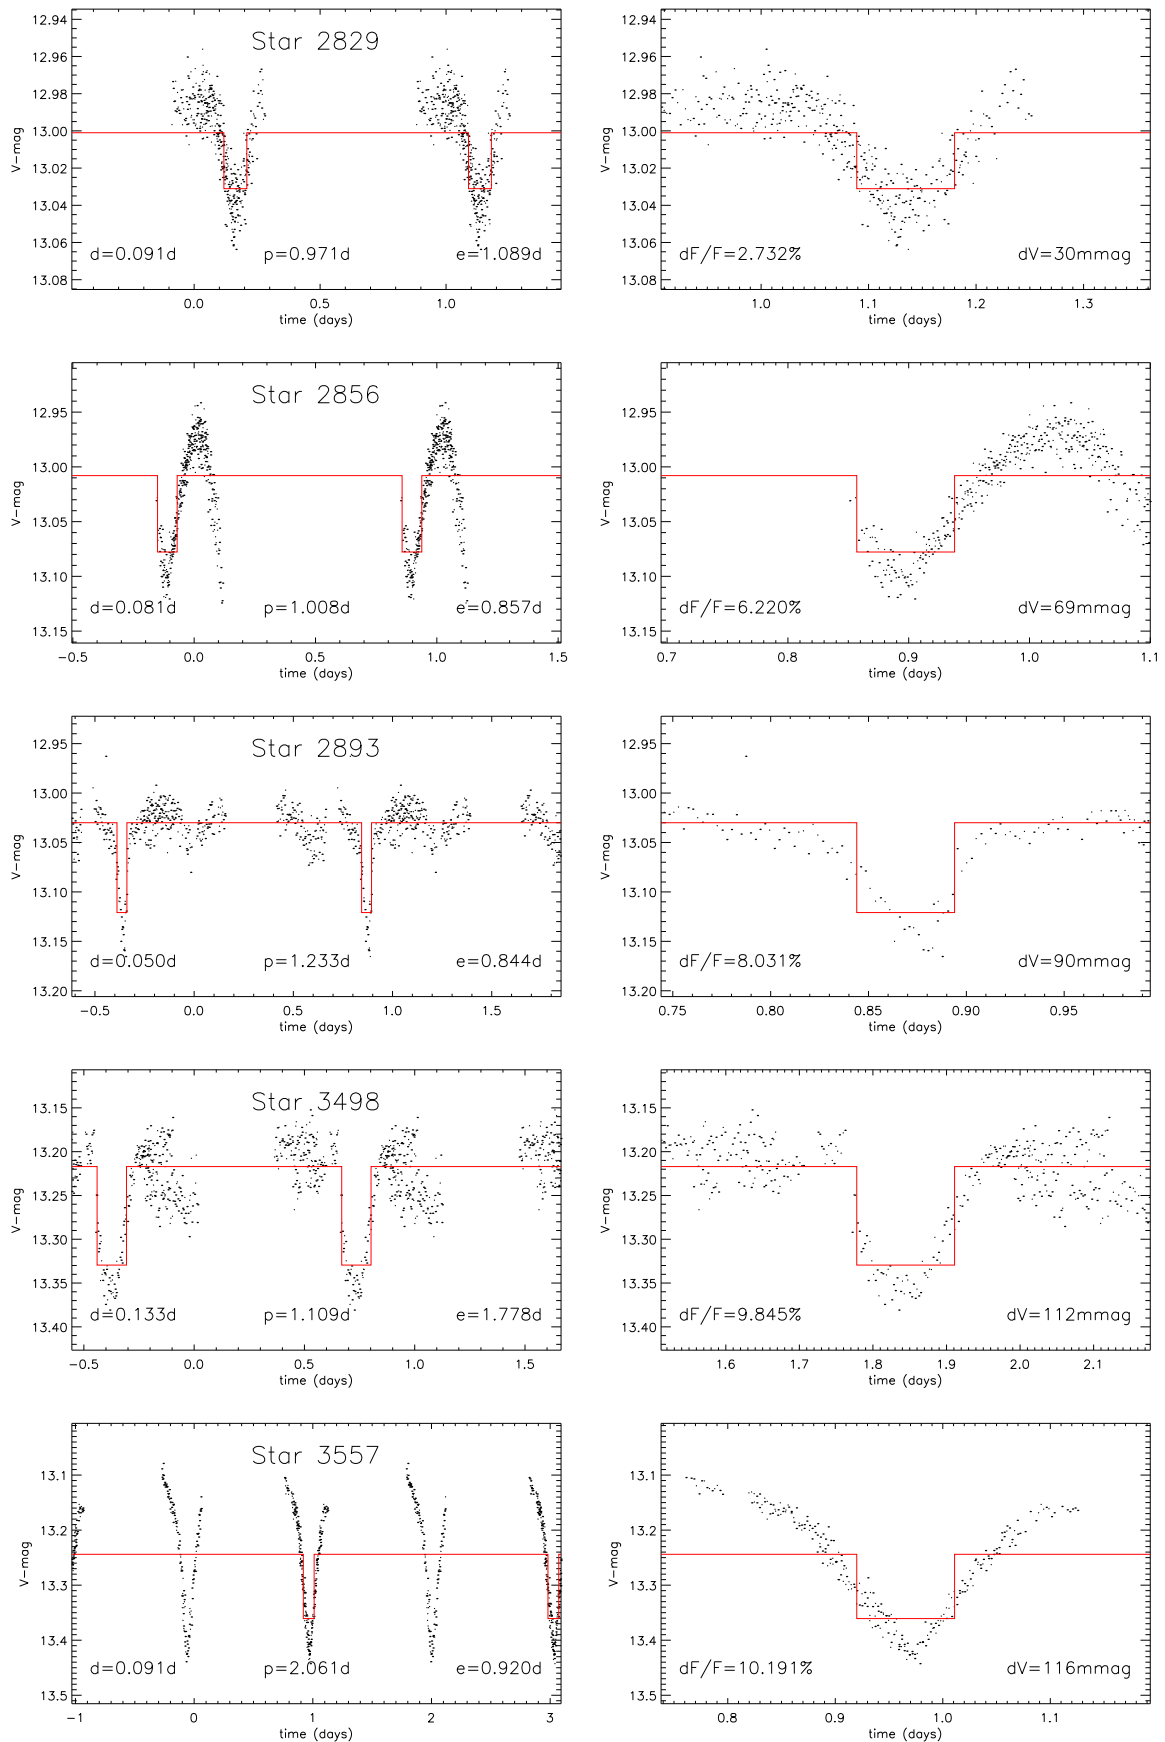


Figure 7.9: Light curves of transit / eclipse candidates – page 4.

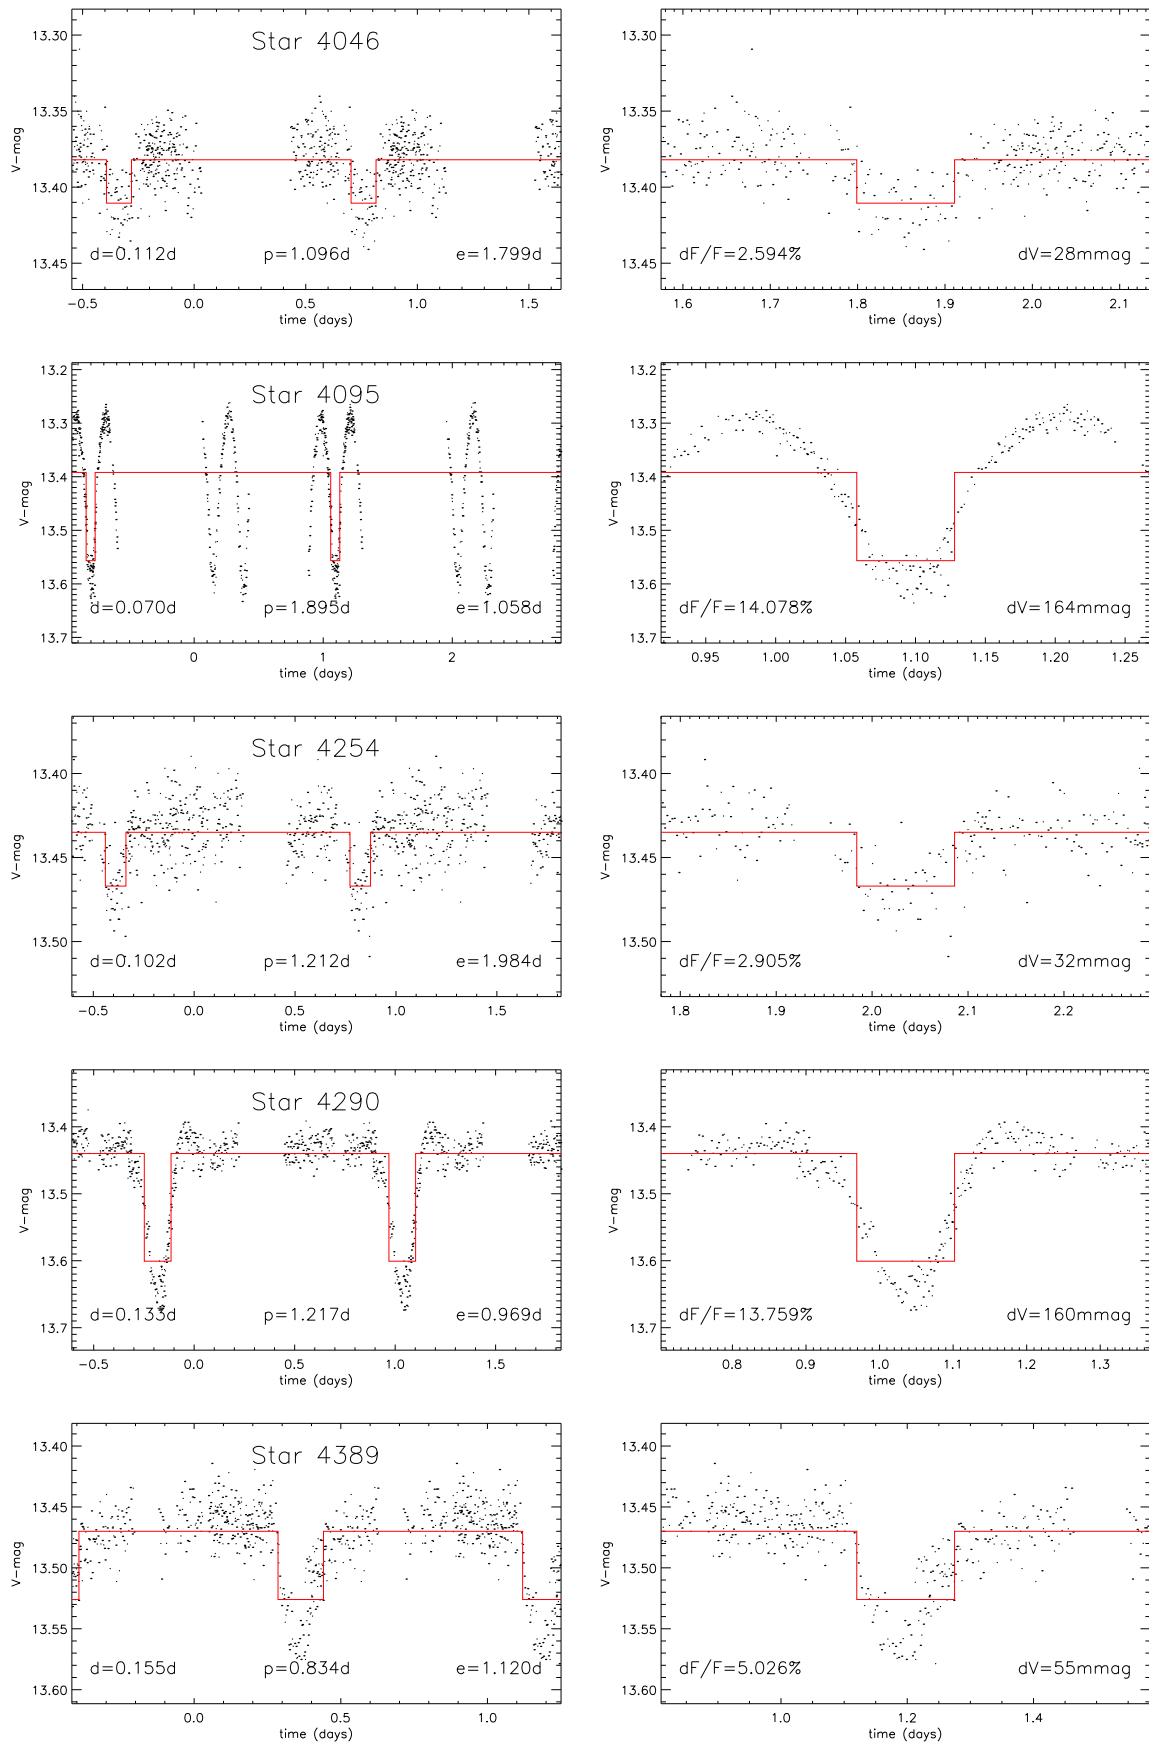


Figure 7.10: Light curves of transit / eclipse candidates – page 5.

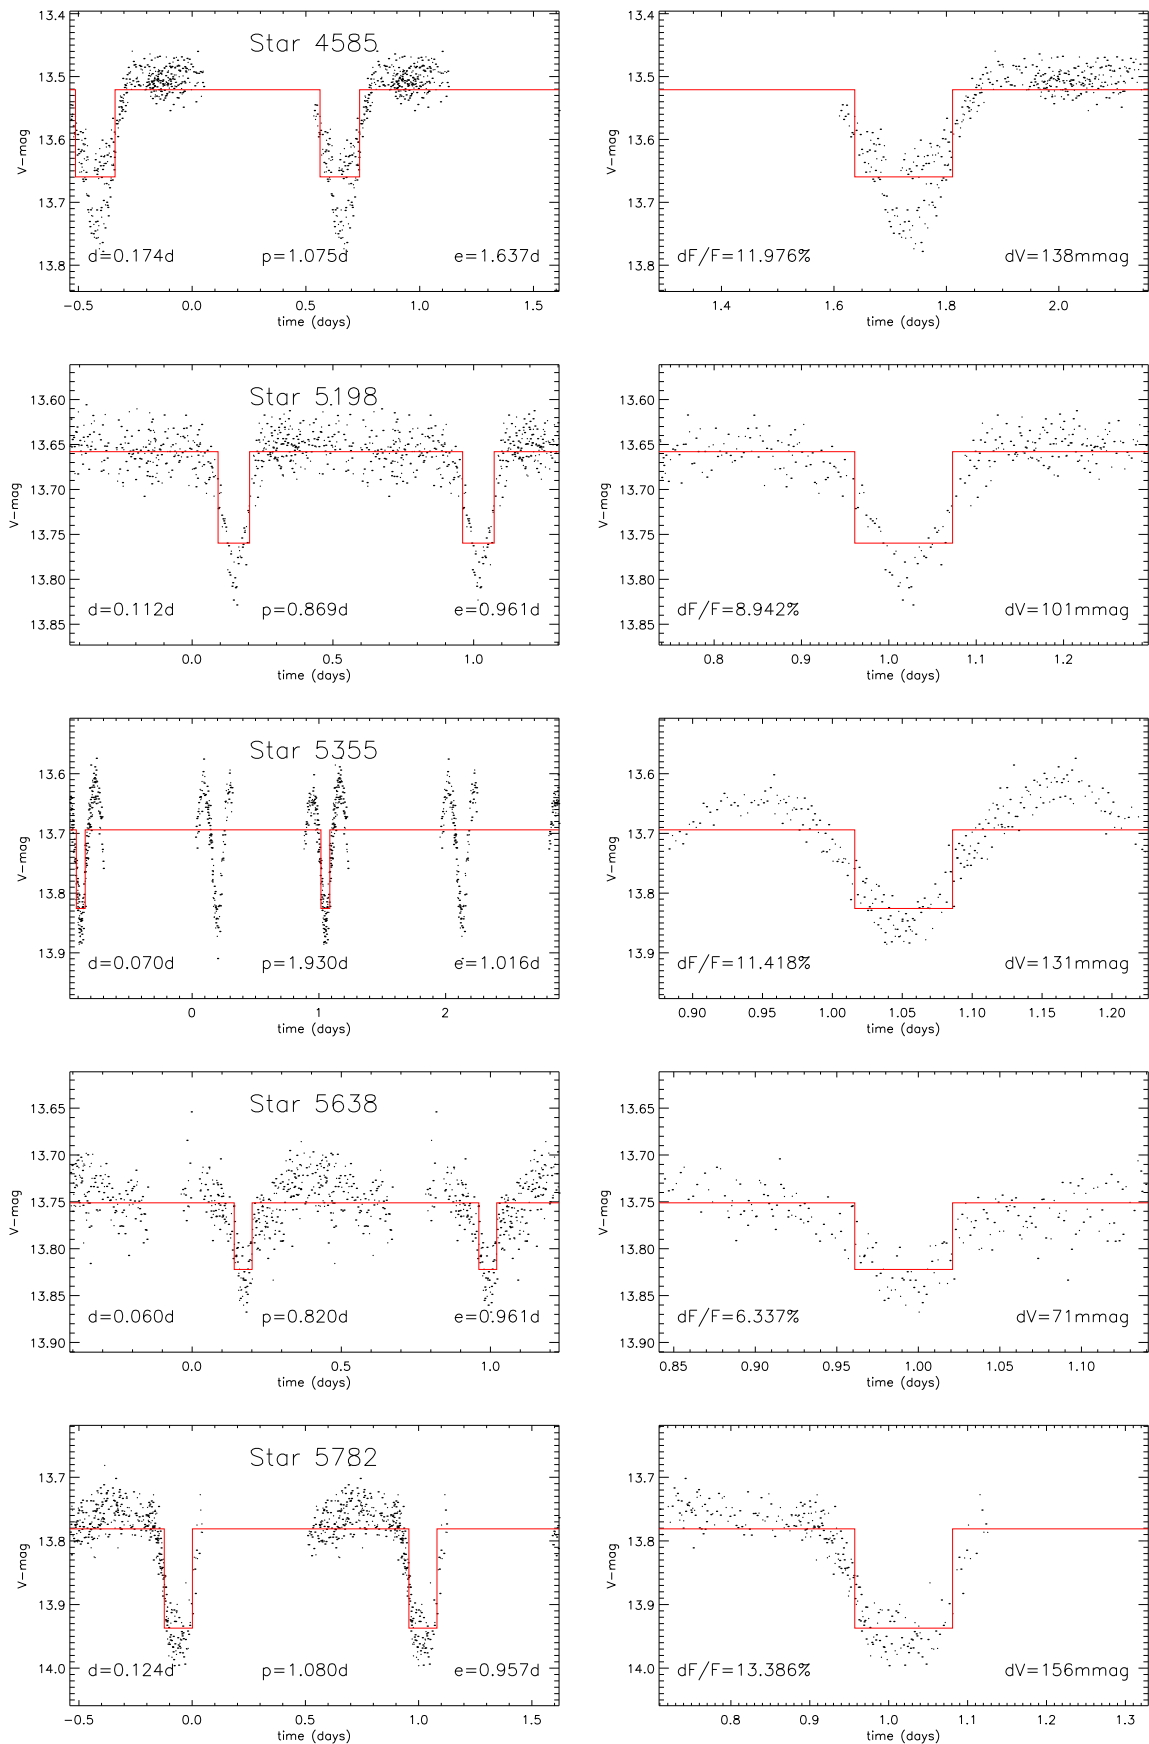


Figure 7.11: Light curves of transit / eclipse candidates – page 6.

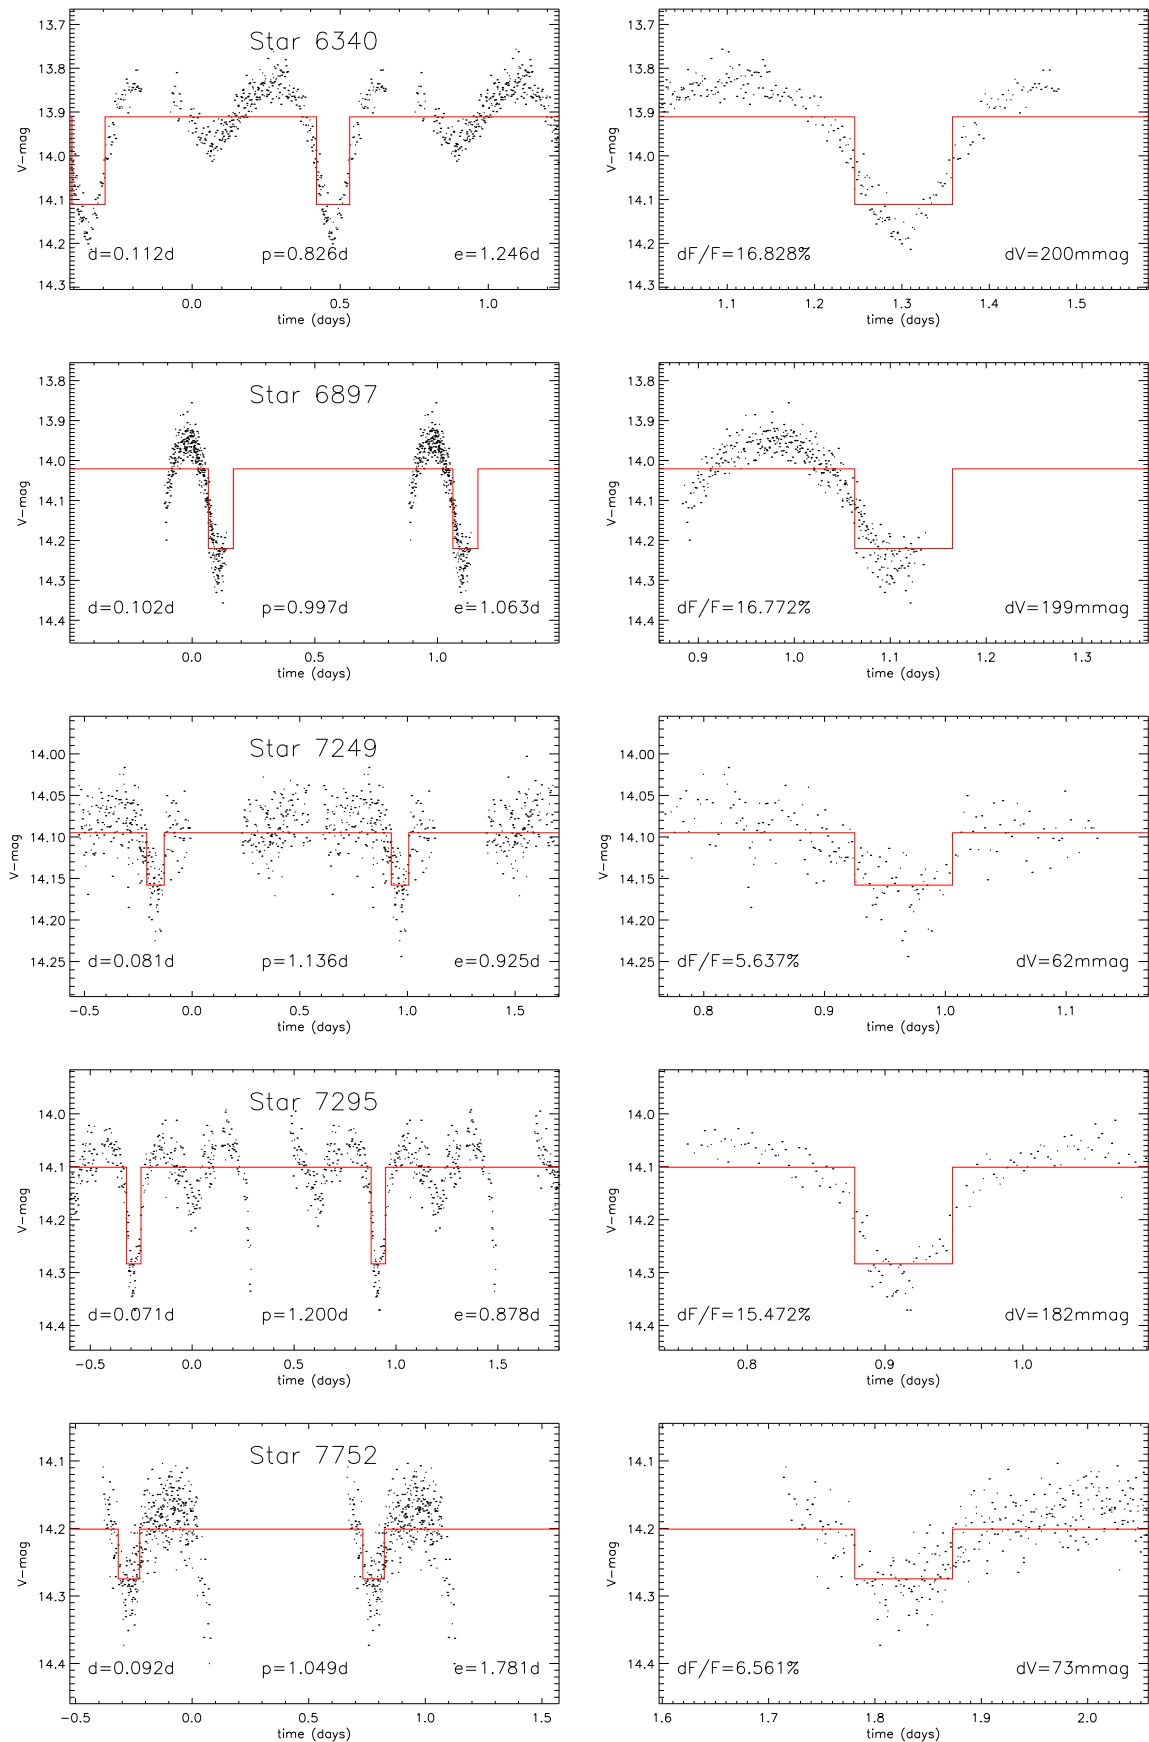


Figure 7.12: Light curves of transit / eclipse candidates – page 7.

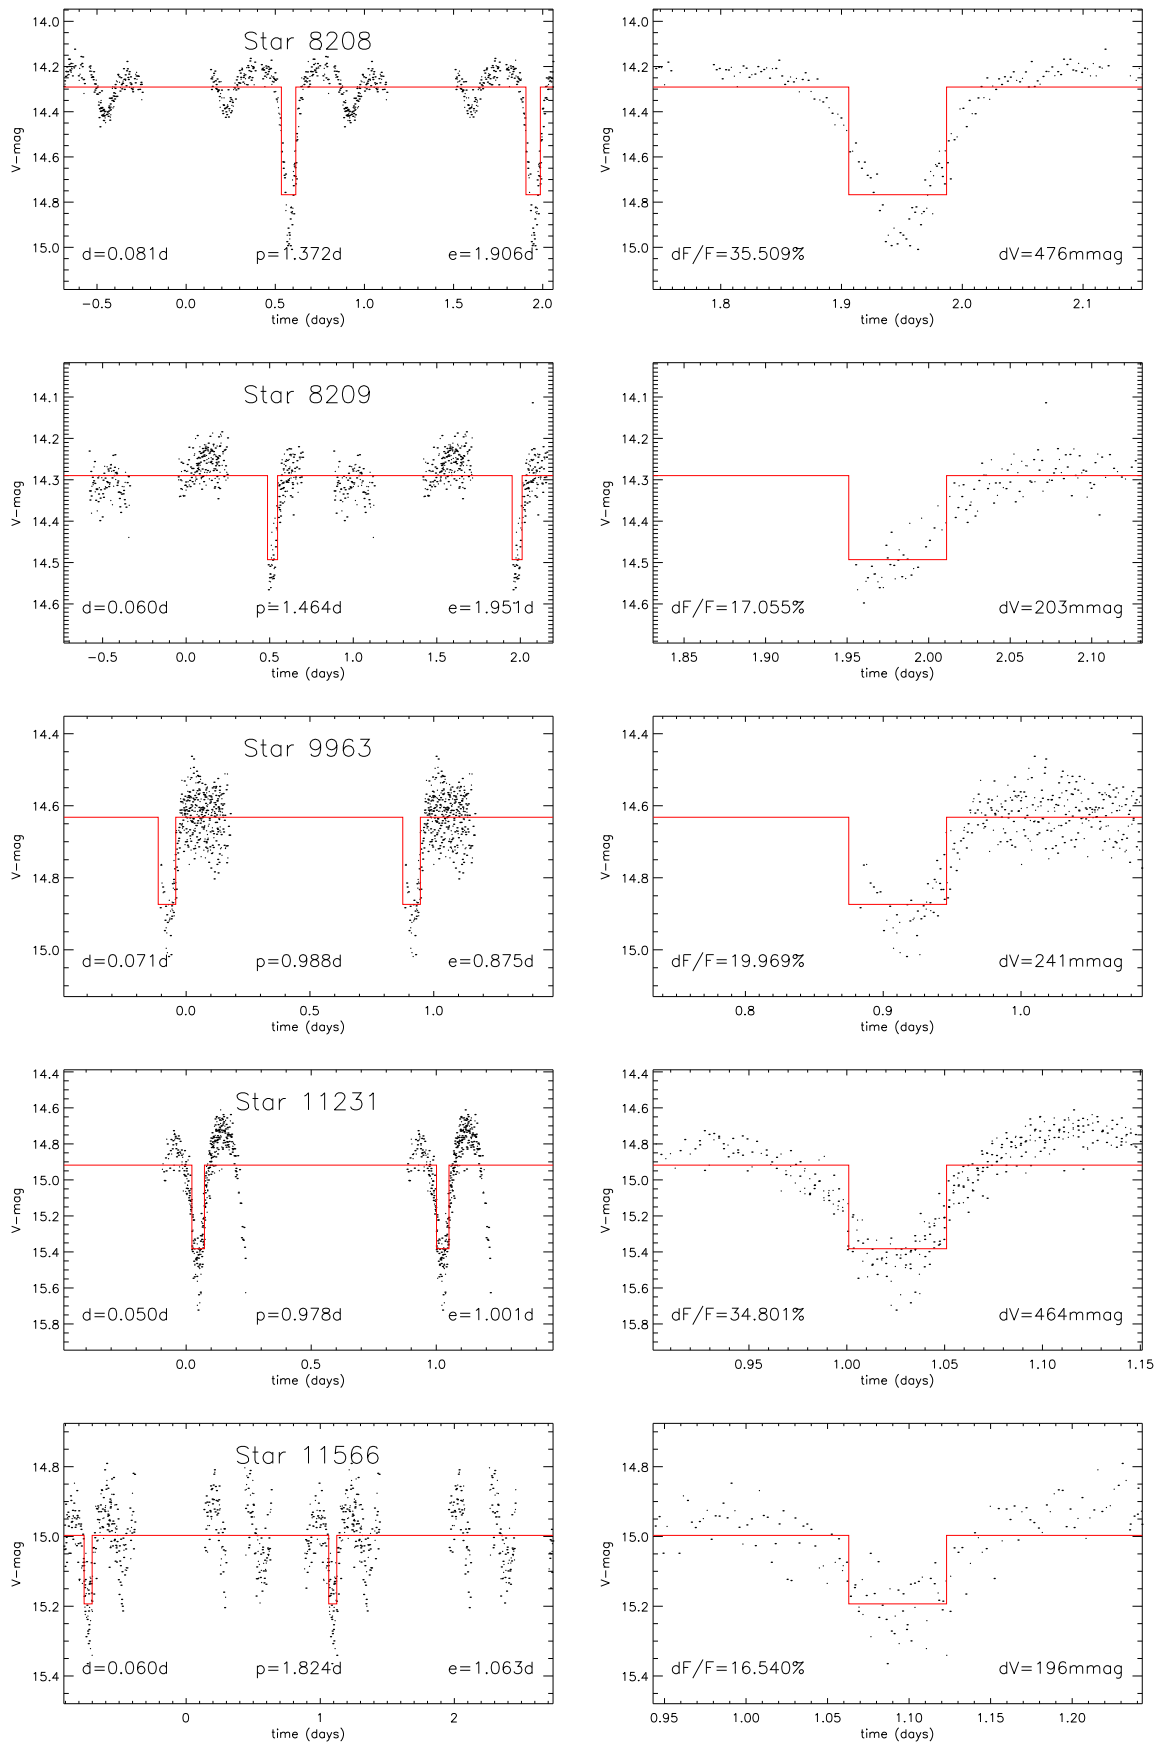


Figure 7.13: Light curves of transit / eclipse candidates – page 8.

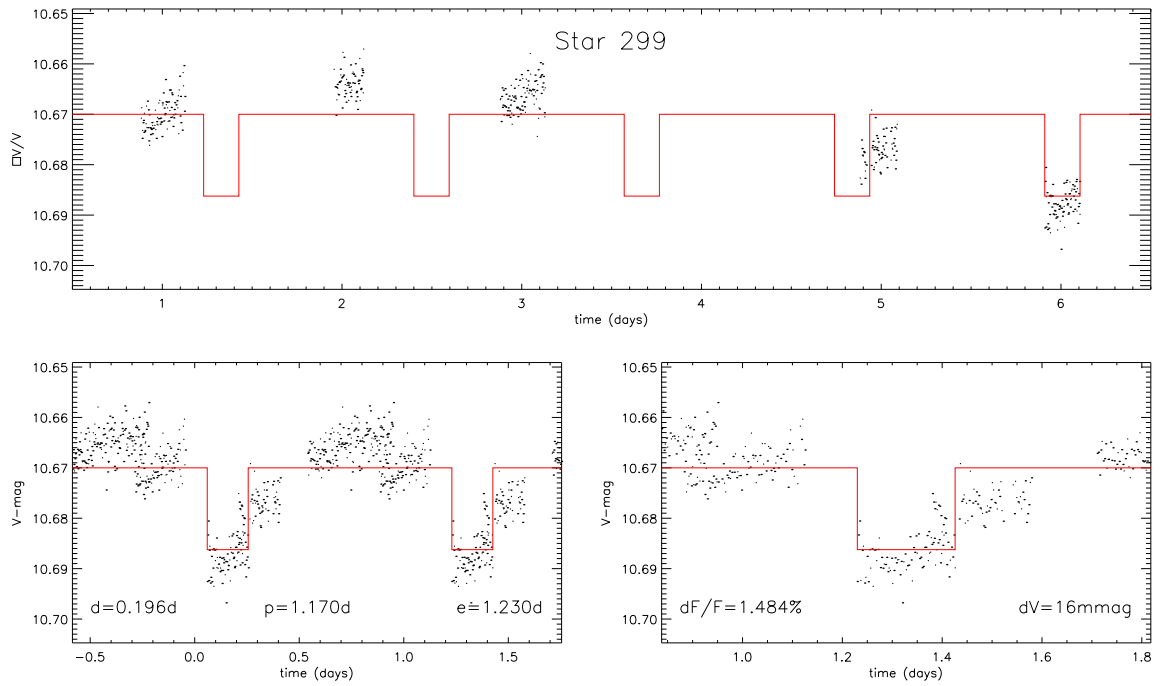


Figure 7.14: Spurious candidate example: long-term variations.




Figure 7.15: Spurious candidate example: night-to-night shifts.




Figure 7.16: Spurious candidate example: night-edge effects.

Chapter 8

Conclusions

The detection and study of planets outside the solar system is one of the most active areas of astrophysics today. Amongst the detection methods currently implemented from the ground, the transit method has a number of advantages, the most notable being the simultaneous monitoring of many thousands of stars and the possibility of measuring the planet radius and inclination directly from the light curve. In combination with the radial velocity method, it yields mass measurements free of the mass-inclination degeneracy, and a direct probe of the planetary mass-radius relation. The first few detections thus made have recently received confirmation, and many more are expected in the coming years. From space, the transit method is expected to be the first to allow the detection of terrestrial and habitable planets.

8.1 Data analysis tools for planetary transit searches

The detection of planetary transits in stellar light curves poses a number of challenges due to the rarity of the transits, their brief and shallow nature, the many thousands of light curves that have to be searched, and the noise sources that affect them, in particular intrinsic low amplitude stellar variability on timescales of tens of minutes to weeks, which all stars are expected to display to some extent.

In response to these challenges, a number of data analysis tools have been developed. These include pre-processing filters to minimise the impact of stellar micro-variability and a transit search algorithm based on chi-squared minimisation with a box-shaped transit model. An empirical micro-variability model was also developed to test these tools and assist in the selection of the best target stars for transit searches, by allowing the simulation of realistic light curves for stars of different spectral and activity level.

8.2 Applications

The micro-variability model, filters and transit search algorithms were applied to estimate the impact of stellar micro-variability on transit detectability in *Eddington* and COROT data by exploring the star-planet parameter space in a systematic manner and identifying promising areas. The results of these simulations have already had some influence on some aspects of the observing strategy and design of these missions. I am planning to carry out more detailed and extensive simulations when the target fields have been selected and more fully characterised, and the instrument models are at a more advanced stage.

These tools were also applied in the context of a large collaboration within the COROT Exo-planet Working Group, as part of the first COROT blind exercise. This exercise helped identify detection limits for COROT. It has also shown that the performance of our filters and transit search algorithm compares well with that of other published methods, and also highlighted areas where room for improvement remain and where either alterations to the existing tools or the development of additional ones are desirable.

A transit search was also run in a 5 night dataset from the the University of New South Wales (UNSW) program on the 0.45 m Automated Patrol Telescope in Siding Springs Observatory. This analysis, though very preliminary, helped bring to light a number of issues not foreseen when testing the tools developed in this thesis on simulated space-based data, in particular the importance of systematics removal and the difficulty of obtaining stable light curves in the presence of daily interruptions and variable atmospheric extinction. Several interesting transit candidates with depths consistent with a planetary origin were also identified, and though will be followed up in the near future.

This work is ongoing, additional data from the APT and from the SuperWASP project begin under investigation at present. These tools will also be applied to a novel ground based transit search project, which aims to find transiting planets orbiting pre-main sequence stars by monitoring star forming regions. The first target for this study is the Orion Nebula Cluster (ONC), using the Wide Field camera on the Isaac Newton telescope (INT), for which the data collection will start in November 2004. In the longer term, the methods will be applied to data from COROT, *Kepler* and if applicable *Eddington*. Their applicability to data from other missions such as GAIA, which will provide high precision photometric monitoring of an even larger number of stars but with much sparser time sampling, will also be investigated.

The micro-variability model will also be used to study and improve our understanding of the causes of stellar micro-variability by comparing it with upcoming

space-based data, in the context of a collaboration within the COROT Additional Program Working Group, focused on understanding the causes of hours-timescale stellar variability.

8.3 Future improvements

There is room for improvement of the aforementioned tools in several areas. This is true in particular for the stellar micro-variability model, where empirical constraints are currently so scarce. This situation is about to change, as data from the MOST satellite starts to become public from the fall of 2004. These data will be used to provide additional constraints and make adjustments to the model as they become available, but it will be limited to a few stars. From 2006, COROT data will provide constraints for a much larger number and wider variety of stars. At the high activity end, the amplitude of variations is larger, and ground-based monitoring, such as the project targeting the ONC from the INT, will provide useful constraints. The possibility of simulating light curves in different bandpasses will also be investigated. Data from the Spectral Irradiance Monitor (SIM) on the SORCE satellite (launched last year), which provides 6-hourly measurements of the irradiance of the Sun as a function of wavelength from the IR to the UV, will be used for this purpose as soon as it becomes available.

Improvements to the filters will include a fuller investigation of the least-squares fitting plus matched filter approach, focusing on the application to particularly active stars and the reconstruction of the stellar signal, and refinements to the way the filter width is chosen in relation to the transit duration for the iterative non-linear filter. However, we will have to wait for the COROT data to become available to know how well these filters will perform on real data.

The performance of the transit detection algorithm itself has been investigated relatively thoroughly on model data. The next step will be to perform Monte Carlo simulations in which artificial transit signals are added to real data (to conserve the noise properties). These will be useful in particular to adjust the candidate selection process, which seems to be the area in which most progress can be expected.

8.4 Additional tools

The applications investigated to date have highlighted a number of areas where the natural progression is to develop a new, additional 'stage' in the modular framework established for the analysis of transit search data.

Light curves from transit search projects constitute a treasure trove of information on stellar variability. A very simple procedure for the identification of periodic variable stars using sine-curve fitting was used in the context of the COROT blind exercise, but much more could be done. At the very least, a procedure should be developed to flag potentially interesting variables – on the basis of a global parameter such as the reduced chi-squared of the light curve – for further study using dedicated tools. The next step is variable star classification, for which a variety of methods are available: comparison to templates, clustering algorithms, self-organising maps, neural networks etc. . . In the first instance, a study of the published literature on these methods will be undertaken to identify the most promising for transit search data.

Another area which has only briefly been touched upon in the present thesis is the detailed analysis of the light curve and the characterisation of transit candidates. It is the first step in the follow up of the candidates, allowing the exclusion of a number of ‘impostor’ signals as well as the derivation of the physical parameters of the system. Measuring the transit parameters in the presence of non-white noise is a non-trivial problem. If filters are used to remove the noise, one must ensure that they do not modify the transit shape. An alternative may be to attempt to reconstruct the transit signal directly from the noisy light curve using prior knowledge of the transit period and epoch.

To summarise, a set of data analysis tools for transit searches has been developed and tested on simulated data. The next step is to apply them to real data. Each new application will no doubt raise new problems and prompt improvements to the existing tools or further developments, as well as hopefully leading the discovery of exo-planets and variable stars. By the time data from the space-based missions which originally motivated this work becomes available, it is hoped that as many as possible of the potential problems will have been identified and dealt with, so that the data may be analysed as speedily and efficiently as possible.

Bibliography

Aigrain, S. & Favata, F. 2002, *A&A*, 395, 625

Aigrain, S., Favata, F., & Gilmore, G. 2004, *A&A*, 414, 1139

Aigrain, S. & Irwin, M. 2004, *MNRAS*, 350, 331

Alonso, R., Brown, T. M., Belmonte, J. A., et al. 2002, in Proc. 1st *Eddington Workshop: Stellar Structure and Habitable Planet Finding*, ESA SP-485, ed. F. Favata, I. W. Roxburgh, & D. Galadi, 245

Alonso, R., Brown, T. M., Torres, G., et al. 2004a, *ApJL*, in press

Alonso, R., Deeg, H. J., Brown, T. M., & Belmonte, J. A. 2004b, in Proc. 2nd *Eddington Workshop: Stellar structure and habitable planet finding*, ESA SP-538, ed. F. Favata & S. Aigrain, 255

Andersen, B. N., Appourchaux, T., Crommelynck, D., et al. 1998, in *IAU Symp.*, Vol. 181, *Sounding solar and stellar interiors*, ed. J. Provost & F. X. Schmider, 147

Andersen, B. N., Leifsen, T., & Toutain, T. 1994, *Sol. Phys.*, 152, 247

Baglin, A. & the COROT Team. 2003, *Adv. Sp. Res.*, 345

Bakos, G., Noyes, R. W., Kovács, G., et al. 2004, *PASP*, 116, 266

Baliunas, S., Sokoloff, D., & Soon, W. 1996, *ApJ*, 457, L99

Baliunas, S. L., Donahue, R. A., Soon, W. H., et al. 1995, *ApJ*, 438, 269

Baraffe, I., Selsis, F., Chabrier, G., et al. 2004, *A&A*, 419, L13

Barnes, J. R., Collier-Cameron, A., Unruh, Y. C., Donati, J. F., & Hussain, G. A. J. 1998, *MNRAS*, 299, 904

Barry, D. C., Cromwell, R. H., & Hege, E. K. 1987, *ApJ*, 315, 264

Bayes, T. 1764, *An essay toward solving a problem in the doctrine of chances*

- Belokurov, V., Evans, N. W., & Du, Y. L. 2003, *MNRAS*, 341, 1373
- Belokurov, V., Evans, N. W., & Le Du, Y. 2004, *MNRAS*, 352, 233
- Bordé, P., Rouan, D., & Léger, A. 2003, *A&A*, 405, 1137
- Borucki, W., Koch, D., Boss, A., et al. 2004, in *Proc. 2nd Eddington Workshop, Stellar structure and habitable planet finding*, ESA-SP 538, ed. F. Favata & S. Aigrain, 177
- Borucki, W. J., Caldwell, D., Koch, D. G., et al. 2001, *PASP*, 113, 439
- Borucki, W. J., Koch, D. G., Dunham, E. W., & Jenkins, J. M. 1997, in *ASP Conf. Ser., Vol. 119, Planets beyond the Solar system and the next generation of space missions*, ed. D. Soderblom, 153
- Boss, A. P. 1996, *Lunar and Planetary Science*, 27, 139
- . 2000, *ApJ*, 536, L101
- Bouchy, F., F. P., Santos, N. C., et al. 2004, *A&A*, submitted, astro-ph/0404264
- Bouvier, J., Forestini, M., & Allain, S. 1997, *A&A*, 326, 1023
- Brault, J. W. & White, O. R. 1971, *A&A*, 13, 169
- Bretthorst, G. L. 1988, *Lecture Notes in Statistics, Vol. 48, Bayesian spectrum analysis and parameter estimation* (New York: Springer-Verlag)
- Brown, T. M. 2001, *ApJ*, 553, 1006
- . 2003, *ApJ*, 593, L125
- Bruntt, H., Grundahl, F., Tingley, B., et al. 2003, *A&A*, 410, 323
- Burke, C. J., Depoy, D. L., Gaudi, B. S., & Marshall, J. L. 2003, in *ASP Conf. Ser. 294: Scientific Frontiers in Research on Extrasolar Planets*, ed. D. Deming & S. Seager, 379
- Burleigh, M. R., Clarke, F. J., & Hodgkin, S. T. 2002, *MNRAS*, 331, L41
- Butler, R. P. & Marcy, G. W. 1996, *ApJ*, 464, L153
- Butler, R. P., Marcy, G. W., Fischer, D. A., et al. 1999, *ApJ*, 526, 916
- Butler, R. P., Marcy, G. W., Williams, E., Hauser, H., & Shirts, P. 1997, *ApJ*, 474, L115
- Carpano, S., Aigrain, S., & Favata, F. 2003, *A&A*, 401, 743
- Cayrel de Strobel, G. 1990, *Memorie della Societa Astronomica Italiana*, 61, 613

- Charbonneau, D., Brown, T. M., Latham, D. W., & Mayor, M. 2000, *ApJ*, 529, L45
- Charbonneau, D., Brown, T. M., Noyes, R. W., & Gilliland, R. L. 2002, *ApJ*, 568, 377
- Chauvin, G., Lagrange, A.-M., Dumas, C. and Zuckerman, B., et al. 2004, *A&AL*, in press
- Chiang, E. I. 2003, *ApJ*, 584, 465
- Cochran, W. D., Hatzes, A. P., Butler, R. P., & Marcy, G. W. 1997, *ApJ*, 483, 457
- Deeg, H. 1999, Universal Transit Modeller,
www.iac.es/galeria/hdeeg/idl.hans.lib/utm/
- Deeg, H. J. 2004, in Proc. 2nd *Eddington* workshop, Stellar structure and habitable planet finding, ESA SP-538, ed. F. F. & S. Aigrain, 231
- Deeg, H. J., Garrido, R., & Claret, A. 2001, *New Astron.*, 6, 51
- Defaÿ, C. 2001, PhD thesis, Université d'Aix Marseille III
- Defaÿ, C., Deleuil, M., & Barge, P. 2001, *A&A*, 365, 330
- Dobbs-Dixon, I., Lin, D. N. C., & Mardling, R. A. 2004, *ApJ*, 610, 464
- Domingo, V., Sanchez, L., Appourchaux, T., & Andersen, B. 1998, in IAU Symp., Vol. 185, New eyes to see inside the Sun and stars, ed. F.-L. Deubner, J. Christensen-Dalsgaard, & D. W. Kurtz, 111
- Doyle, L. R., Deeg, H. J., Kozhevnikov, V. P., et al. 2000, *ApJ*, 535, 338
- Drake, A. J. 2003, *ApJ*, 589, 1020
- Dreizler, S., Rauch, T., Hauschildt, P., et al. 2002, *A&A*, 391, L17
- Duncan, D. K., Vaughan, A. H., Wilson, O. C., et al. 1991, *ApJS*, 76, 383
- Eggenberger, A., Udry, S., & Mayor, M. 2004, *A&A*, 417, 353
- Favata, F. 2004, in Proc. 2nd *Eddington* workshop, Stellar structure and habitable planet finding, ESA SP-538, ed. F. F. & S. Aigrain, 3
- Favata, F. & the *Eddington* Science Team. 2000, *Eddington* Assessment Study Report, ESA-SCI(2000)8
- Fischer, D., Valenti, J. A., & Marcy, G. W. 2004, in Proc. IAU Symp. 219: Stars as Suns: activity, evolution and planets, ed. A. K. Dupree

- Fligge, M., Solanki, S. K., & Unruh, Y. C. 2000, *A&A*, 353, 380
- Franck, S., von Bloh, W., Bounama, C., et al. 2002, in *ASP Conf. Ser.*, Vol. 269, *The evolving Sun and its influence on planetary environments*, ed. B. Montesinos, A. Giménez, & E. F. Guinan, 261
- Frink, S. 2003, in *Proc. Toward other Earths: Darwin/TPF and the search for extrasolar terrestrial planets*, ESA SP-539, ed. M. Fridlund & T. Henning, 413
- Frohlich, C., Andersen, B., Appourchaux, T., et al. 1997, *Sol. Phys.*, 170, 1
- Garcia-Lopez, R. J., Rebolo, R., Beckman, J. E., & McKeith, C. D. 1993, *A&A*, 273, 482
- Gilliland, R. L. & Brown, T. M. 1992, *PASP*, 104, 582
- Gilliland, R. L., Brown, T. M., Guhathakurta, P., et al. 2000, *ApJ*, 545, L47
- Gilliland, R. L., Brown, T. M., Kjeldsen, H., et al. 1993, *AJ*, 106, 2441
- Gilmore, G. 1979, *MNRAS*, 187, 389
- Giménez, A., Mas-Hesse, M. J., Jamar, C., et al. 1999, *Astrophysical Letters Communications*, 39, 347
- Goldreich, P. & Sari, R. 2003, *ApJ*, 585, 1024
- Goldreich, P. & Tremaine, S. 1979, *ApJ*, 233, 857
- . 1980, *ApJ*, 241, 425
- Gonzalez, G. 1997, *MNRAS*, 285, 403
- . 1998, *A&A*, 334, 221
- Gray, D. F. & Nagel, T. 1989, *ApJ*, 341, 421
- Gregory, P. C. 1999, *ApJ*, 520, 361
- Gregory, P. C. & Loredó, T. J. 1992, *ApJ*, 398, 146
- Guinan, E. F., McCook, G. P., Wright, S., & Schneider, J. 1996, *BAAS*, 28, 1312
- Guinan, E. F. & Ribas, I. 2002, in *ASP Conf. Ser.*, Vol. 269, *The evolving Sun and its influence on planetary environments*, ed. B. Montesinos, A. Giménez, & E. F. Guinan, 85
- Guis, V. & Barge, P. 2004, *A&A*, in prep.

- Harvey, J. W. 1985, in *Future missions in Solar, heliospheric and space plasma physics*, ESA SP-235, ed. E. J. Rolfe & B. B., 199
- Henry, G. W., Baliunas, S. L., Donahue, R. A., Fekel, F. C., & Soon, W. 2000b, *AJ*, 531, 415
- Henry, G. W., Baliunas, S. L., Donahue, R. A., Soon, W. H., & Saar, S. H. 1996, *ApJ*, 474, 503
- Henry, G. W., Marcy, G. W., Butler, R. P., & Vogt, S. S. 2000a, *ApJ*, 529, L41
- Hidas, M. G., Ashley, M. C., Webb, J. K., et al. 2003, in *Proc. IAU Symp. 219: Stars as Suns: Activity, Evolution and Planets*, ed. A. K. Dupree
- Hidas, M. G., Ashley, M. C. B., Webb, J. K., et al. 2005, *MNRAS*, in prep.
- Hiremath, K. M. 2002, *A&A*, 386, 674
- Hoaglin, D. C., Mostellar, F., & Tukey, J. W., eds. 1983, *Understanding robust and exploratory data analysis* (New York Chichester: John Wiley)
- Horne, K. 2003, in *Proc. ASP Conf. Ser. 294: Scientific frontiers in research on extrasolar planets*, ed. D. Deming & S. Seager, 361
- Jenkins, J. M. 2002, *ApJ*, 575, 493
- Jenkins, J. M., Caldwell, D. A., & Borucki, W. J. 2002, *ApJ*, 564, 495
- Jenkins, J. M., Doyle, L. R., & Cullers, D. K. 1996, *Icarus*, 119, 244
- Kasting, J. F., Whitmire, D. P., & Reynolds, R. T. 1993, *Icarus*, 101, 108
- Kawaler, S. D. 1989, *ApJ*, 343, L65
- Kay, S. 1998, *Fundamentals of statistical signal processing: Detection theory* (Upper Saddle River: Prentice-Hall PTR)
- Kjeldsen, H., Bedding, T. R., Frandsen, S., & Dall, T. H. 1999, *MNRAS*, 303, 579
- Kjeldsen, H. & Tingley, B. 2004, in *Proc. 2nd Eddington Workshop: Stellar structure and habitable planet finding*, ESA SP-538, ed. F. Favata & S. Aigrain, 335
- Koen, C. & Lombard, F. 2002, in *Proc. 1st Eddington Workshop, Stellar structure and habitable planet finding*, ESA SP-485, ed. F. Favata, I. W. Roxburgh, & D. Galadi, 159
- Konacki, M., Torres, G., Jha, S., & Sasselov, D. 2003a, *Nature*, 421, 507
- Konacki, M., Torres, G., Sasselov, D. D., & Jha, S. 2003b, *ApJ*, 597, 1076

- Konacki, M., Torres, G., Sasselov, D. D., et al. 2004, *ApJ*, 609, L37
- Königl, A. 1991, *ApJ*, 370, L39
- Kotredes, L., Charbonneau, D., O'Donovan, F. T., & Looper, D. L. 2003, *American Astronomical Society Meeting*, 203
- Kovács, G., Bakos, G., & Noyes, R. W. 2004, *MNRAS*, 644
- Kovács, G., Zucker, S., & Mazeh, T. 2002, *A&A*, 391, 369
- Krishnamurthi, A., Pinsonneault, M. H., Barnes, S., & Sofia, S. 1997, *ApJ*, 480, 303
- Krivova, N. A., Solanki, S. K., Fligge, M., & Unruh, Y. C. 2003, *A&A*, 399, L1
- Léger, A., Selsis, F., Sotin, C., et al. 2004, *Icarus*, 169, 499
- Lammer, H., Selsis, F., Ribas, I., et al. 2003, *ApJ*, 598, L121
- Lanza, A., Rodonò, M., & Pagano, I. 2004, *A&A*, in press
- Lanza, A. F., Rodonò, M., Pagano, I., Barge, P., & Llebaria, A. 2003, *A&A*, 403, 1135
- Latham, D. W. 2003, in *Proc. ASP Conf. Ser. 294: Scientific Frontiers in Research on Extrasolar Planets*, ed. D. Deming & S. Seager, 409
- Lissauer, J. J. 1993, *ARA&A*, 31, 129
- Lister, T. & the SuperWASP collaboration. 2004, in *Proc. 13th Cambridge Workshop on Cool Stars, Stellar Systems and the Sun*, ed. F. Favata (ESA Publications division), in press
- Lockwood, G. W., Skiff, B. A., & Radick, R. R. 1997, *ApJ*, 485, 789
- Lockwood, G. W., Thompson, D. T., Radick, R. R., et al. 1984, *PASP*, 96, 714
- Mallén-Ornelas, G., Seager, S., Yee, H. K. C., et al. 2003, *ApJ*, 582, 1123
- Mandel, K. & Agol, E. 2002, *ApJ*, 580, L171
- Marcy, G. W. & Butler, R. P. 1995, *BAAS*, 27, 1379
- . 1996, *ApJ*, 464, L147
- . 1998, *ARA&A*, 36, 57
- Martín, S., Alonso, R., Rodríguez, E., et al. 2004, in *Proc. 2nd Eddington Workshop: Stellar structure and habitable planet finding*, ESA SP-538, ed. F. Favata & S. Aigrain, 349

- Marín, S. & Rodríguez, E. 2002, in Proc. 1st *Eddington* Workshop: Stellar structure and habitable planet finding, ESA SP-485, ed. F. Favata, I. W. Roxburgh, & D. Galadi, 303
- Mayor, M. & Queloz, D. 1995, *Nature*, 378, 355
- Messina, S. & Guinan, E. F. 2002, *A&A*, 393, 225
- Messina, S., Pizzolato, N., Guinan, E. F., & Rodonò, M. 2003, *A&A*, 410, 671
- Messina, S., Rodonò, M., & Guinan, E. F. 2001, *A&A*, 366, 215
- Middelkoop, F. 1992, *A&A*, 107, 31
- Mizerski, T. & Bejger, M. 2002, *Acta Astronomica*, 52, 61
- Moutou, C., Aigrain, S., Auvergne, M., et al. 2004a, *A&A*, in prep.
- Moutou, C., Pont, F., Bouchy, F., & Mayor, M. 2004b, *A&A*, in press
- Noyes, R. W. 1983, in *IAU Symp.*, Vol. 102, Solar and stellar magnetic fields: origins and coronal effects, 133
- Noyes, R. W., Hartmann, L. W., Baliunas, S. L., Duncan, D. K., & Vaughan, A. H. 1984, *ApJ*, 279, 763
- Papaloizou, J. & Lin, D. N. C. 1984, *ApJ*, 285, 818
- Paulson, D. B., Saar, S. H., Cochran, W. D., & Hatzes, A. P. 2002, *AJ*, 124, 572
- Pepe, F., Mayor, M., Queloz, D., et al. 2004, *A&A*, 423, 385
- Pepper, J., Gould, A., & Depoy, D. L. 2003, *Acta Inf.*, 53, 213
- Perryman, M. A. C. 2000, *Rep. Prog. Phys.*, 63, 1209
- Perryman, M. A. C., Brown, A. G. A., Lebreton, Y., et al. 1998, *A&A*, 331, 81
- Pollack, J. B. 1984, *ARA&A*, 22, 389
- Pollack, J. B., Hubickyj, O., Bodenheimer, P., et al. 1996, *Icarus*, 124, 62
- Pont, F., Bouchy, F., Queloz, D., et al. 2004, *A&AL*, in press
- Posson-Brown, J., Latham, D. W., Stefanik, R. P., et al. 2000, *BAAS*, 32, 1483
- Prosser, C. F., Shetrone, M. D., Dasgupta, A., et al. 1995, *PASP*, 107, 211
- Radick, R. R., Lockwood, G. W., Skiff, B. A., & Baliunas, S. L. 1998, *ApJS*, 118, 239
- Radick, R. R., Lockwood, G. W., Skiff, B. A., & Thompson, D. T. 1995, *ApJ*, 452, 332

- Radick, R. R., Thompson, D. T., & Lockwood, G. W. et al. 1987, *ApJ*, 321, 459
- Rasio, F. A. & Ford, E. B. 1996, *Science*, 274, 954
- Rauer, H., Eislöffel, J., Erikson, A., et al. 2004, *PASP*, 116, 38
- Rhode, K. L., Herbst, W., & Mathieu, R. D. 2001, *AJ*, 122, 3258
- Robichon, N. & Arenou, F. 2000, *A&A*, 355, 295
- Robin, A. C., Reylé, C., Derrière, S., & Picaud, S. 2003, *A&A*, 409, 523
- . 2004, *A&A*, 416, 157
- Rutten, R. G. M. 1984, *aa*, 130, 353
- Saar, S. H., Butler, R. P., & Marcy, G. W. 1998, *ApJ*, 498, L153
- Santos, N. C., Bouchy, F., Mayor, M., et al. 2004, *A&A*, in prep.
- Santos, N. C., Israelian, G., Mayor, M., Rebolo, R., & Udry, S. 2003, *A&A*, 398, 363
- Schatzman, E. 1962, *Annales d'Astrophysique*, 25, 18
- Schou, J. & Buzasi, D. L. 2001, in *SOHO 10/GONG 2000 Workshop: Helio- and Asteroseismology at the Dawn of the Millennium*, ESA SP-464, ed. P. L. Pallé, 391
- Seager, S. & Mallén-Ornelas, G. 2003, *ApJ*, 585, 1038
- Seleznyov, A. D., Solanki, S. K., & Krivova, N. A. 2003, in *Earths: DARWIN/TPF and the Search for Extrasolar Terrestrial Planets*, ESA SP-539, 589–+
- Selsis, F. 2002, in *Proc. 1st European Workshop on Exo-Astrobiology*, ESA SP-518, ed. H. Lacoste, 365
- Sirko, E. & Paczyński, B. 2003, *ApJ*, 592, 1217
- Skumanich, A. 1972, *ApJ*, 171, 565
- Soon, W., Frick, P., & Baliunas, S. 1999, *ApJ*, 510, L135
- Stassun, K. G. & Terndrup, D. 2003, *PASP*, 115, 505
- Street, R. 2002, PhD thesis, University of St Andrews
- Street, R. A., Horne, K., Lister, T. A., et al. 2002, *MNRAS*, 330, 737
- . 2003, *MNRAS*, 340, 1287
- Tamuz, O., Mazeh, T., & Zucker, S. 2005, *MNRAS*, in press

- Terquem, C. 2004, Halting planetary migration, E-proceedings of the 3rd COROT Planet Workshop: Close-in exoplanets: the star-planet connection
- Tingley, B. 2003a, *A&A*, 403, 329
- . 2003b, *A&A*, 408, L5
- Tinney, C. G., McCarthy, C., Jones, H. R. A., et al. 2002, *MNRAS*, 332, 759
- Torres, G., Konacki, M., Sasselov, D. D., & Jha, S. 2004, *ApJ*, 609, 1071
- Torres, G., Stefanik, R. P., & Latham, D. W. 1997, *ApJ*, 474, 256
- Trampedach, R., Christensen-Dalsgaard, J., Nordlund, A., & Stein, R. F. 1998, in *The First MONS Workshop: Science with a Small Space Telescope*, ed. H. Kjeldsen & T. R. Bedding (Aarhus Universiteit), 59
- Udalski, A., Paczynski, B., Zebrun, K., et al. 2002a, *Acta Inf.*, 52, 1
- Udalski, A., Pietrzynski, G., Szymanski, M., et al. 2003, *Acta Inf.*, 53, 133
- Udalski, A., Szewczyk, O., Zebrun, K., et al. 2002c, *Acta Inf.*, 52, 317
- Udalski, A., Zebrun, K., Szymanski, M., et al. 2002b, *Acta Inf.*, 52, 115
- Udry, S. 2000, in *Etoiles doubles, Ecole CNRS de Goutelas XXIII*, ed. D. Egret, J. L. Halbwachs, & J. M. Hameury
- Udry, S., Mayor, M., Clausen, J. V., et al. 2003a, *A&A*, 407, 679
- Udry, S., Mayor, M., & Santos, N. C. 2003b, *A&A*, 407, 369
- Van Hamme, W. 1993, *ApJ*, 106, 2096
- Vidal-Madjar, A., Désert, J.-M., Lecavelier des Etangs, A., et al. 2004, *ApJ*, 604, L69
- Vidal-Madjar, A., Lecavelier des Etangs, A., Désert, J.-M., et al. 2003, *Nature*, 422, 143
- Walker, G., Matthews, J., Kuschnig, R., et al. 2003, *PASP*, 115, 1023
- Weber, E. J. & Davis, L. J. 1967, *ApJ*, 148, 217
- Weiss, W. W., Aerts, C., Aigrain, S., et al. 2004, in *Proc. 2nd Eddington Workshop, Stellar structure and habitable planet finding*, ESA SP-538, ed. F. Favata & S. Aigrain, 435
- Wichman, R. 1998, *Nightfall*,
www.lsw.uni-hedelberg.de/~rwichman/Nightfall.html
- Wolszczan, A. & Frail, D. A. 1992, *Nature*, 355, 145

Young, A. T. 1967, *AJ*, 72, 747

Zucker, S. & Mazeh, T. 2002, *ApJ*, 568, L113