UNDERSTANDING AND REDUCING DISPERSED FIXED-DELAY INTERFEROMETRIC DATA FOR EXTRASOLAR PLANET SEARCHES

By

JULIAN CHRISTOPHER VAN EYKEN

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

UNIVERSITY OF FLORIDA

2007

© 2007 Julian Christopher van Eyken

To my parents (look what I did!); and to Gradstu, who helped me more than they know.

ACKNOWLEDGMENTS

I would like to express my gratitude to all those at Penn State, the University of Florida, and elsewhere who have helped me in putting this work together, and who have helped to keep me sane: to Jian and the whole ET instrument team for keeping ET afloat and giving me the opportunity to work on a great project; to my committee members, Robert Buchler, Steve Eikenberry, Elizabeth Lada, and Ata Sarajedini, for their helpful input; to Dimitri, Brian, and Suvrath, my partner in crime, for proofreading, and much helpful discussion and exchange of ideas; to the staff of the Kitt Peak and Apache Point observatories, for their patience and willingness to help; to Larry Ramsey, for a good dose of sanity and perspective; to Steinn Sigurdsson, for being so encouraging – and getting me into this whole mess in the first place; to Tamara and (again) Suvrath, the other two musketeers, for their patient acceptance while I offloaded all my culture shock and stress; to all those at Penn State who brought me over here and turned me into a proto-astronomer; and to the grad students at UF astro and my contemporaries at Penn State: there can be few finer crowds, and I wish there was room to name every every single one of you. (*All* y'all. And thank you, incidentally, for clearing all my forgotten biology experiments out of the fridge as I became increasingly out of touch with the outside world during the last few months. I did notice.) I am also tremendously grateful for all the financial support I have received, from the generous Penn State fellowship; from the SPIE scholarship program; and from the NASA JPL Michelson Fellowship program.

Finally, I would especially like to thank Nora and Fr. Guy, for helping keep my feet on the ground while I was trying to see beyond the clouds. I am genuinely not sure I could have done this without you.

To God: for adventures in distant lands that I would once never have dreamed of; for surrounding me with such wonderful friends; and for a Universe vaster, stranger and more fantastical than any human fiction.

TABLE OF CONTENTS

page

ACKNOWLEDGMENTS 4

LIST OF TABLES 8

LIST OF FIGURES 9

LIST OF SYMBOLS AND ABBREVIATIONS 12

ABSTRACT 13

CHAPTER

1 INTRODUCTION 15

1.1 Background 15
1.2 Exoplanet Research 16
1.3 The Search for Exoplanets 18
1.4 The Radial Velocity Technique 20
1.5 The ET Program and the DFDI Concept 23
1.5.1 The Need for a New Instrument 23
1.5.2 The DFDI Principle 24
1.5.3 A Brief History 25
1.5.4 The ET project 27

2 INSTRUMENT PRINCIPLES AND THEORY 28

2.1 Formation of a Fringing Spectrum 28
2.2 Fringe Phase and Visibility 31
2.3 The Interferometer Comb 35
2.4 From Phase to Velocity 38
2.5 Calculating the Interferometer Delay 40
2.6 Handling a Fiducial Reference Spectrum 43
2.6.1 The Addition Approximation 45
2.6.2 An Alternative: Combined-Beam Reference 48
2.7 Photon Error Propagation 49
2.7.1 Photon Error for Multiplied Reference 49
2.7.2 Photon Error for Reference Spectrum in Addition 52

3 INSTRUMENT HARDWARE 55

3.1 Subsystem Overview 55
3.2 The Single-Object ET at KPNO 58
3.2.1 Test Run, 2002 59
3.2.2 Upgrades, 2004 59
3.2.3 The current KPNO ET 60

3.3 The Multi-Object Keck ET at APO . . . 61

4 THE DATA REDUCTION PIPELINE . . . 65

4.1 Observing Procedure and Calibrations . . . 66
4.2 User Front-End . . . 67
4.3 Fibre Mapping . . . 70
4.4 Data Reduction . . . 71
4.4.1 Preprocessing . . . 72
4.4.1.1 Bias/dark subtraction and flatfielding . . . 72
4.4.1.2 Bad pixel masks and cosmic ray rejection . . . 72
4.4.1.3 Trimming . . . 73
4.4.1.4 Slant correction . . . 73
4.4.1.5 Illumination correction . . . 76
4.4.1.6 Filtering . . . 77
4.4.2 Fringe Fitting . . . 78
4.4.3 Radial Velocity Measurement . . . 80
4.4.3.1 Reference extraction . . . 80
4.4.3.2 Bulk shift compensation . . . 84
4.4.3.3 An improved approach . . . 85
4.4.4 Barycentric Correction . . . 86

5 SOURCES OF ERROR . . . 88

5.1 Error Bars . . . 88
5.2 Fringe Fitting . . . 91
5.3 Illumination Correction . . . 91
5.4 Slant Correction . . . 95
5.5 Flux Centroiding . . . 97
5.6 Barycentric Correction . . . 98
5.7 Contaminating Spectra . . . 98
5.8 Moonlight Contamination . . . 102
5.9 Residual Interferometer Comb . . . 103
5.10 The Addition Approximation . . . 104
5.11 Other Sources of Error . . . 108
5.12 Planet Detection Limits . . . 110

6 RESULTS FROM THE KPNO ET . . . 112

6.1 Confirmation of 51 Peg b . . . 112
6.1.1 Observations . . . 112
6.1.2 Data Analysis . . . 112
6.1.3 Results . . . 114
6.1.4 Discussion . . . 116
6.2 Discovery of ET-1 . . . 119
6.2.1 Early Measurements . . . 120
6.2.1.1 Observations and data analysis . . . 120

6.2.1.2 RV results . . . 122
6.2.1.3 Photometry . . . 126
6.2.2 Followup and Confirmation . . . 128
6.2.3 Discussion . . . 134
6.3 Addendum: 51 Peg b, ET-1, and the Addition Approximation . . . 134

7 RESULTS FROM THE MULTI-OBJECT KECK ET AT APO . . . 138

7.1 Prototype . . . 139
7.2 May 2006 Results . . . 140
7.3 November 2006 Results . . . 142
7.4 Discussion . . . 146

8 MEASURING STELLAR ROTATION WITH A DFDI . . . 150

8.1 Background . . . 150
8.2 Theoretical Predictions . . . 152
8.2.1 The Broadened Line Profile . . . 152
8.2.2 Fringe Visibility . . . 153
8.2.3 Predicted Relation . . . 156
8.3 Simulations . . . 159
8.3.1 Description . . . 159
8.3.2 Results and Dependence of $v \sin i$ on Stellar Parameters . . . 160
8.3.3 Comparison with Theoretical Predictions . . . 164
8.4 Observations . . . 165
8.5 Conclusions . . . 169

9 CONCLUSIONS AND CONTRIBUTIONS . . . 171

9.1 Work Remaining . . . 171
9.1.1 RV measurement . . . 171
9.1.2 Stellar rotation . . . 173
9.2 Future Possibilities . . . 173
9.3 Summary and Current Status . . . 174
9.4 Contributions . . . 176
9.5 Credits . . . 178

APPENDIX

A THE SPECTROGRAPH RESPONSE FUNCTION AND THE LSF . . . 180

B FRINGE FORMATION: AN ALTERNATIVE VIEWPOINT . . . 182

REFERENCES . . . 184

BIOGRAPHICAL SKETCH . . . 191

LIST OF TABLES

Table page

3-1 Keck ET instrument specifications 63

6-1 RV measurements for 51 Peg 116

6-2 RV measurements for η Cas 118

6-3 Mean photon limiting error estimation for 51 Peg and η Cas observations 118

6-4 Orbital parameters for HD 102195 130

6-5 Complete radial velocities for HD 102195 131

6-6 Stellar parameters for HD 102195 133

8-1 Dependencies of mean visibility, $\overline{\gamma}$, and $v \sin i$ on stellar parameters 163

LIST OF FIGURES

Figure page

2-1 Dispersed interferometer schematic 29

2-2 A representation of the summation (thick line) of fringes due to each individual wavelength of white light (thin lines) 34

2-3 Interferogram showing the coherence envelope due to a rectangular band pass . 36

2-4 Simulated interferometer comb 37

3-1 Optical design of ET at the KPNO 2.1m telescope 61

3-2 Optical design of the Keck ET interferometer 63

3-3 Optical design of the spectrograph for the Keck ET 64

4-1 Example screenshot from the ET pipeline graphical user interface 69

4-2 Example screenshot from the ET electronic observing log 70

4-3 Example non-fringing ThAr frame from the multi-object Keck ET 74

4-4 Slant correction of a non-fringing ThAr spectrum 75

4-5 Pre-processing steps 77

5-1 Phase and visibility errors due to curve-fitting alone 92

5-2 Fringe fitting errors for poor illumination correction 94

5-3 Effect of uniform artificial spectrum shift in the dispersion direction 96

5-4 Fringe along one channel due to target (upper curve), and contaminating low flux fringe (lower curve) 99

5-5 Vector representation of the summation of the fringes due to the target source and background contamination 100

5-6 Simulations of moonlight contamination 103

5-7 Vector representation of the summation of the true complex visibility and the error term due to the addition approximation 106

5-8 Analytically calculated expected error due to the addition approximation 108

5-9 Simulations showing the addition approximation error 109

5-10 Radial velocity semi-amplitude, K, for different minimum planet masses 111

6-1 Small section of raw fringing spectrum of 51 Peg 114

6-2 Plot of RV measurements for 51 Peg 115
6-3 Plot of RV measurements for η Cas 117
6-4 51 Peg RV measurements in December 2004 with the upgraded ET 120
6-5 36 UMa (known RV stable star) short-term precision measurements with the upgraded ET 121
6-6 Example section of raw spectrum taken with upgraded ET 122
6-7 Known planet-bearing star, 55 Cnc, measured with current KPNO ET 123
6-8 RV stable star, 36 UMa, measured over a few days with current KPNO ET . . . 124
6-9 Best fit Keplerian orbit for early ET-1 RV measurements 125
6-10 As for figure 6-9, but phase folded on the best-fit period 126
6-11 Lomb-Scargle periodogram for early ET-1 data 127
6-12 Early photometry of HD 102195 128
6-13 Folded combined radial velocities for HD 102195 129
6-14 Updated periodogram for HD 102195 132
6-15 Differential barycentric corrections for early 51 Peg measurements 136
6-16 Differential barycentric corrections for HD 102195 measurements 137
7-1 Results from the 20-object prototype 141
7-2 Example raw data frame from early Keck ET 143
7-3 Example solar data from early Keck ET 144
7-4 Known planet-bearing stars from early Keck ET data 145
7-5 Small selection of different example search star results from early Keck ET data 146
7-6 Keck ET measurements of HD 209458 from November 2006 147
7-7 Keck ET measurements of HIP 14810 (TYC 1231 1727 1) from November 2006 . 147
7-8 Representative selection of day-sky RV data from Keck ET, November 2006 . . 148
8-1 Top-hat response function centred on a single Gaussian absorption line 155
8-2 Gaussian best-fit approximation to a normalised rotational broadening profile . 157
8-3 Analytically predicted visibility vs. $v \sin i$ curve 159

8-4 Fringe visibility dependency on different stellar parameters from simulations . . 161

8-5 Contours of visibility vs. $v \sin i$ from simulations 162

8-6 Simulated data overplotted on the theoretical curve from figure 8-3 164

8-7 Real data and overplotted simulations for a well-matched subset of observed targets 167

8-8 All eight targets with overplotted visibility contours from simulations 168

LIST OF SYMBOLS AND ABBREVIATIONS

$\mathcal{F}[\ldots]_d$	—	Fourier transform evaluated at interferometer path difference d (or in some cases, time delay τ).
$\widehat{\ldots}\vert_d$	—	Shorter notation for Fourier transform, for convenience.
$\Re\{\ldots\}$	—	Represents the real part of a complex function.
$\widehat{\ldots}\vert_d$	—	Shorter notation for Fourier transform, for convenience.
$\ldots \otimes \ldots$	—	Used to denote convolution.
$[X^r, X^\phi]$	—	Notation used to represent in polar coordinates a general vector $\mathbf{X}$ of length, X^r, and phase, X^ϕ.
Γ	—	The scaling factor between change in fringe phase and change in radial velocity.
d_0	—	The fixed value of the interferometer delay.
M_J	—	Unit of the mass of Jupiter, conveniently approximately one thousandth of a solar mass.
$M_\odot$	—	Unit of a solar mass, 1.989×10^{33} g.
FWHM	—	Full-width half-maximum.
GUI	—	Graphical user interface.
LSF	—	Line-spread function.
PSF	—	Point-spread function.
PZT	—	Piezoelectric transducer.
PMT	—	Photomultiplier tube.
RMS	—	Root mean square.
RV	—	Radial velocity.
S/N	—	Signal to noise ratio (usually per pixel).

Abstract of Dissertation Presented to the Graduate School
of the University of Florida in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy

UNDERSTANDING AND REDUCING DISPERSED FIXED-DELAY
INTERFEROMETRIC DATA FOR EXTRASOLAR PLANET SEARCHES

By

Julian Christopher van Eyken

December 2007

Chair: Dr. Jian Ge
Major: Astronomy

The dispersed fixed-delay interferometer represents a new instrument concept for high-precision radial velocity surveys for extrasolar planets. A combination of an interferometer in series with a medium resolution spectrograph, it has the potential for performing multi-object surveying down to faint magnitude limits, where previous radial velocity techniques have been limited to observing only one target at at time. The sample of relatively bright stars that can quickly be surveyed using traditional radial velocity techniques with current technology is becoming exhausted, and the radial-velocity planet discovery rate is beginning to level out. Because of the large sample of extrasolar planets needed to perform statistical analyses to be able to understand aspects of planetary formation and evolution, such a multi-object instrument represents the next logical step in instrumentation for extrasolar planet research. As a useful by-product, the instrument has the potential to provide simple measurements of the projected rotational velocities of stars.

The development of this instrument has necessitated the development of new data reduction procedures to efficiently and reliably turn the raw data into meaningful results, and a careful consideration of precision levels and sources of error. This in turn has required fleshing out a detailed understanding of the physical principles of the instrument.

The single-object Exoplanet Tracker installed at Kitt Peak National Observatory and the multi-object W. M. Keck Exoplanet-Tracker at Apache Point Observatory are the first

two fully-fledged astronomical radial velocity instruments to have been built based on this new technique. The former served to prove the concept with a successful confirmation of the known planet, 51 Peg b. We were later able to use it to discover a new planet, HD 102195 b, or ET-1. The latter has also demonstrated detection of known planets, and has been the workhorse for a pilot survey in preparation for a planned major wide-field survey over the next decade.

CHAPTER 1
INTRODUCTION

1.1 Background

It has become almost a tradition that articles on the subject of extrasolar planet searches begin with a tally of the number of planets found to date. At the time of writing, a total of 268 extrasolar planets (or 'exoplanets') have been discovered in 230 planetary systems, of which 26 are multiple planet systems, according to the regularly updated online Extrasolar Planets Encyclopaedia.[1] That the exact count still remains of such interest is a testament to the difficulty of detecting these objects, and to the attractiveness of the challenge.

Until less than two decades ago, the Solar System contained the only known planets known to humankind, and their uniqueness in the Universe remained an open question. The last two decades have seen a dramatic shift in that perspective. The first exoplanet detection was achieved in 1992, and was actually a detection of a multi-planet system, orbiting a pulsar, PSR 1257+12, by Wolszczan & Frail (1992). Using exquisitely precise pulse timing measurements, they were able to detect the tiny shift in the distance of the pulsar from the Earth due to the pulsar's reflex motion as its planets orbited. The first unambiguous 'traditional' planetary companion to a main-sequence solar-type star was detected in 1995, orbiting 51 Pegasi. Using the radial velocity (RV) technique, Mayor & Queloz (1995) detected the tiny Doppler shift in the spectrum of the star due to its reflex 'wobble' caused by its orbiting companion. More planet detections quickly followed (e.g., Marcy & Butler 1996; Butler & Marcy 1996; Butler et al. 1997), and the field of exoplanet research was born. A recent compilation of the known exoplanets was given in Butler et al. (2006).

[1] Maintained by Jean Schneider, `http://exoplanet.eu/`

The fascination with exoplanet research is perhaps in part a reflection of the perennial human appetite for exploration; and perhaps also in part due to the fact that it provides us with some human connection between our own world and the Universe at large. Learning about exoplanets can start to tell us about their ubiquity and formation (e.g. Ida & Lin 2004a,b, 2005; Armitage 2007), which in turn can begin to tell us about the formation of our own Solar System, our own Earth, and ultimately our origins and place within the Universe. The discovery of exoplanets has provided an important step in the quest to find life – if it exists – elsewhere in the Universe (e.g. Seager et al. 2005; Segura et al. 2005; Raymond et al. 2006), and has been an important catalyst for the birth of the field of astrobiology (Morrison 2001). The now apparent ubiquity of planets fulfils a major requirement for the existence of life, as far as we understand it, beyond our own world.

1.2 Exoplanet Research

From the very first discovery, the diversity of exoplanets has been a surprise, belying expectations from the start. The properties of many of the planetary systems found so far are very different to our own, and a surprising diversity is seen (see Perryman 2000; Marcy et al. 2005; Udry et al. 2007, for reviews of the field). Around 1% of the stars surveyed, for example, harbour 'hot Jupiters', very close-in gas-giant planets with minimum masses on the order of a Jupiter mass or more, and orbital periods of only a few days ($< 10\,\mathrm{d}$) (Udry et al. 2007). With the exception of the extremely close-in planets where tidal interaction with the host star is expected to have circularised the orbits, the majority of planets are also on very eccentric orbits, with a median eccentricity of 0.26 (Udry et al. 2007), unlike our own system where they are all on highly circularised orbits. A very wide range of minimum planet masses has also been found, from the lower limit of detection – around a Neptune-mass and recently even several–Earth-masses – up to the rare instances where the minimum mass is $\sim 10\,M_{\mathrm{J}}$, right on the boundary where the distinction between

planet and brown dwarf[2] becomes blurred. For those planets found in multiple systems, many also show planet-planet interactions, some showing resonances in their orbital period ratios (Kley et al. 2004, and references therein), others interacting more strongly in chaotic systems (e.g. Correia et al. 2005).

Finally, as survey time baselines lengthen and sensitivity improves, systems similar to our own Solar System ('Solar System analogues') are just becoming accessible (Marcy et al. 2002). In order to make a solid detection of an exoplanet, it is generally necessary to observe the candidate target for the duration of at least one orbital period. Largely for this reason, Solar System analogues are only now becoming accessible to current detection techniques: Jupiter would present the most easily detectable of our planets were we to view it from another stellar system, but the baseline of observations since the start of major planet surveys is only now becoming long enough to match its 12-year orbital period.

This diversity has raised many questions in exoplanet research. Some of the current major questions include:

- How ubiquitous are exoplanets? From current surveys, $\sim$ 6% of solar-type (i.e. generally late F- to early M-type main sequence) stars have been found to harbour planets down to a completeness limit of $0.5\,M_{\rm J}$, and around 1% harbour hot Jupiters. Of the known planet bearing stars, $\sim$ 12% have known multiple planet systems, and this may in fact be the norm. As surveys increase in sensitivity, it is beginning to appear that even the majority of solar-type stars may harbour planets (Udry et al. 2007).

- How do properties of the host star affect those of companion planets? Planet presence is found to be strongly correlated with host star metallicity (Fischer & Valenti 2005); it is also starting to appear that more massive stars may possibly lead to more massive planets (Sato et al. 2007; Ida & Lin 2005).

[2] Essentially failed stars – defined by an IAU working group as objects massive enough to ignite deuterium burning, but not massive enough to burn hydrogen, with the boundary between planets and brown dwarfs set at $13\,M_{\rm J}$.

- What is the mass distribution of exoplanets? From the lower limit of masses currently detectable ($< 1\,M_{\rm J}$), the distribution tails off as roughly the inverse of the mass (Marcy et al. 2005) up to the brown dwarf mass boundary at $13\,M_{\rm J}$. There is a pronounced absence of low mass companions in the mass range $\sim$ 20–60 $M_{\rm J}$, known as the 'brown dwarf desert', beyond which point hydrogen burning stars begin to take over. This is taken to suggest different formation mechanisms for stars and planets.

- What is the orbital period distribution? There is found to be a pileup in the period distribution of planets at periods of about 3 d. Very few planets are found in orbits shorter than this period. Going out to longer periods, there is a slow increase in occurrence again with increasing period. In single-planet systems with periods less than 100 d, only planets of mass less than $\sim 2\,M_{\rm J}$ have been found.

- What is the effect of stellar binarity on planet formation? Owing to the added difficulty involved in finding planets around multiple star systems, they tend to have been excluded from surveys, and so this remains somewhat of an open question (Udry et al. 2007)

- What gives rise to the high eccentricities of exoplanet orbits? If planets form in protostellar debris discs, one might naively expect simple circular orbits to result, and yet most planets are on surprisingly eccentric orbits. (Tremaine & Zakamska 2004).

- What do these questions, along with other correlations between planet parameters, tell us about the formation and evolution of planets? Each piece of statistical information gathered from planet surveys ultimately provides clues to guide and constrain the theoretical models that tell us about planetary formation (e.g. Ida & Lin 2004a,b, 2005; Armitage 2007).

1.3 The Search for Exoplanets

Actual direct detection of exoplanets by imaging has been somewhat of a holy grail, but until now has remained elusive because of the extreme contrast between the planet and host star brightnesses, and their close proximity on the sky. Technology is beginning to catch up to the challenge, with several detections recently reported (Beuzit et al. 2007, and references therein). Most techniques, however, depend on indirect observations of the host star to infer the presence of a companion planet.

One of the earliest claims to the detection of a planet was made by van de Kamp (1963) using astrometry, orbiting Barnard's star, although this claim was later discredited by Gatewood & Eichhorn (1973). A planet-bearing star executes a very small reflex

motion as its planet orbits, since both planet and star are in fact orbiting their common centre of mass. The astrometric technique depends on measuring this wobble as a change in position of the star on the plane of the sky, and is sensitive to long orbital periods, more massive planets, and lower mass stars (of which Barnard's Star was an example, assumed by van de Kamp to have a mass of $0.15\,M_{\odot}$). The positional accuracy required in the measurements are typically extremely small however, and difficult to achieve from the ground owing to the Earth's atmosphere (Perryman 2000). To date, the technique has not seen any success in discovering planets (see Sozzetti (2005) for a full review of the field).

Even more indirect, but to this point more successful, is the gravitational lensing technique (e.g. Gould et al. 2006), where the proper motion of a faint planet-bearing star causes it to pass in front of another much brighter star, gravitationally lensing the background star's light towards the Earth and causing a brief brightening in magnitude. Small excursions in the expected shape of the lightcurve of the background star during the event betray the existence of a planetary companion. The events are only one-time detections which cannot be revisited, however, and the host stars are typically extremely faint, making them very hard to follow up with other techniques and limiting the information that can be gained about the planet.

As early as the 1950s, Struve (1952) pointed out that a fraction of exoplanets, those which transit their parent stars, should be detectable through photometry of the host star during the transit. The transit method has the advantage that it is relatively easy to target many stars at once in a wide field, and so what it loses in terms of numbers of planets which transit, it gains in survey numbers; it also is uniquely able to determine the radius of the planet, and its orbital inclination. The transit method has been slow to take off but has recently been rapidly gaining ground (Charbonneau et al. 2007).

The original pulsar timing method employed by Wolszczan has since been relatively unproductive. However, the same technique has been applied to similarly exotic pulsating white dwarfs, and is showing some promise (Mullally & Winget 2006; Mullally et al.

2006). Even more recently, a first planet has been claimed using timing of pulsations of a post–red-giant star, a dramatic new extension to the range of star types which have been found to host planets (Silvotti et al. 2007).

Of all the techniques for finding exoplanets, however, by far the highest yield has come from the radial velocity technique, and it is this that the ET instruments depend on.

1.4 The Radial Velocity Technique

The radial velocity technique is somewhat complementary to the transit technique. In itself, it can determine the orbital period, semi-major axis, eccentricity, and the minimum mass, $M \sin i$, where M is the planet mass and i the orbital inclination to the line of sight. Transit measurements can determine i, however, so in those instances where the planet transits the host star, it becomes possible to pin down the actual mass of the planet by combining with the RV data. Since the RV technique is biased towards finding planets with orbital planes that lie in the line of sight, it is also suited to detecting planets that may be suitable for transit followup.

Like the astrometric technique, the radial velocity technique relies on the reflex motion of the host star as its planet orbits. However, in this case, it is the radial velocity along the line of sight to the star that is measured.

When light from a star is dispersed in a spectrograph, a spectrum is generally seen that broadly resembles a black-body function, cut with many narrow absorption lines. These are caused by absorption in the photosphere of the star by the various atomic elements present, at the wavelengths of their respective atomic transitions. The depths of the lines are highly varied, depending on the opacity of their respective transitions and the depth in the photosphere – and hence the black-body temperature – at which they are formed. For solar-type stars, the equivalent widths of the lines in the optical band are typically on the order of 0.1 Å or more (depending on the rotational velocity of the star), with a handful of particularly strong metal lines on the order of a few to ten times wider. The line width is determined primarily by Doppler broadening, caused by the thermal

motion of the atoms in the gas, and by rotational broadening, where the line profile starts to represent an integrated sampling of the different projected velocities over the surface of the stellar disc (Gray 1992).

When the radial velocity of the star changes, a bulk Doppler shift in the wavelength of the absorption lines results, and it is this that enables the measurement of changes in radial velocity (specifically, the information on the Doppler shift is in the slopes of the line profiles). Along with the transit method, the radial velocity technique was proposed by Struve (1952) as the primary method for finding exoplanets.

Exoplanet radial velocity surveys have traditionally depended on recording very high resolution spectra to obtain well resolved absorption lines, and either cross correlating the spectra with reference template spectra, or fitting functions to the line profiles themselves to measure the positions of the centroids. Best internal precisions have typically reached down to the $3\,\mathrm{m\,s^{-1}}$ level (Butler et al. 1996; Vogt et al. 2000), and are now reaching as low as $1\,\mathrm{m\,s^{-1}}$ or better (Pepe et al. 2005). (For comparison, a Jupiter analogue in a circular orbit around a solar-type star would cause sinusoidal radial velocity variations with an amplitude of about $12.5\,\mathrm{m\,s^{-1}}$.)

The very high levels of precision required for planet detection and the difficulty of directly measuring absolute wavelengths means that some kind of stationary reference spectrum is invariably used as a calibration. Early RV measurements were limited in their precision largely because of differences between the source and reference illumination profiles entering the instrument (Griffin & Griffin 1973). Furthermore, the reference was often either passed along a separate beam path through the instrument, or taken at a different time: in both cases, differential instrument drifts (e.g. due to thermal flexure or changes in air pressure) would affect source and calibration differently. Griffin & Griffin (1973) made the first attempts to overcome these problems by using a simultaneous superposed reference: in this case, they proposed using the telluric lines imposed by the Earth's atmosphere. In this way, the reference spectrum passed through the instrument

in a way absolutely identical to the source spectrum, thus eliminating any question of differential effects.

The atmosphere, however, being beyond the observer's control, presents its own difficulties. The telluric spectrum can vary with changing atmospheric wind speeds, changing levels of water vapour, and also with zenith angle. Campbell & Walker (1979) overcame these problems by inserting a glass cell filled with hydrogen fluoride (HF) gas into the optical path of the instrument as a fiducial reference, superimposing a well controlled set of lines in the range 8670–8770 Å and achieving a precision of $15\,\mathrm{m\,s^{-1}}$. Of the 21 stars they had observed over 12 years, none were found to harbour any planets at their detection limits (Walker et al. 1995), though this result is consistent with the now-known fraction of planet-bearing stars (Udry et al. 2007).

Hydrogen fluoride also presented a number of problems as a reference, however, not the least of which was that the gas is corrosive and highly toxic. The now-favoured iodine reference was first adopted by Marcy & Butler (1992). Substituting a glass cell filled with heated iodine vapour rather than HF had the advantage of allowing for a much smaller cell length, and created a reference spectrum that covered the much broader wavelength range of 5000–6300 Å. Since broader wavelength coverage amounts to more flux, one could search much fainter targets, where previous magnitude limits had significantly limited the number of stars available for search.

The years surrounding these developments also saw the birth of the optical fibre in astronomy. Optical fibres allowed the instrument to be completely separated from the telescope, so that it could be situated in a better controlled and more stable environment. Furthermore, the optical scrambling properties of fibres could be used to erase the changing spatial structure of the beam profile before entering the instrument, so that a consistent illumination was always presented (Heacox 1986). It thus became possible for the first time to consider using a parallel reference beam running alongside the main science beam for planet hunting. This was the approach adopted for the ELODIE

instrument (Baranne et al. 1996), using a Thorium Argon (ThAr) emission spectrum as a reference, and it was ELODIE that was used to first discover 51 Peg b (Mayor & Queloz 1995).

In time Marcy's group was able to push their RV precision down to the $3\,\mathrm{m\,s^{-1}}$ level (Butler et al. 1996) using the iodine technique. There are now a number of other active planet surveys using the RV method, including the Lick, Keck and AAT program (Marcy et al. 2005); the CORALIE (Udry et al. 2000, and references therein), ELODIE (Baranne et al. 1996; da Silva et al. 2006), and HARPS (Pepe et al. 2000, 2005) programs; and the McDonald Observatory (Cochran & Hatzes 1993) and AFOE (Brown et al. 1994) programs, among others. Precisions down to better than $1\,\mathrm{m\,s^{-1}}$ are now being reached by the HARPS group, using parallel ThAr (with the option of using iodine as an alternative). A review of radial velocity discoveries is given in Udry et al. (2007).

1.5 The ET Program and the DFDI Concept

1.5.1 The Need for a New Instrument

Despite the achievements in planet detection, more planets are still needed to constrain formation and evolutionary models. This is in part because of the diversity of planet properties that has been found; but it is also because, to date, the majority of surveys have not used well-defined unbiased target lists, making it difficult to perform robust statistical analyses. To this point, the primary concern has been finding planets in the first place, and most surveys suffer from completeness issues or biases toward planet detection (e.g. da Silva et al. 2006). Armitage (2007) concludes that there is still a strong need for large uniform surveys to enlarge the statistical sample available: of all the planets known, he was only able to find a uniform subsample of 22 that satisfied the requirements needed for a statistical comparison with models.

A few thousand stars have been searched between the various RV surveys, including most stars down to visual magnitude ~ 8, but instrument light throughput issues have made it less easy to go much fainter than this, and RV detection rates have been levelling

off in the last few years. Although the rate of detections from transit surveys is now quickly taking off, transit surveys can only detect the small fraction of planets which happen to eclipse their parent stars, ($\sim$ 10% probability for hot Jupiters, from geometrical considerations – Kane et al. 2004). Furthermore, the complementary information gained from RV detections remains of great value. There is therefore a strong case for finding a technique which enables RV surveying down to fainter magnitudes and at considerably faster speeds than have been achieved over the last decade. The goal of the ET program described in this work is to satisfy this requirement.

1.5.2 The DFDI Principle

The Exoplanet Tracker (ET) project is built on a new type of fibre-fed RV instrument based on the 'dispersed fixed-delay interferometer' (DFDI), a combination of a Michelson interferometer followed by a low or medium resolution post-disperser (also referred to by Erskine et al. (2006) as an externally dispersed interferometer, or 'EDI'). The effective resolution of the instrument is determined primarily by the interferometer, so the post-dispersing spectrograph can be of much lower resolution than in traditional techniques, and consequently can have higher throughput (Ge 2002; Ge et al. 2003b,a). The technique is closely related to Fourier transform spectroscopy: the post-disperser effectively creates a continuum of very narrow bandpasses for the interferometer, increasing the interference fringe contrast. All the information needed is contained in the fringe phase and visibility (see chapter 2). It turns out that since we are only interested in the Doppler shift of the lines, measurements are required at only one value of interferometer delay (hence 'fixed delay').

The cost of the instrument is comparatively low, and most importantly, it operates in a single-order mode. Where traditional echelle spectrograph techniques operate by spreading a single stellar spectrum over an entire CCD detector in multiple orders, here the spectrum only takes up one strip along the detector. Spectra from multiple stars can therefore be lined up at once on a single detector (Ge 2002). In combination with a wide

field multi-fibre telescope, this makes multi-object RV planet surveying possible for the first time (Mahadevan et al. 2003).

1.5.3 A Brief History

The idea of using the combination of a Michelson interferometer with a postdisperser was first proposed for precision Doppler planet searches by D. J. Erskine in 1997, at Lawrence Livermore National Laboratory (Erskine & Ge 2000; Ge 2002; Ge et al. 2002; Erskine 2003). The same approach is being followed by Erskine et al. (2006) in the infra-red, in an attempt to find planets around late-type stars. A similar approach is discussed by Mosser et al. (2003) for asteroseismology and the measurement of stellar oscillations; more recently the technique has also been adopted for the USNO dFTS instrument (Hajian et al. 2007) (in this last case, the interferometer delay can also be varied so that high resolution spectra can be reconstructed).

The idea of dispersed interferometry is by no means new. Michelson himself recognised the use of interferometers for spectroscopy (Michelson 1903), and even proposed combining a disperser in series with a Michelson interferometer. In this case the disperser, a prism, was placed *before* the interferometer, allowing only a narrow bandwidth of light to enter the interferometer in the first place, but the underlying physics is essentially equivalent. In what was probably the first realisation of a DFDI, Edser & Butler (1898) placed a Fabry-Perot type interferometer in front of a spectrograph to produce dispersed fringes (effectively an interferometer comb - see chapter 2), which they used as a fiducial reference for measuring the wavelengths of spectral lines. Such dispersed fringes were later to become known as 'Edser-Butler fringes' (Lawson 2000).

The various combinations of interferometers with dispersers came into use somewhat later in the field of astronomy. Geake et al. (1959) placed a Fabry-Perot interferometer before a spectrograph to allow the spectrograph slit to be widened and hence increase the throughput. P. Connes developed the SISAM technique for spectroscopy, replacing the mirrors in a Michelson interferometer with gratings to select a small bandpass and so

increase fringe visibility for precise spectroscopy (detailed in a review of interference spectroscopy by Jacquinot 1960). The later SHS (Harlander et al. 1992) and HHS (Frandsen et al. 1993; Douglas 1997) techniques shared the use of internally dispersed Michelson interferometers, but used the dispersion instead for the purpose of scanning a range of interferometer delays, obviating the need for moving parts.

Barker & Hollenbach (1972) were able to measure the velocity history of a projectile in the laboratory in reflected laser light, in an early example of the use of true fixed-delay interferometry for velocimetry. The use of a Michelson interferometer for astronomical RV measurements was proposed shortly afterwards by Gorskii & Lebedev (1977) and Beckers & Brown (1978). Forrest & Ring (1978) also proposed using a Michelson interferometer with a fixed delay for high-precision Doppler measurements of single spectral lines for the detection of stellar oscillations.

During the 80s, Connes (1985) proposed a novel technique using laser tracking of a Fabry-Perot interferometer, which in turn tracks the Doppler shifts of stellar spectral lines for high precision measurements. More recently, McMillan et al. (1993) have used a Fabry-Perot interferometer combined with a cross-dispersed echelle spectrograph for precision Doppler measurements (see also McMillan et al. 1994). Fixed-delay Michelson interferometers with very narrow band passes have also been used for producing Doppler *images* by Shepherd et al. (1985) (the WAMDII instrument) for measuring upper atmospheric winds, and Harvey & The GONG Instrument Team (1995) (the GONG project) for Doppler measurements across the solar disc. Here, fringe phase measurements were made over the field of an image to give Doppler measurements at each point in the field of view.

Many of these interferometric instruments suffered from the limitation of having an extremely narrow bandpass, tending to limit their application to only bright targets. The DFDI technique used in the ET instruments allows for an arbitrarily wide bandpass, limited only by the spectrograph capabilities, while still retaining the high resolution

spectral information needed for precision velocity measurements. The first such DFDI measurements were made at the Lawrence Livermore National Laboratory and the Lick 1m telescope between 1997 and 1999, and were reported in Erskine & Ge (2000) and Ge et al. (2002). The ET project began shortly after.

1.5.4 The ET project

The ET project was developed at Penn State University and the University of Florida, beginning in 2000. Early lab tests were performed at Penn State, and prototype test runs were conducted at the McDonald Observatory Hobby-Eberly Telescope in late 2001, and at the Palomar 200 inch telescope in early 2002 (the results from those runs are not discussed here, but are reported in Ge et al. 2003b; Mahadevan 2006).

Two ET instruments have now been built: the single-object prototype ET, permanently installed at the KPNO 2.1m telescope in 2003 after a temporary test run in August 2002, and the multi-object Keck ET, installed at the APO Sloan 2.5m telescope in March 2005, and then upgraded and moved to a more stable location later that year (see chapter 3).

Proof of concept was achieved using the KPNO ET with the first DFDI planet detection, a confirmation of the companion to 51 Pegasi (van Eyken et al. 2004a). Our first planet discovery, HD 102195b (ET-1) was also later made using this instrument (Ge et al. 2006a). The multi-object Keck ET is a full scale instrument being developed to satisfy the survey requirements laid out in section 1.5.1, and it is hoped that it will have a dramatic impact on the planet discovery rate over the coming decade.

CHAPTER 2
INSTRUMENT PRINCIPLES AND THEORY

Although various forms of the dispersed fixed-delay interferometer (DFDI) have been used before, the concept, particularly in its specific application to exoplanet finding, is rather new. Much of the work in understanding the data from the instrument has therefore involved coming to a full understanding of the physics of the instrument itself. Some related theory is discussed in a number of sources, for example Goodman (1985); Erskine & Ge (2000); Lawson (2000); Erskine (2003); Ge (2002); Ge et al. (2002); Mosser et al. (2003); van Eyken et al. (2003). An attempt is made here to draw together, expand on, and more precisely state the theoretical material needed for data reduction, and to provide a somewhat complete overview of the physics underlying the instrument's working. This provides us with the fundamental mathematics necessary for writing a data reduction pipeline.

2.1 Formation of a Fringing Spectrum

Figure 2-1 shows a highly simplified schematic of a dispersed fixed-delay interferometer (DFDI), consisting of the two main components, a fibre-fed Michelson interferometer and a disperser, followed by a detector. Light input from the fibre is split into two paths along the arms of the interferometer and then recombined at the beamsplitter. The output is fed to the disperser, represented for convenience as a prism, though generally this will be a spectrograph. An etalon is placed in one of the interferometer arms to create a fixed optical path difference (or 'delay'), $d = d_0$, between the two arms, while allowing for adequate field widening (Hilliard & Shepherd 1966). d_0 is typically on the order of mm. In practice, an iodine vapour cell can also be placed in the optical path before the interferometer to act as a fiducial reference (section2.6).

Inputting a wide beam of monochromatic light into the instrument with both interferometer mirrors exactly perpendicular to the light travel path will give either a bright or a dark fringe at the output of the interferometer, as shown in figure 2-1A,

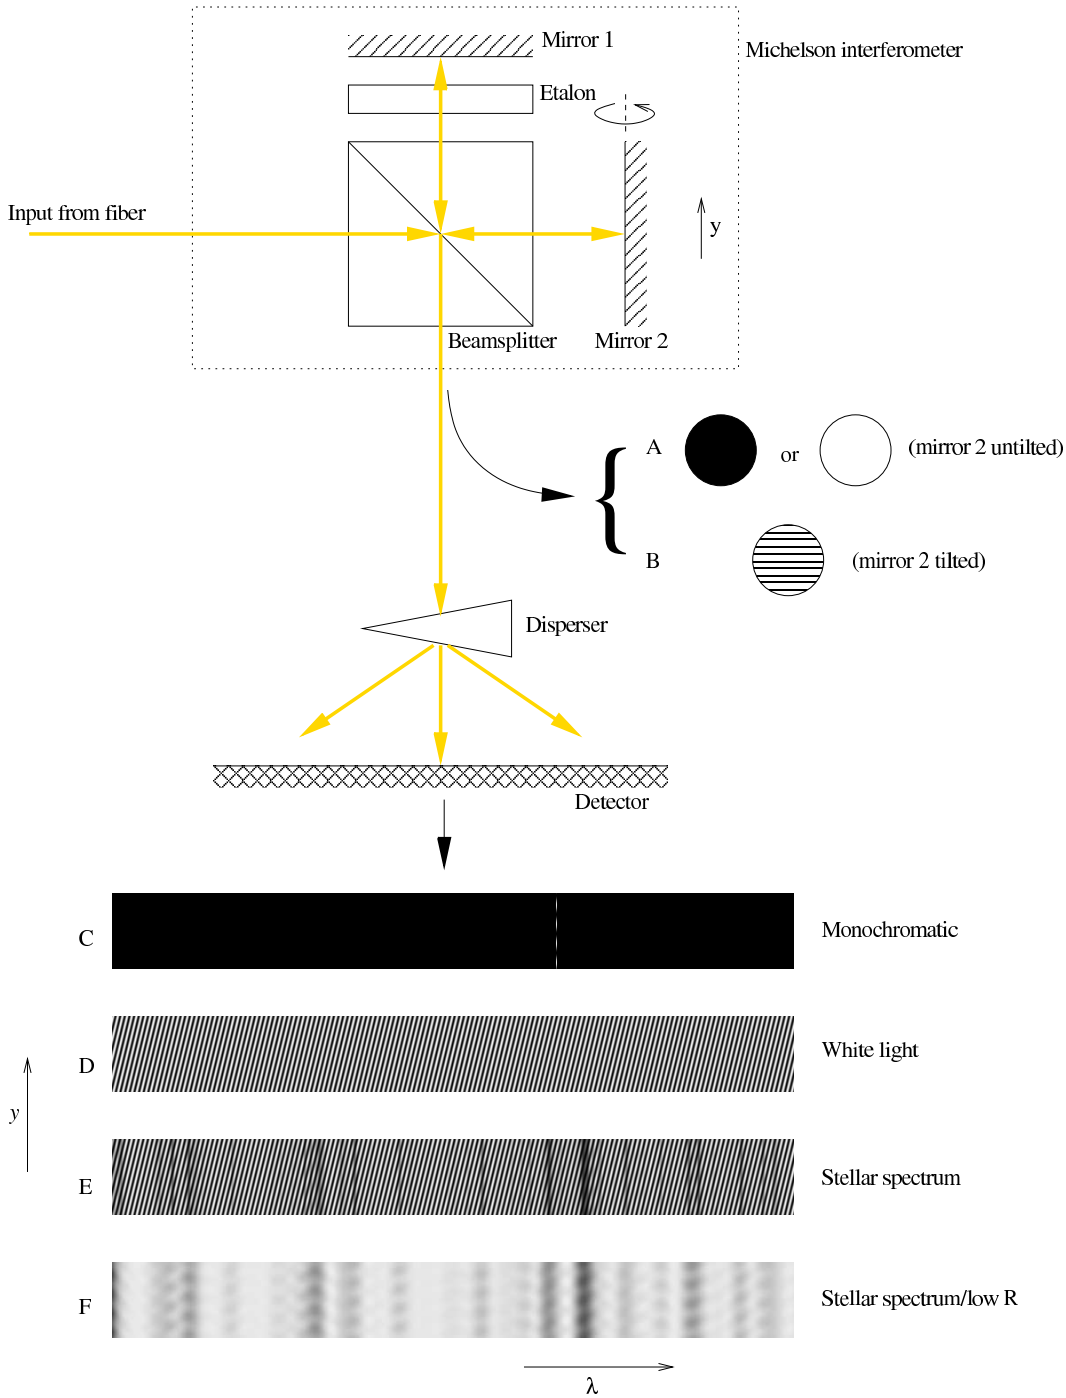

Figure 2-1. Dispersed interferometer schematic. y corresponds to position in the slit direction, and λ indicates wavelength in the dispersion direction. A) Output from interferometer alone with monochromatic light input, and mirror 2 untilted. B) The same with mirror 2 tilted along the axis in the plane of the page, as shown. C) Image on detector with monochromatic light. D) Detector image with white light input. E) Image with stellar spectrum input. F) As for E but at low resolution.

depending on whether the exact path difference d between the two arms corresponds to constructive or destructive interference. If we were to scan the mirror in the arm without the etalon (mirror 2) back and forth, the intensity at the interferometer output would vary sinusoidally as a function of d. If we now tilt this mirror along the axis in the plane of the page, we effectively scan a small range of delays along the y direction (i.e. perpendicular to the axis of the tilt and in the plane of the mirror, corresponding to the slit direction in the spectrograph). Hence we would see a series of parallel bright and dark fringes, now varying sinusoidally as a function of y.

Consider first a very high (actually infinite) resolution spectrograph disperser for the sake of argument: following the beam through to the detector plane would result in a single emission line with fringes along the slit direction, as in figure 2-1C. Switching the input spectrum to white light, which can be thought of as a continuum of neighbouring delta functions, leads to a similar fringe pattern on the detector at every wavelength channel. Due to the fact that, in terms of number of wavelengths, the optical path difference is different for different wavelengths, each fringe is slightly offset in phase from its neighbours (and very slightly different in period). This gives rise to the series of parallel lines known as the interferometer 'comb', shown in figure 2-1D. Going further and inputting a stellar spectrum into the instrument would simply give the product of the stellar spectrum and the comb, as in figure 2-1E. Finally, changing to the real case of a low or medium resolution spectrograph as for an ET-type instrument, the comb is no longer (or barely) resolved, and we see a spectrum like that in figure 2-1F. Such a spectrum is sometimes referred to as a spectrum "channelled with fringes", also known as Edser-Butler fringes (Edser & Butler 1898; Lawson 2000; Ge 2002). The remaining fringes contain high spatial frequency Doppler information that has been heterodyned down to lower spatial frequencies by the interferometer comb (Mahadevan 2006; Erskine 2003). It is this heterodyning that allows for the use of a low-resolution spectrograph at low dispersion, and is the key to the DFDI technique.

2.2 Fringe Phase and Visibility

Each wavelength channel j on the detector has an associated sinusoidal fringe running along the slit direction. A given fringe has an associated phase and visibility, where visibility is a measure of the contrast in the fringe, defined as the ratio of the amplitude of the fringe to its central (mean) flux value, or equivalently, $(I_{\max} - I_{\min})/(I_{\max} + I_{\min})$, where $I_{\max}$ and $I_{\min}$ are the maximum and minimum flux values in the fringe (Michelson 1903). Here we introduce the concept of a 'whirl' (Erskine & Ge 2000). The phase and velocity for a fringe can together be thought of as representing a vector, with the visibility representing the magnitude. An array of these representing a full spectrum of wavelength channels is called a whirl. By 'wavelength channel', we mean specifically an infinitesimally wide strip along the slit direction at pixel position j, where j is not necessarily an integer. The whirl is what we measure directly from a fringing spectrum and contains the information relevant to velocity determination. Vector operations such as addition, subtraction and scalar products can be performed on these whirls just as for the individual vectors (Erskine & Ge 2000).

To understand what determines the values of the phase and visibility for a fringe, we can consider the contribution from each wavelength of light to a particular wavelength channel on the detector. Each contributing wavelength has passed through the interferometer, and for an ideal interferometer, will contribute a sinusoid of 100% visibility like that in figure 2-1C. The intensity of these sinusoids on the detector can each be described by $\Re\{1 + \exp(\mathrm{i}2\pi d/\lambda)\}$, where d varies linearly with y along the length of the slit, and $\Re$ represents the real part of the argument. Since the spectrograph has finite resolution, the line spread function (LSF) of the spectrograph will mean that a narrow band of such wavelengths will contribute to any given resolution element. The measured fringe is a continuous summation of those sinusoids, weighted by the shape of the spectrum contributing to that channel, Q. Switching from wavelength to wavenumber $\kappa \equiv 1/\lambda$, this

summation can be expressed as:

$$I(d) \quad = \quad \Re\left\{\int Q(\kappa)(1+\mathrm{e}^{\mathrm{i}2\pi\kappa d})\ \mathrm{d}\kappa\right\} \quad = \quad \int Q(\kappa)\ \mathrm{d}\kappa + \Re\left\{\int Q(\kappa)\mathrm{e}^{\mathrm{i}2\pi\kappa d}\ \mathrm{d}\kappa\right\}, \quad (2\text{–}1)$$

where $I(d)$ is the measured intensity along the slit direction. The first term on the right hand side is simply the total integrated flux in the channel, which must be real valued. The second term can immediately be identified as the real part of a Fourier transform, $\Re\{\mathcal{F}[Q]_d\}$, with delay as the conjugate variable to wavenumber, and shows the close relationship between DFDI instruments and Fourier transform spectroscopy (Jacquinot 1960).

Normalising by dividing through by the total flux, we can define the complex quantity $\boldsymbol{\gamma}$ such that

$$I_{\mathrm{norm}}(d) = 1 + \Re\left\{\frac{\mathcal{F}[Q(\kappa)]_d}{\int Q(\kappa)\ \mathrm{d}\kappa}\right\} = 1 + \Re\{\boldsymbol{\gamma}\}, \quad (2\text{–}2)$$

where

$$\boxed{\boldsymbol{\gamma} \equiv \gamma\mathrm{e}^{\mathrm{i}\phi_\gamma} \equiv \frac{\mathcal{F}[Q(\kappa)]_d}{\int Q(\kappa)\ \mathrm{d}\kappa}} \quad (2\text{–}3)$$

is the 'complex degree of coherence' (Goodman 1985), and describes the phase, ϕ_γ, and amplitude, γ of the normalised fringes (i.e. the visibility), as a function of d. $\boldsymbol{\gamma}$ is referred to here as the *complex visibility.* More rigorous derivations of this can be found in Goodman (1985, ch. 5) and Lawson (2000), but this simplification is adequate for the explanation here.

In order to understand the actual form of the fringes seen in DFDI, it is important to realise that the nature of the instrument is such that for any given wavelength channel, the contributing spectrum Q has a very narrow passband (for our instruments, $\Delta\lambda/\lambda \sim 1\,\text{Å}/5000\,\text{Å} = 2\times10^{-4}$). We imagine Q as being equal to a function Q_0 of characteristic width $\Delta\kappa$ and centred at zero wavenumber, which has been shifted in wavenumber so that its centre falls at wavenumber $\kappa = \overline{\kappa}$, and $Q(\kappa) = Q_0(\kappa - \overline{\kappa})$. By the Fourier shift theorem we can write:

$$\mathcal{F}[Q]_d = \mathrm{e}^{-\mathrm{i}2\pi d\overline{\kappa}}\mathcal{F}[Q_0]_d. \quad (2\text{–}4)$$

The right hand side shows two components. The exponential term represents a linear phase variation with delay, varying on scales of the period $1/\overline{\kappa}$. By the reciprocal scaling property of Fourier transforms, the second term, the Fourier transform, can be expected to vary on scales of the reciprocal of the width of Q_0, that is, on scales of $1/\Delta\kappa$. Since $1/\Delta\kappa \gg 1/\overline{\kappa}$, the real part of the right hand side represents a sinusoidal fringe of frequency $\overline{\kappa}$ modulated by a slow variation in both phase and amplitude, so that substituting into equation 2–2 we can write:

$$I_{\mathrm{norm}}(d) = 1 + \frac{\Re\left\{\mathrm{e}^{-\mathrm{i}2\pi d\overline{\kappa}}\mathcal{F}[Q_0]_d\right\}}{\int Q(\kappa)\ \mathrm{d}\kappa} \tag{2–5}$$

If we define:

$$\mathbf{V}(d) \equiv V(d)\mathrm{e}^{\mathrm{i}\phi_V(d)} \equiv \frac{\mathcal{F}[Q_0]_d}{\int Q(\kappa)\ \mathrm{d}\kappa} = \mathrm{e}^{\mathrm{i}2\pi d\overline{\kappa}}\boldsymbol{\gamma}(d), \tag{2–6}$$

we can rewrite equation 2–5 as:

$$\begin{aligned} I_{\mathrm{norm}}(d) &= 1 + \Re\left\{\mathbf{V}(d)\mathrm{e}^{-\mathrm{i}2\pi d\overline{\kappa}}\right\} \\ &= 1 + V(d)\cos(-2\pi d\overline{\kappa} + \phi_V(d)). \end{aligned} \tag{2–7}$$

This clearly shows the form of the fringe. Over large ranges of d, the fringe appears like a 'carrier wave', given by the cosine term, that is slowly modulated in phase and amplitude by an envelope $\mathbf{V}$ (also known as a 'coherence envelope', Lawson 2000). Over the length of the slit direction on the detector, we sample only a very small range of delays, so that $d_0 - \Delta d/2 \leq d \leq d_0 + \Delta d/2$, where d_0 is determined by the interferometer etalon, as before. Over this range, the variation in $\mathbf{V}$ is negligible, so we see only a uniform sinusoid. In measuring the phase and visibility of the fringe, we essentially make a measurement of $\mathbf{V}$ at the fixed delay $d = d_0$. The phase offset of the sinusoid is determined by the argument of $\mathbf{V}$, ϕ_V. The measured (absolute) fringe visibility is simply the amplitude of the normalised fringe, V.

We note that, as shown by equation 2–6, $\mathbf{V}$ and $\boldsymbol{\gamma}$ are very closely related, the only difference being a phase offset, which, for a given wavelength channel j at wavenumber $\overline{\kappa}_j$

and fixed delay $d = d_0$, is constant – that is to say, $V = \gamma$ and $\phi_V = \phi_\gamma + 2\pi d_0 \overline{\kappa}_j$. Since the instrument is to be used purely for differential measurements, the zero point from which phases are measured is somewhat arbitrary and has no physical significance: we are concerned with *changes* in phase over time, which will affect both $\mathbf{V}$ and $\boldsymbol{\gamma}$ in the same way. For the analyses presented hereafter, the difference between $\mathbf{V}$ and $\boldsymbol{\gamma}$ is therefore not of great significance, and either can equally well be thought of as the complex visibility. However, for the sake of consistency, $\boldsymbol{\gamma}$ is generally intended by the term.

Figure 2-2 demonstrates conceptually how the summation of 100% visibility fringes due to a narrow band of wavelengths sums to give a modulated sinusoid. A plot is shown of the measured fringe intensity against delay, a plot known as an *interferogram*. In this case the contributing spectrum is white light through a rectangular bandpass, i.e. a simple top-hat function. The Fourier transform of a top-hat function centred at $\kappa = 0$ is a sinc function, and this is the shape of the modulation seen.

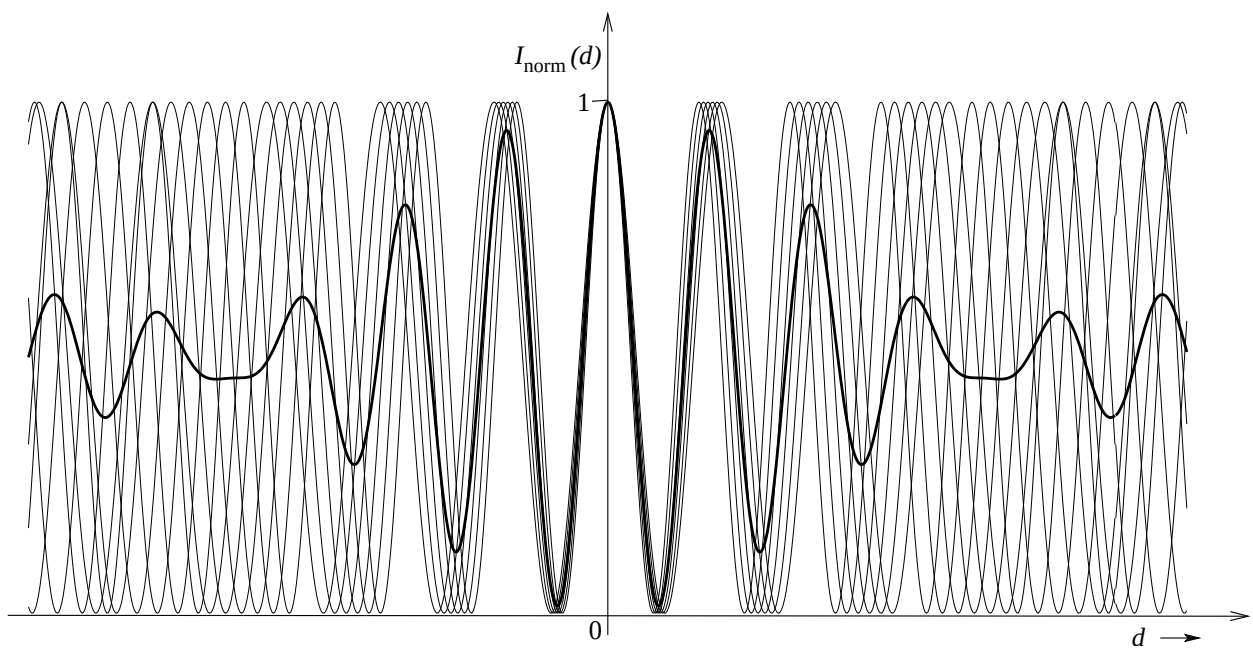

Figure 2-2. A representation of the summation (thick line) of fringes due to each individual wavelength of white light (thin lines) passing through a narrow rectangular bandpass (thin lines), resulting in a sinc-modulated sinusoid. (After Lawson, P. R. 2000, in Principles of Long Baseline Stellar Interferometry, ed. P. R. Lawson (Pasadena: NASA JPL), ch. 8, fig. 8.2.)

In figure 2-3, the amplitude of the modulating coherence envelope, $\mathbf{V}$, is shown more explicitly, and we see how measuring the fringe over a narrow range of delays around Δd around d_0 gives an approximately uniform sinusoid. This corresponds directly to the image seen along the length of the slit direction in a given wavelength channel on the detector (see e.g. 2-1F). Again, here the case is shown for white light through a rectangular bandpass, so that $\mathbf{V}(d)$ is a sinc function, with zeroes at $d = n/\Delta\kappa$ ($n \in \mathbb{Z}^+$), which modulates a sinusoid of period $1/\overline{\kappa}$. Since the LSF is at least theoretically an image of the slit, a top-hat is a reasonably good representation of the LSF, and therefore also of the response function of a wavelength channel (see appendix A). In practice the passband, $\Delta\kappa/\overline{\kappa}$, will be much narrower than indicated in the figure, so that the variation of $\mathbf{V}$ will be much slower compared to the sinusoid, and the sinusoid itself highly uniform over Δd.

For a more complicated input spectrum, such as that from a star, the coherence envelope will generally also have a more complicated shape, though the variations will still be slow in d. Each wavelength channel will have its own unique piece of spectrum contributing to it, and therefore each will have its own particular phase and visibility. It is this that gives rise to the varied patterns of fringes that are seen in the final fringing stellar spectra (e.g. figure 2-1F).

An alternative approach to understanding the formation of the fringing spectrum is to think of it as an infinite-resolution fringing spectrum, given by the product of the input spectrum and an interferometer transmission function, convolved with the LSF of the spectrograph. This is the approach adopted by Erskine (2003) and Mahadevan (2006). A derivation relating the two approaches is outlined in appendix B.

2.3 The Interferometer Comb

We have now laid the groundwork to easily understand the form of the interferometer comb as it is seen (if the instrument is appropriately configured) on the detector. The spectrum $Q_j(\lambda)$ contributing to a wavelength channel j is the product of the spectrograph response function $w_j(\lambda)$ for that channel, and the complete input spectrum, $P(\lambda)$. For

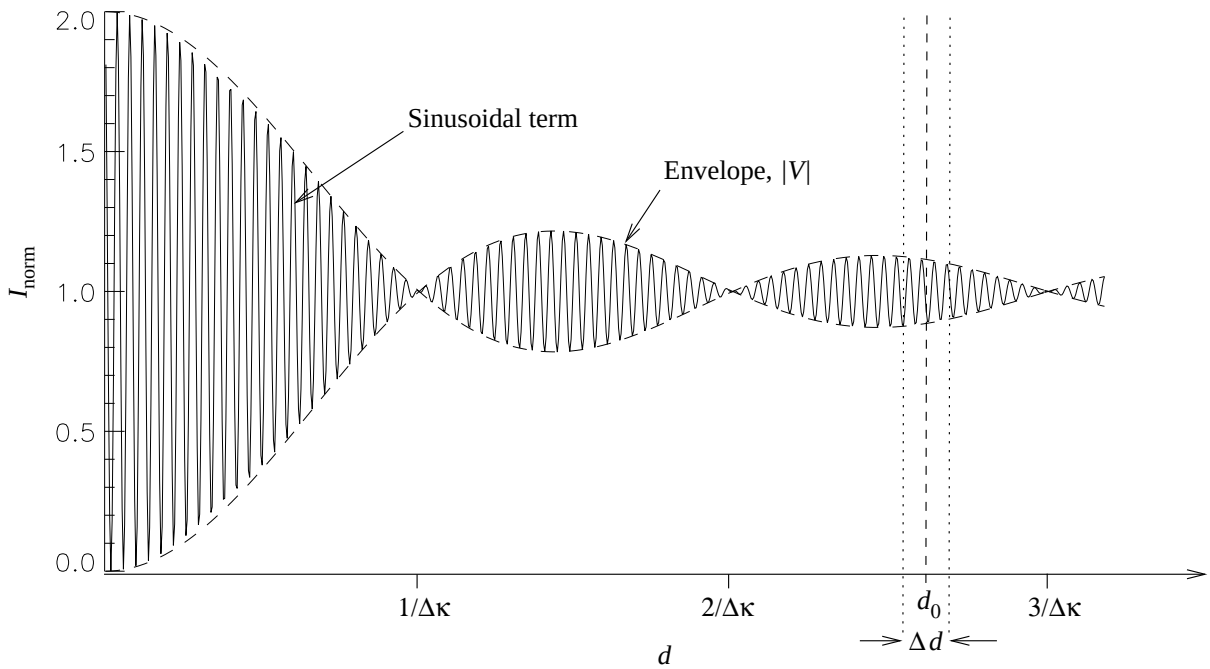

Figure 2-3. Interferogram showing the coherence envelope due to a rectangular band pass modulating the sinusoidal fringe. Along the slit direction of a fringing spectrum, a very small part of the interferogram is sampled over the range $d_0 \pm \Delta d/2$.

an idealised infinite resolution spectrograph, w_j is a delta function. Shining white light into the instrument, so that $P(\lambda) = 1$, leads to Q_j also being a delta function for all j. By equation 2–6, the coherence envelope, $\mathbf{V}(d)$ is the normalised Fourier transform of this delta function shifted to $d = 0$, so that $\mathbf{V}(d) = 1$ at all delays. Equation 2–7 then gives the very simple form of the resulting interferogram:

$$I_{\rm norm} = 1 + \Re\{\mathrm{e}^{\mathrm{i}2\pi d\kappa}\} \tag{2–8}$$

where we have stopped representing κ as a mean value since the width of the channel is negligible, and κ is now effectively used interchangeably with j as an index to position in the dispersion direction on the detector. This 'infinite resolution' interferogram is depicted in figure 2-4, plotting contours of intensity on the detector as a function of wavelength $\lambda = 1/\kappa$ on the x-axis versus delay in the y direction. Since λ maps linearly to x position on the detector (at least for an ideal spectrograph), and delay maps linearly

to y position, this also represents the image that would be seen on the detector if the full ranges were sampled down to zero wavelength and zero delay. The box in the figure schematically represents the segment of the interferogram that we actually observe with the instrument: a series of tilted parallel fringes, with a very slow wavelength dependency. For clarity, the figure is not to scale. In practice, the delay is fixed to a much larger value so that the fringes are observed at much higher order, n, and the wavelengths observed are much longer, so that any real observed comb is much denser and more uniform, and the wavelength dependency much smaller.

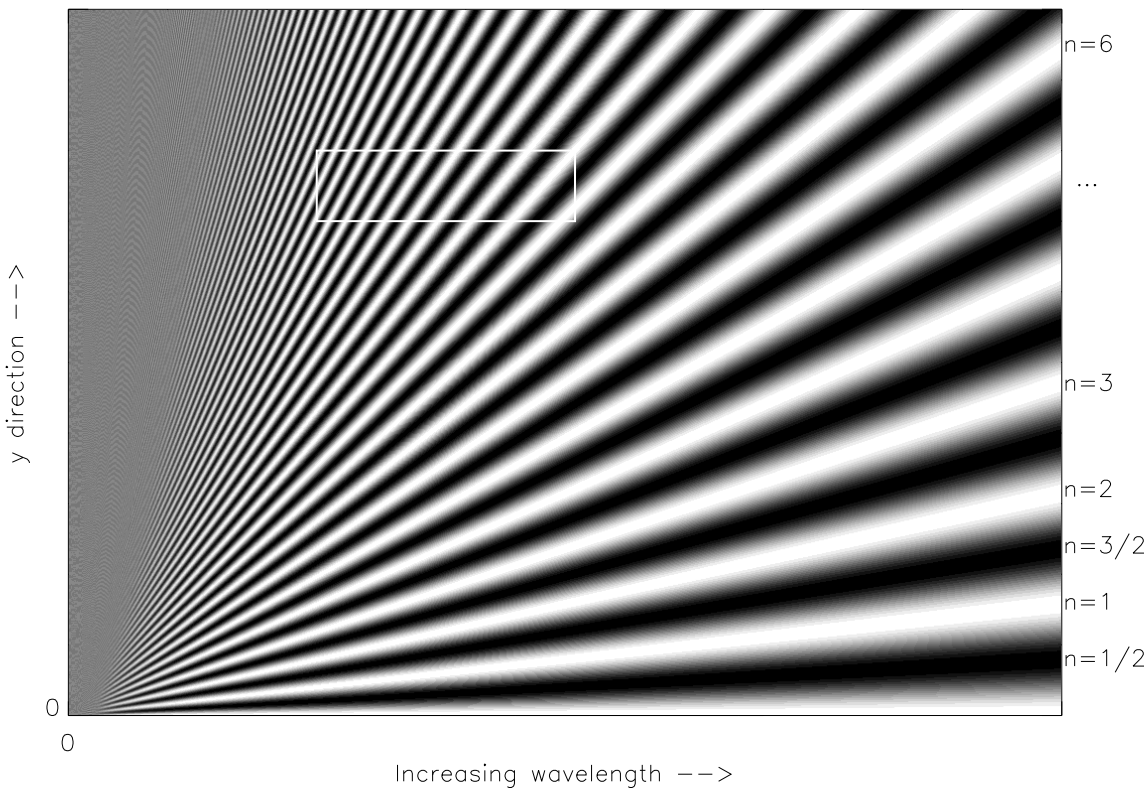

Figure 2-4. Simulated interferometer comb. Setting a large interferometer delay and choosing the wavelength range over which the spectrum is observed selects a 'window' in the comb (shown schematically) where the fringes are approximately parallel. The orders of some of the fringes, n, are shown down the right hand side. (In practice the 'window' chosen is at much longer wavelength and much higher order.)

In reality, of course, the spectrograph does not have infinite resolution. If we widen the delta-function response function so that w is now a top-hat function, we approach the

real situation: now $Q(\lambda)$ becomes a top-hat as well, and we have exactly the situation depicted in figure 2-3, where the fringes are modulated by a slow sinc envelope. In the segment of the interferogram imaged by the detector, the envelope is close to constant in value over the length of the slit. Since the top-hat function is symmetric and real, its sinc-shaped Fourier transform is also real, so there is no change in phase of our interferogram. The visibility is simply reduced according to the width of w, $\Delta\kappa$, and according to where the interferometer delay d_0 is set. We can see from this that by appropriately choosing the delay and spectrograph slit width we can null out the interferometer comb by finding a minimum in the envelope. Early experiments changing the slit width and delay with ET prototypes did indeed show this kind of sinc-type variation in the comb visibility.

2.4 From Phase to Velocity

To recap, in general, for a given channel j on the detector, the complex visibility of the measured fringe is given as in equation 2–3 (or see Goodman 1985, ch. 5) by:

$$\boxed{\boldsymbol{\gamma} = \frac{\mathcal{F}[P_\kappa w_{\kappa j}]_{d=d_0}}{\mathcal{F}[P_\kappa w_{\kappa j}]_{d=0}} = \frac{\mathcal{F}[P_\nu w_{\nu j})]_{\tau=\tau_0}}{\mathcal{F}[P_\nu w_{\nu j}]_{\tau=0}}} \tag{2–9}$$

where $\boldsymbol{\gamma}$ is the complex visibility (or complex degree of coherence), a vector quantity whose phase represents the phase of the measured fringe, and whose magnitude (from 0 to 1) represents the absolute visibility of the measured fringe; $\mathcal{F}[\ldots]$ represents a Fourier transform evaluated at interferometer path difference d (usually referred to here simply as 'delay'), or time delay τ, where $d = c\tau$ and c is the speed of light; P is the input spectrum; and w_j is the response function for that particular channel on the detector. We take d to be fixed at a value d_0 (for the purposes of the calculations here, the small difference in d across the length of a sinusoidal fringe is of no consequence). Subscripts are added to explicitly indicate functions of wavenumber, κ, or optical frequency, $\nu = c\kappa$: we note that the equation is completely equivalent in κ space with d as the conjugate variable, or in ν space with τ as the conjugate variable. In general the form being used will be

implicit from the context, so we usually drop these subscripts. We have also replaced the integral over the flux in the denominator with the Fourier transform at zero delay, which is mathematically equivalent (this important fact is made use of a number of times later on in this analysis). All the necessary mathematics for determining Doppler shifts and for dealing with the combination of the star and fiducial reference spectra (see section 2.6) derive from this formula.

The key to the DFDI radial velocity technique is the fact that Doppler shifts of the spectrum result in directly proportionate phase shifts of the fringes. This is a direct consequence of the Fourier shift theorem (D. J. Erskine, private communication). If the spectrum shifts such that $P(\kappa) \rightarrow P'(\kappa) = P(\kappa + \Delta\kappa)$, and we correctly follow the shift in the dispersion direction so that we now compare to the wavelength channel corresponding to $w_{j+\Delta j} = w_j(\kappa + \Delta\kappa)$ (assuming that the response function maintains the same form in nearby channels, and noting that Δj is not necessarily an integer), then the shift theorem gives:

$$\boldsymbol{\gamma}' = \frac{\mathcal{F}[P(\kappa + \Delta\kappa)w_j(\kappa + \Delta\kappa)]_{d=d_0}}{\mathcal{F}[P(\kappa + \Delta\kappa)w_j(\kappa + \Delta\kappa)]_{d=0}} = e^{\mathrm{i}2\pi\Delta\kappa d_0}\frac{\mathcal{F}[P(\kappa)w_j(\kappa)]_{d=d_0}}{\mathcal{F}[P(\kappa)w_j(\kappa)]_{d=0}} = e^{\mathrm{i}2\pi\Delta\kappa d_0}\boldsymbol{\gamma}. \qquad (2\text{–}10)$$

In other words, we have a phase shift of $\Delta\phi = 2\pi d_0 \Delta\kappa$, and by comparing the measured phase of the new fringes $\boldsymbol{\gamma}'$ with the previously unshifted ones, $\boldsymbol{\gamma}$, it is thus possible, in this simple case where there is no superposed reference spectrum, to derive the Doppler shift *without any explicit knowledge of the underlying high resolution spectrum, or of the spectrograph LSF.* Using the Doppler shift equation $\Delta\kappa/\kappa \approx -\Delta v/c$, where v represents velocity, conventionally positive in the direction away from the observer, we can write:

$$\Delta\phi = 2\pi d_0 \Delta\kappa = -\frac{2\pi d_0 \kappa \Delta v}{c} = -\frac{2\pi d_0}{c\lambda}\Delta v \equiv \frac{\Delta v}{\Gamma}, \qquad (2\text{–}11)$$

where, Γ, the phase-velocity scaling factor which gives the proportionality between phase shift and velocity shift, is defined as:

$$\boxed{\Gamma \equiv -\frac{c\lambda}{2\pi d_0}.} \tag{2–12}$$

2.5 Calculating the Interferometer Delay

The interferometer delay, d_0, is determined by the etalon in the interferometer. The best precision that can be obtained in RV measurements is a tradeoff between maximising the phase-velocity scale Γ (so that a large phase shift for a small change in velocity) and maximising the visibility of the fringes (since higher visibility means more accurate measurements of the fringes). Since the visibility of the fringes is determined by the match between d_0 and the typical spectral line widths to be observed, an optimal value of d_0 can be chosen to give the best precision for the expected typical targets for the survey (Ge 2002). This is set at design time, and remains fixed for the instrument.

Annual variations in the RV of a star due to the Earth's motion can be as large as $60\,\mathrm{km\,s^{-1}}$, even for an RV-stable target.[1] If we are to consider approaching precisions on the order of $1\,\mathrm{m\,s^{-1}}$ we therefore need to know Γ to at least one part in 60,000. Since Γ depends on the interferometer delay (equation 2–12), determining Γ is synonymous with measuring the delay.

To a first approximation, the delay can be calculated from the properties of the etalon. The etalon has two effects: first, it produces a path length difference due to the fact that the wavelength of the light travelling through it is reduced by a factor $1/\eta$, where

[1] This figure simply comes from the peak-to-peak change in orbital velocity of the Earth along the line of sight towards a target in the plane of the Earth's orbit over the course of a year, which dominates strongly over the velocity due to diurnal rotation on annual timescales.

η is its index of refraction. This gives a path difference compared to air of:

$$2t\eta - 2t, \tag{2–13}$$

where t is the thickness of the etalon. The factor of two is introduced since the light must pass through the etalon twice as it travels toward and then away from the mirror.

The second effect of the etalon is to reduce the apparent distance between the mirror and the etalon, by an amount $t(1 - 1/\eta)$ where t is the thickness of the etalon. Since the virtual image formed must coincide with the image of the unimpeded mirror in the other interferometer arm, the mirror in the etalon arm must be moved back by this distance to compensate. Hence the path length in wavelengths is increased by an amount:

$$2t(1 - 1/\eta), \tag{2–14}$$

where again the factor of two is due to passing twice through the etalon.

Adding these two effects (2–13 and 2–14) we obtain the total path difference due to the etalon:

$$d_0 = 2t(\eta - 1/\eta). \tag{2–15}$$

This derivation depends on the assumption that there is negligible dispersion in the etalon glass, i.e., that η is close to independent of wavelength over the wavelength range of interest. Dispersion can in fact be a significant effect, but the assumption should be good to a few percent (Barker & Schuler 1974, D. J. Erskine, private communication), enough for an initial estimate.

A more precise measure of the delay can be determined simply by counting fringes in the interferometer comb. We know from equation 2–8 that the phase of the comb varies as

$2\pi d\kappa = 2\pi d/\lambda$.[2] Differentiating with respect to wavelength:

$$\frac{\partial \phi}{\partial \lambda} = \frac{1}{2\pi}\frac{\partial n}{\partial \lambda} = -\frac{2\pi d}{\lambda^2}, \tag{2–16}$$

where n is the fringe order, giving:

$$d = -\lambda^2 \frac{\partial n}{\partial \lambda}. \tag{2–17}$$

In other words, by counting the fringe density over wavelength, we can immediately calculate d_0, and hence Γ. Since there is a λ^{-2} dependence in $\partial n/\partial \lambda$, care needs to be taken to account for the dependence properly when determining the fringe density at a given wavelength. This may be more easily done in wavenumber space instead, since the fringe density is uniform with wavenumber, and $d = \partial n/\partial \kappa$.

In practice, counting fringes is often not easy, since the comb is often barely resolved (usually by design). As long as the comb is not under-sampled on the detector, this can be overcome by temporarily using a narrower slit in the spectrograph, since in principle the delay should only need to be determined once. Even so, it is usually possible in practice only to count over a few hundreds to a thousand or two fringes, giving at best an accuracy on the order of one part in 1000. Over a $60\,\mathrm{km\,s^{-1}}$ variation, this is still only good to the $60\,\mathrm{m\,s^{-1}}$ level.

Other methods of measuring the delay are under investigation, but to date, the method of choice has been simply to observe known stable reference stars over the time baseline of interest and use their known apparent changes in velocity due to the Earth's motion to calibrate Γ. Provided the reference stars are genuinely stable, and they are positioned in the sky such that their barycentric motions are large, this technique will

[2] Although this is for a comb at infinite resolution, the same variation will hold true at lower resolutions: a spectrograph response function broader than a delta function will only reduce the visibility of the interferogram, and possibly add an overall phase offset to the entire interferogram (provided that the shape of the response function is uniform across the detector).

provide an accuracy in the determination of Γ at least equal to the intrinsic RV stability of the stars.[3]

2.6 Handling a Fiducial Reference Spectrum

The extremely high sensitivity of the instrument means that numerous instrumental effects can masquerade as velocity shifts. Tiny changes in the interferometer cavity due to thermal flexure, for example, will shift the interferometer comb, appearing as shifts in the fringe pattern. (Note that from equation 2–11, a change in d_0 conveniently has mathematically exactly the same effect as a change in velocity, Δv.) The image can also shift on the detector in both the slit and the dispersion directions. To account for these instrumental artifacts, an absorption reference is inserted into the optical path – in the case of the ET instruments, a glass cell filled with iodine vapour maintained at a fixed temperature. In this way an iodine spectrum can be multiplied with the stellar spectrum. Since the iodine is stationary with respect to the instrument, its spectrum will track instrument shifts, which can then be subtracted from the measured stellar shift to reveal the star's intrinsic motion. To do this, for each target to be observed, two fringing 'template' spectra are taken, one being pure star with no reference in the beam path, and the other pure reference – i.e. a pure iodine spectrum taken by shining a tungsten continuum lamp through the cell. These templates are then used to separate out the stellar and iodine components of the combined star-iodine data (referred to here as 'data' or 'measurement' frames, as distinct from 'template' frames). A formalism is required to extract the iodine and stellar spectra from the combined spectrum. In order to proceed, we define the following symbols:

- j – the pixel number in the dispersion direction which identifies the column along which a fringe is measured in the slit direction, corresponding to a single wavelength channel. Strictly speaking, the wavelength channel is infinitesimally wide on the

[3] In fact the issue can be further complicated by other systematic errors in the measurements, especially the iodine addition approximation, discussed in section 2.6.1.

detector, so that j need not necessarily be an integer. Since the spectrum is oversampled, however, it is often a reasonable simplification to think of the entire pixel column representing an infinitesimal sample in the dispersion direction (see appendix A).

- $\mathbf{M}(j)$ – the complex visibility vector (ie. phase and absolute visibility) for a fringe at wavelength channel j in a single Doppler measurement frame of combined star-iodine data, an ensemble of such values for a spectrum across all j comprising a 'whirl.'

- $\mathbf{S}(j)$ – the measured complex visibility for the star template at channel j.

- $\mathbf{I}(j)$ – the measured complex visibility for the iodine template at channel j.

- $\mathcal{M}(\lambda) \equiv C_{\mathrm{m}}(\lambda)M(\lambda)$ – the input spectrum for a combined star/iodine data frame, where C_{m} is the continuum function and includes an overall normalisation factor for M, which is the continuum-corrected normalised spectral density, equal to one wherever there are no absorption lines. C_{m} is assumed constant to a good approximation over the scale of the width of the response function w (see below) and instrument LSF, and $0 \leq M \leq 1$.

- $\mathcal{S}(\lambda) \equiv C_{\mathrm{s}}(\lambda)S(\lambda)$ – the same for the star template spectrum.

- $\mathcal{I}(\lambda) \equiv C_{\mathrm{i}}(\lambda)I(\lambda)$ – the same for the iodine template spectrum.

- $s(\lambda), i(\lambda)$, such that $S \equiv 1 - s,\ I \equiv 1 - i;\ 0 \leq (s, i) \leq 1$.

- $w(j, \lambda)$ – the response function at position j on the detector, by which is meant the spectrum that contributes to an infinitesimally wide wavelength channel at the detector plane if perfect continuum light is passed through the instrument. (Note that w is very closely related to the instrument LSF – see appendix A)

- d – the interferometer delay, fixed to a value of $d = d_0$, as usual.

- Γ – phase/velocity scaling constant, also as usual.

We assume for now the case where there is no Doppler or instrument shift in either phase or in the dispersion direction, for both star and iodine components, and no photon shot noise. Here the aim is simply to reconstruct the data whirl from the two template whirls. Once this is achieved, it is conceptually a relatively trivial step to allow for shifted and noisy data: the template whirls need only to be shifted iteratively in phase and in the

dispersion direction until a best fit solution is found, allowing the intrinsic stellar Doppler shift to be directly calculated. This can be done using any standard least-squares method.

Following equation 2–9, the complex visibility measured at detector channel j for the templates and combined star-iodine data can be written exactly as:

$$\mathbf{S} = \frac{\mathcal{F}[\mathcal{S}w]_{d_0}}{\mathcal{F}[\mathcal{S}w]_0} = \frac{[\widehat{\mathcal{S}} \otimes \widehat{w}]|_{d_0}}{[\widehat{\mathcal{S}} \otimes \widehat{w}]|_0}, \tag{2–18}$$

$$\mathbf{I} = \frac{\mathcal{F}[\mathcal{I}w]_{d_0}}{\mathcal{F}[\mathcal{I}w]_0} = \frac{[\widehat{\mathcal{I}} \otimes \widehat{w}]|_{d_0}}{[\widehat{\mathcal{I}} \otimes \widehat{w}]|_0}, \tag{2–19}$$

$$\mathbf{M} = \frac{\mathcal{F}[\mathcal{M}w]_{d_0}}{\mathcal{F}[\mathcal{M}w]_0} = \frac{\mathcal{F}[\mathcal{S}\mathcal{I}w]_{d_0}}{\mathcal{F}[\mathcal{S}\mathcal{I}w]_0} = \frac{[\widehat{\mathcal{S}} \otimes \widehat{\mathcal{I}} \otimes \widehat{w}]|_{d_o}}{[\widehat{\mathcal{S}} \otimes \widehat{\mathcal{I}} \otimes \widehat{w}]|_0}. \tag{2–20}$$

The problem lies in expressing equation 2–20 in terms of 2–18 and 2–19. This is made difficult by the convolutions, which appear to require knowledge of the template spectra at all possible values of the delay d in order to be evaluated. The nature of ET is such that we measure it only at one value, d_0. An additive approximation is currently used to address this problem, which is described in section 2.6.1.

2.6.1 The Addition Approximation

It is possible to rewrite the input spectrum as:

$$\begin{aligned}
\mathcal{M} &= A\mathcal{S}\mathcal{I} \\
&= AC_{\mathrm{s}}C_{\mathrm{i}}SI \equiv C'SI \\
&= C'(1-s)(1-i) \\
&= C'(1-s+1-i-1+si) \\
&= C'(S+I-1+si),
\end{aligned} \tag{2–21}$$

where A is a scaling constant to allow for difference in total flux level between the templates and data, and $C' \equiv AC_{\mathrm{s}}C_{\mathrm{i}}$ is a constant over the width of the response function. If we assume that either s or i or both $\ll 1$, then the 'crosstalk' term, si, can be neglected. Since i and s essentially represent line depths, this means that we are assuming either very

shallow lines, or no significant overlap between lines in the two different spectra. Keeping the cross talk term in place for now for completeness, we can continue, substituting equation 2–21 in equation 2–20:

$$\mathbf{M} = \frac{\mathcal{F}[\mathcal{M}w]_{d_0}}{\mathcal{F}[\mathcal{M}w]_0} = \frac{[\widehat{Sw} + \widehat{Iw} - \widehat{w} + \widehat{siw}]|_{d_0}}{[\widehat{Sw} + \widehat{Iw} - \widehat{w} + \widehat{siw}]|_0}. \tag{2–22}$$

The factor C' has cancelled because it is constant over the width of the response function, and therefore can be taken outside the Fourier transforms. The denominator of this equation represents a normalisation, corresponding to the total flux in channel j on the detector. The term $\widehat{w}|_{d_0}$ in the numerator is due to the interferometer comb, since if white light is passed through the instrument, then $S = I = 1$, and the cross talk term vanishes. We are then left with:

$$\mathbf{M}_{\text{continuum}} = \widehat{w}|_{d_0}/\widehat{w}|_0, \tag{2–23}$$

which describes the interferometer comb. As expected, the properties of the comb are determined purely by the response function.

Substituting equations 2–18 and 2–19 into 2–22 we can write:

$$\mathbf{M} \quad = \quad K_{\text{s}}\mathbf{S} + K_{\text{i}}\mathbf{I} \quad + \quad \frac{-\widehat{w}|_{d_0} + \widehat{siw}|_{d_0}}{\widehat{Sw}|_0 + \widehat{Iw}|_0 - \widehat{w}|_0 + \widehat{siw}|_0}, \tag{2–24}$$

where the scalar quantities K_{s} and K_{i} are given by:

$$K_{\text{s}} \equiv \frac{\widehat{Sw}|_0}{\widehat{Sw}|_0 + \widehat{Iw}|_0 - \widehat{w}|_0 + \widehat{siw}|_0}, \qquad K_{\text{i}} \equiv \frac{\widehat{Iw}|_0}{\widehat{Sw}|_0 + \widehat{Iw}|_0 - \widehat{w}|_0 + \widehat{siw}|_0}. \tag{2–25}$$

Hence we see that we can now represent the combined star+iodine data in terms of a linear combination of the measured star and iodine templates, along with an error term.

The fraction on the right in equation 2–24 contains two terms in the numerator, the comb term, $\widehat{w}|_{d_0}$, and a cross talk term, $\widehat{siw}|_{d_0}$. It is in principle possible to arrange the instrument such that at delay $d = d_0$ the interferometer comb has zero visibility, by choosing the delay and slit width so that $\mathbf{M}_{\text{continuum}}$ is at a zero point of $\widehat{w}$. Alternatively, it is possible to low-pass Fourier filter the data image before measuring the whirls,

essentially simulating a lower spectrograph resolution. In either case, we assume that $\widehat{w}|_{d_0} \to 0$. If we now also neglect all the cross talk terms si following from equation 2–21, we finally have the whirl addition approximation, which we can write:

$$\mathbf{M} \approx K_{\mathrm{s}}\mathbf{S} + K_{\mathrm{i}}\mathbf{I}. \tag{2–26}$$

K_{s} and K_{i} represent scaling factors that need to be allowed for in the absolute visibilities of the two templates in order to obtain a good solution. Remembering that the evaluation of a Fourier transform at $d = 0$ represents the total integrated area under the function, we can try to gain a handle on the expected sizes of these scaling factors. To the extent that the total area under Sw and Iw is not much less than that under w (i.e. that the area in discrete absorption lines is small, or $\int_{\Delta w} s\,\mathrm{d}\lambda \ll 1$ and $\int_{\Delta w} t\,\mathrm{d}\lambda \ll 1$, where Δw is a representative width of the response function), equation 2–25 shows that we would expect $K_{\mathrm{s}}, K_{\mathrm{i}} \approx 1$. As far as this approximation holds good, and to the extent that K_{s} and K_{i} are approximately constant across all channels j, it is then a simple matter to allow for Doppler and instrument drift by allowing the template whirls to rotate in phase and shift in the dispersion direction as a function of j, and minimising χ^2 in the residuals to find the best fit solution for the measured data $\mathbf{M}$.

However, we note that in fact the assumptions made regarding K_{s} and K_{i} are actually not likely to be terribly good – particularly given that the iodine fiducial cell typically absorbs a total of $\sim$ 40% of the incident light. Furthermore, there is little reason to assume that K_{s} and K_{i} should be constant from channel to channel. We can recast them, rewriting equation 2–25 as:

$$K_{\mathrm{s}} = \frac{\widehat{Sw}|_0}{\widehat{Mw}|_0} = \frac{\mathcal{F}[(\mathcal{S}/C_{\mathrm{s}})w]_0}{\mathcal{F}[(\mathcal{M}/C_{\mathrm{m}})w]_0} = \frac{C_{\mathrm{m}}}{C_{\mathrm{s}}}\frac{\widehat{\mathcal{S}w}|_0}{\widehat{\mathcal{M}w}|_0}, \qquad K_{\mathrm{i}} = \frac{C_{\mathrm{m}}}{C_{\mathrm{i}}}\frac{\widehat{\mathcal{I}w}|_0}{\widehat{\mathcal{M}w}|_0}. \tag{2–27}$$

We now see that they are calculable in terms of measurable quantities, namely the total fluxes in each channel j for the templates and the data. The continuum functions $C_{\mathrm{m}}, C_{\mathrm{s}}$ and C_{i} are in principle derivable by simply continuum fitting a smooth function to the

measured fluxes. Since the functional forms of K_s and K_i are a relatively recent discovery, the two scaling factors have been considered constant across all channels to date in the pipeline, and this has worked quite well: apparently the values average out across the spectrum so that the effects of the assumption are not too drastic. However, a fully rigorous treatment requires their inclusion even in the absence of any cross talk term.

2.6.2 An Alternative: Combined-Beam Reference

Even allowing for the visibility scaling factors, the above analysis is useful only in as far as the approximation that the cross talk, si, is very small holds well. It is starting to appear, however, that this approximation is in fact not accurate enough for our purposes (see section 5.10). If an exact solution can be found, it would solve the addition approximation problem. Finding a better approximation may also be a possibility, but we have been unable to find either at the time of writing.

One possible solution to the problem is to actually physically superpose a fiducial ThAr or tungsten-illuminated iodine spectrum on top of the stellar target spectrum, for example by splicing two fibres into one, one coming from the telescope and one from the reference lamp. In this case, the two spectra now combine additively instead of multiplicatively. We can then write:

$$\mathcal{M} = A_\mathrm{s}\mathcal{S} + A_\mathrm{i}\mathcal{I}$$

where A_s and A_i are scaling factors to allow for flux differences between the templates and data (note that two such factors are now required). Once again, following equation 2–9 we can now write:

$$\begin{aligned}
\mathbf{M} &= \frac{\mathcal{F}[\mathcal{M}w]_{d_0}}{\mathcal{F}[\mathcal{M}w]_0} \\
&= \frac{\mathcal{F}[(A_\mathrm{s}\mathcal{S} + A_\mathrm{i}\mathcal{I})w]_{d_0}}{\mathcal{F}[(A_\mathrm{s}\mathcal{S} + A_\mathrm{i}\mathcal{I})w]_0} \\
&= \frac{A_\mathrm{s}\widehat{\mathcal{S}w}|_{d_0} + A_\mathrm{i}\widehat{\mathcal{I}w}|_{d_0}}{A_\mathrm{s}\widehat{\mathcal{S}w}|_0 + A_\mathrm{i}\widehat{\mathcal{I}w}|_0} \\
&= K'_\mathrm{s}\mathbf{S} + K'_\mathrm{i}\mathbf{I}, \qquad (2\text{–}28)
\end{aligned}$$

where:

$$K_{\mathrm{s}}' = \frac{A_{\mathrm{s}}\widehat{\mathcal{S}w}|_0}{A_{\mathrm{s}}\widehat{\mathcal{S}w}|_0 + A_{\mathrm{i}}\widehat{\mathcal{I}w}|_0}, \qquad K_{\mathrm{i}}' = \frac{A_{\mathrm{i}}\widehat{\mathcal{I}w}|_0}{A_{\mathrm{s}}\widehat{\mathcal{S}w}|_0 + A_{\mathrm{i}}\widehat{\mathcal{I}w}|_0}, \tag{2–29}$$

or:

$$K_{\mathrm{s}}' = A_{\mathrm{s}}\frac{\widehat{\mathcal{S}w}|_0}{\widehat{\mathcal{M}w}|_0}, \qquad K_{\mathrm{i}}' = A_{\mathrm{i}}\frac{\widehat{\mathcal{I}w}|_0}{\widehat{\mathcal{M}w}|_0}. \tag{2–30}$$

We see that we now have an exact expression for $\mathbf{M}$, with the difference being that we now need to take into account the flux scaling factors A_{s} and A_{i}, where previously the flux scaling factor cancelled.

It is interesting to note that if we multiply through both sides of equation 2–28 by the denominator, $\widehat{\mathcal{M}w}|_0$ (which represents the total flux along the channel in the combined data), we essentially find we have an expression which is a summation of flux $\times$ visibility terms. Since visibility is defined as $(I_{\mathrm{max}} - I_{\mathrm{min}})/(I_{\mathrm{max}} + I_{\mathrm{min}})$, where I_{max} and I_{min} are the maximum and minimum fringe intensities, then multiplying by total flux in the channel gives a quantity equal to the amplitude of the fringe. Hence equation 2–28 is really simply summing fringe amplitudes, and is exactly what we expect when the two input spectra are combined additively: the resulting image on the detector should simply be a direct intensity summation of the respective images that would be obtained individually.

2.7 Photon Error Propagation

2.7.1 Photon Error for Multiplied Reference

It is useful to have an estimate of the level of error expected purely from photon shot noise to be able to assess instrument performance. The photon error in the phase measurement (and hence velocity measurement) from a single wavelength channel can be estimated following Ge (2002). This gives essentially:

$$\sigma_j \approx \frac{1}{\pi\sqrt{2}} \cdot \frac{c\lambda}{d\gamma_j\sqrt{F_j}} \tag{2–31}$$

where σ_j is the error in velocity due to channel j alone, λ is the wavelength of channel j, d is the optical delay, γ_j is the visibility of the fringe, and F_j is the total flux in the channel.[4]

To calculate the expected error in a radial velocity measurement for a single data frame, assuming the standard instrument configuration where an iodine spectrum multiplies the input stellar spectrum, we consider the resulting data spectrum as consisting of two components, a star component, and an iodine component. The calculated phase shift due to intrinsic target Doppler shift, $\Delta\phi$ is given by:

$$\Delta\phi = \langle \phi_{\mathrm{sm},j} - \phi_{\mathrm{st},j} \rangle \; - \; \langle \phi_{\mathrm{im},j} - \phi_{\mathrm{it},j} \rangle, \tag{2–32}$$

where $\langle \ldots \rangle$ represents a weighted mean over j, $\phi_{\mathrm{sm},j}$ and $\phi_{\mathrm{im},j}$ represent the phases for the star and iodine components of the combined star/iodine data ('measurement') frame and $\phi_{\mathrm{st},j}$ and $\phi_{\mathrm{it},j}$ are the phases measured in the separate pure star and iodine templates. For convenience, we immediately map these phases to corresponding 'velocity' measurements by multiplying both sides by Γ (though with the caveat that a velocity measurement of a single channel in a single spectrum has no physical meaning in itself until it is differenced with another spectrum):

$$\Delta v = \langle v_{\mathrm{sm},j} - v_{\mathrm{st},j} \rangle - \langle v_{\mathrm{im},j} - v_{\mathrm{it},j} \rangle, \tag{2–33}$$

Using σ with corresponding subscripts to represent the various errors in this equation, we might expect a total photon error in Δv to be given by:

$$\sigma^2_{\mathrm{photon}} = \left[E_j \left(\sqrt{\sigma^2_{\mathrm{sm}} + \sigma^2_{\mathrm{st}}} \right) \right]^2 + \left[E_j \left(\sqrt{\sigma^2_{\mathrm{im}} + \sigma^2_{\mathrm{it}}} \right) \right]^2, \tag{2–34}$$

[4] The small difference in the numerical factor in the denominator ($\pi\sqrt{2}$ vs. 4) is due to using the RMS slope of the fringe, rather than the mean absolute slope used in Ge (2002). Monte Carlo simulations of sinusoid fits suggest that the RMS slope gives more accurate results.

where $E_j(\sigma)$ represents the standard statistical error in a weighted mean:

$$E_j(\sigma) \equiv \frac{1}{\sqrt{\sum_j 1/\sigma_j^2}}. \tag{2–35}$$

In practice, the two template terms in equation 2–34 are neglected, for two reasons. The first is simply because in general the templates will have significantly higher flux than the data frame: the iodine template can be taken with arbitrarily high flux since it is obtained with a quartz lamp as a source; and the stellar template is usually deliberately taken with higher flux than the data so that it does not compromise the entire data set. The second reason is a little more subtle. All RV measurements with this kind of instrument are differential, measured relative to the two templates which effectively set the zero point of the measurements for the star and iodine, as seen in equation 2–33. Since this 'zero point' is the same for every RV measurement, any error in the zero point will not contribute to the RMS scatter in a set of measurements which uses the same templates. Since photon errors go as $1/\sqrt{\text{flux}}$, the remaining terms, $E_j(\sigma_{\rm sm})$ and $E_j(\sigma_{\rm im})$, can be estimated by scaling the respective template terms by the flux difference between the templates and data, giving:

$$\begin{aligned} \sigma^2_{\rm photon} &= [E_j(\sigma_{\rm sm})]^2 + [E_j(\sigma_{\rm im})]^2 \\ &\approx \frac{\overline{F}_{\rm st}}{\overline{F}_{\rm m}}\,[E_j(\sigma_{\rm st})]^2 + \frac{\overline{F}_{\rm it}}{\overline{F}_{\rm m}}\,[E_j(\sigma_{\rm it})]^2 \end{aligned} \tag{2–36}$$

where $\overline{F}_{\rm st}$, $\overline{F}_{\rm it}$ and $\overline{F}_{\rm m}$ represent the mean fluxes across the whole star template, iodine template and data frame respectively.[5] Alternatively, explicitly substituting equation 2–35

[5] This assumes that the flux ratio terms remain the same from channel to channel, so that an overall mean scaling can be applied. This is not strictly accurate (e.g. if line depths are very deep and broad, or the star and iodine continuum functions are very different), but is taken to work to a reasonable approximation. It is possible that the assumption needs revisiting. In the event that a more accurate calculation is needed, it is a simple enough matter to introduce channel-dependent flux ratios for each element j within the summations.

into equation 2–36:

$$\boxed{\sigma_{\text{photon}} = \sqrt{\frac{\overline{F}_{\text{st}}}{\overline{F}_{\text{m}}} \cdot \frac{1}{\sum_j 1/\sigma_{\text{st},j}^2} + \frac{\overline{F}_{\text{it}}}{\overline{F}_{\text{m}}} \cdot \frac{1}{\sum_j 1/\sigma_{\text{it},j}^2}},} \tag{2–37}$$

where $\sigma_{\text{st},j}$ and $\sigma_{\text{it},j}$ are obtained using equation 2–31.

It should be noted, however, that these formulae for the photon limit are for the values expected *given* the fringe visibility that was obtained. Various instrument effects – for example defocus – can reduce the visibility from its optimum and hence reduce the photon limiting precision from its optimum.

As an example, a typical data point taken from a KPNO 2.1m ET run (single-object) in January 2006 for 51 Peg ($V = 5.49\,\text{mag}$, 10 min exposure) with good flux gives mean signal/noise ratios (S/N) per pixel for star template, iodine template, and data frame of 106.6, 129.0 and 86.6 respectively. These values give photon errors for the star and iodine components of 5.6 and $4.5\,\text{m}\,\text{s}^{-1}$ respectively, which when added in quadrature give a total photon error of $7.2\,\text{m}\,\text{s}^{-1}$. This is for only one of the two output beams of the KPNO instrument: a weighted average over the two beams gives a final photon error of $4.8\,\text{m}\,\text{s}^{-1}$. It is interesting to note that the error due to the iodine reference is in fact comparable to that due to the star, since the signal in the iodine component of the data frame is intrinsically limited by the magnitude of the target being observed. This is possibly an additional argument for pursuing the method of combining star and reference beams additively rather than by the normal insertion of an absorption cell into the stellar beam path.

2.7.2 Photon Error for Reference Spectrum in Addition

In the case of the reference spectrum being combined additively, rather than multiplicatively, the photon errors must be calculated differently. However, we can follow a somewhat similar approach. Again, we consider the errors due to star and iodine components of the combined star+iodine data, and neglect the errors due to the

templates, so that, as for equation 2–36:

$$\sigma^2_{\text{photon}} = [E_j(\sigma_{\text{sm}})]^2 + [E_j(\sigma_{\text{im}})]^2, \tag{2–38}$$

where E_j is again defined as in equation 2–35. The individual components $\sigma_{\text{sm},j}$ and $\sigma_{\text{im},j}$ must be reevaluated, however, since the photon noise from the two separate sources will now combine additively (for example, if one of the sources is considerably brighter than the second, its photon noise will dominate over the signal in the second). We can think of an effective visibility for the two components in the combined data, $\gamma_{\text{sm},j}$ and $\gamma_{\text{im},j}$. Remembering that fringe amplitude is given by the product of the visibility and the mean flux in the fringe, we can write:

$$\gamma_{\text{sm},i} = \frac{\gamma_{\text{st},j} A_S F_{\text{st},j}}{F_{\text{m},j}} \quad ; \quad \gamma_{\text{im},j} = \frac{\gamma_{\text{it},j} A_I F_{\text{it},j}}{F_{\text{m},j}}, \tag{2–39}$$

where $\gamma_{\text{st},j}$ and $\gamma_{\text{it},j}$ are the fringe visibilities for channel j in the star and iodine templates respectively, $F_{\text{st},j}$ and $F_{\text{it},j}$ are the mean fluxes across the channel for the same, and $F_{\text{m},j}$ is the mean flux in the channel for the data (measurement) frame. A_{s} and A_{i} are scaling factors that allow for flux differences between the templates and the respective data components, as in section 2.6.2 (we take these to be independent of channel). Substituting these effective visibilities in equation 2–31 gives:

$$\begin{aligned} \sigma_{\text{sm},j} &= \frac{1}{\pi\sqrt{2}} \cdot \frac{c\lambda\sqrt{F_{\text{m},j}}}{d\gamma_{\text{st},j} A_{\text{s}} F_{\text{st},j}} \\ \sigma_{\text{im},j} &= \frac{1}{\pi\sqrt{2}} \cdot \frac{c\lambda\sqrt{F_{\text{m},j}}}{d\gamma_{\text{it},j} A_{\text{i}} F_{\text{it},j}}. \end{aligned} \tag{2–40}$$

Using these we can now evaluate equation 2–38 to obtain an estimate of the photon limiting error, so that:

$$\boxed{\sigma_{\text{photon}} = \frac{1}{\pi\sqrt{2}} \cdot \frac{c\lambda}{d} \sqrt{\left[E_j\left(\frac{\sqrt{F_{\text{m},j}}}{\gamma_{\text{st},j} A_{\text{s}} F_{\text{st},j}}\right)\right]^2 + \left[E_j\left(\frac{\sqrt{F_{\text{m},j}}}{\gamma_{\text{it},j} A_{\text{i}} F_{\text{it},j}}\right)\right]^2}.} \tag{2–41}$$

In this case, we do not attempt to assume channel independent flux ratios. This is because for additively combined references it becomes possible to consider using emission spectra (e.g. a ThAr lamp) as the reference, rather than the usual iodine absorption spectrum. Clearly the flux ratio between data and reference template frames is very different for regions where there are no reference emission lines compared to those where emission lines are present. It is therefore not reasonable to take the flux terms outside the summation in the error combination function E_j.

CHAPTER 3
INSTRUMENT HARDWARE

Two DFDI-type instruments have been built at Penn State University and the University of Florida: the single-object prototype installed at the Kitt Peak National Observatory 2.1m telescope, known as the Exoplanet Tracker (ET); and the full multi-object survey instrument designed to observe 60 objects at once, installed at the 2.5m Sloan Deep Sky Survey (SDSS) telescope at Apache Point Observatory and known as the W. M. Keck ET owing to major funding from the Keck Foundation. The instrument design and construction was performed primarily by engineers and other ET team members, so is not the focus of this thesis. However, a brief description of the main instrument subsystems relevant to data acquisition are provided here, along with a brief summary of the technical specifications for completeness.

3.1 Subsystem Overview

The ET instruments are both built on broadly the same model. Optical fibres feed light from the telescope to the instrument, which is built on a vibration-isolated optical bench, and kept in a thermally stabilised environment. In addition to the Michelson interferometer and the spectrograph outlined in chapter 2, there are several additional subsystems:

Calibration lamps. Tungsten continuum and ThAr emission line calibration lamps can be easily switched in and out in place of the telescope feed to provide the necessary calibration exposures for the data reduction pipeline.

Iodine cell. A motor-controlled sealed glass cell containing iodine vapour can be remotely switched into the beam path between the fibre input to the instrument and the interferometer to provide a fiducial reference (see section 2.6). The cell is actively heated to maintain a fixed temperature of $60 \pm 0.1\,^{\circ}\mathrm{C}$ so that the absorption spectrum remains stable.

Active phase locker. The delay in the interferometer is locked against thermal drifts by adjusting the position of one of the interferometer mirrors via an active feedback loop. A laser is fed through the interferometer at a slight angle to the main beam path, picked off at the output with fold mirrors, and fed to an ordinary video camera, where the image is sent to a computer. When the instrument is aligned properly, the image is simply a series of horizontal parallel fringes (like that shown schematically in figure 2-1B). The interferometer mirror in the arm without the etalon is set on a PZT actuated mount. Using software written in Labview by Curtis DeWitt and later upgraded by Pengcheng Guo, the computer feeds back continuously to the PZT actuators, adjusting the mirror position to keep the laser fringes locked at the same phase, so that the interferometer path difference in principle remains constant under any environmental changes or instrument drifts.

The facility also exists for the instrument operator to dial the phase of the fringes manually from the PZT control software, which applies piston motion to the mirror and changes the delay by tiny amounts, either for calibration purposes, or so that the original delay can be recovered in the event that it is lost. Utilising the fact that there is a small wavelength dependency in the phase change as a function of delay (equation 2–11), it is possible to compare 'before' and 'after' measurements of pure iodine spectra in the event of phase loss by using the data reduction software to establish the number of phase wraps that may have occurred, and thereby break the modulo 2π phase degeneracy. In this way exactly the same interferometer delay can be re-established, to within a very small fraction of a wavelength.

The same software can also be used to 'jitter' the mirror, scanning it rapidly back and forth over a range of some tens of fringes. Over any exposure longer than a few seconds, this washes out the fringes sufficiently to produce a standard non-fringing spectrum, and is used for creating various calibration frames for the pipeline (see section 4.1).

The mirror has three degrees of freedom: piston motion and two axes of tilt. Earlier versions of the software locked the fringes by monitoring the phase, period and angle of the laser fringes on the video camera, giving the three necessary constraints on the three degrees of freedom. (In fact, the earliest versions actually monitored only the phase itself.) Monitoring the fringe period was found not to provide sufficient locking capability along its respective mirror tilt axis, however, especially for the purposes of the multi-object instrument, since the width of the laser beam only samples a very small part of the mirror and does not give much of a 'lever arm' for making accurate corrections.

During upgrades to the Keck ET instrument, two laser beams were fed through the interferometer to sample the mirror surface at two widely separated positions, giving a much longer effective baseline over which to measure the mirror's tilt in the slit direction (i.e. in the direction with the tilt axis along the dispersion direction). Actively locking the phase over both of these effectively locks this tilt to much better accuracy. In the case of Keck ET, the enhanced locking is important as the slit direction is the direction along which all the spectra are stacked up: changes in the tilt in this direction lead to large changes in the delay d_0 for different spectra, and therefore changes in the respective phase-velocity scales Γ. Changes in the other tilt axis are less critical, and are still controlled by locking the apparent angle of the laser fringes on the video camera.

If necessary, three laser beams could in principle be used to sample the mirror at three of the corners, giving the largest possible baseline over both axes. Locking the phase of all three would lock the three degrees of freedom to the best accuracy, but in practice it is difficult to fit this many beams through the beamsplitter.

An alternative that is currently being pursued by Suvrath Mahadevan is to use a monolithic 'passive interferometer', where there is no active feedback. Instead the complete interferometer, including mirrors, is constructed as a single unit from materials whose thermal expansion properties are chosen such that they compensate each other for

temperature changes, so that the length of the delay should be relatively insensitive to temperature (Mahadevan 2006).

PMT. A photomultiplier tube (PMT) is mounted directly behind the VPH grating on the axis of the input beam from the collimator, and picks off some of the extraneous light that is lost to the first order. The PMT enables the operator to monitor the flux live during the exposure and ensure that the focus is optimal and that guiding, seeing conditions, and sky clarity are adequate.

The flux is also continuously recorded during the length of an exposure, triggered automatically by the opening and closing of the instrument shutter. This allows for flux weighting of the centre-time of the exposure to account for temporary cloud cover, loss of guiding, etc. (This has still to be fully implemented in the pipeline - see section 5.5.)

Operator interface. Full instrument control is provided remotely by two computers. The CCD computer controls image acquisition, running software provided by the manufacturers of the detector. This computer also runs the PZT recording software (adapted by Pengcheng Guo), and the electronic observing log software, (written by Craig Warner) which is written specifically for the ET instruments and stores the logs in a format suitable for automatic reading by the data reduction pipeline. A second computer is used to operate and monitor the phase locking system, with a graphical interface written in Labview (written by Curtis DeWitt and Pengcheng Guo).

3.2 The Single-Object ET at KPNO

The KPNO ET is situated in a small insulated room in the basement of the 2.1 m telescope, and can be fed either by the 2.1 m telescope itself or the 0.9 m coudé feed telescope on the roof of the same building. Fans and a thermostat-controlled space heater maintain the temperature at a constant $24.0 \pm 0.1\,^{\circ}\mathrm{C}$. An inner enclosure built around the instrument on the bench itself provides a second layer of thermal stabilisation. As a prototype instrument, the KPNO ET has been through a number of incarnations and upgrades, beginning with its first trial installation in late 2002. A complete description of

the current set up is found in Mahadevan (2006). A brief summary of the setup at various points is described here.

3.2.1 Test Run, 2002

The KPNO ET's earliest form was as a temporary set up during an engineering run in August 2002. The instrument was built largely from donated parts and off-the-shelf components at minimal cost, and installed over a period of around 5 days before observing began. Details were given in Ge et al. (2003b). This was the setup first used to confirm the known planet around 51 Peg (section 6.1).

A Michelson-type interferometer was employed, with one mirror tilted by a few wavelengths, and with a 4mm BK7 glass plate inserted in one arm, giving a total optical path difference between the arms of 7 mm. The interferometer output was fed into a spectrograph of Czerny Turner design with two parabolic mirrors and a first-order reflection grating, and an adjustable entrance slit. The spectrograph operated at $f/7.5$ and the final operating resolution was measured at $R = 4540$. Using a KPNO 1k×3k back-illuminated CCD, we obtained a wavelength coverage of 270 Å centred around 5445 Å. The image was spread over $\sim$ 300 pixels in the slit direction, giving a total of around 12 fringe periods. The $f/8$ telescope beam was fed into a 200 μm fibre, matching a 2.5″ stellar image. Due to focal ratio degradation, the output focal ratio of the fibre was $f/6$, which was converted to $f/7.5$ to feed the spectrograph. This étendue mismatch led to $\sim$ 40% photon loss at the slit.

3.2.2 Upgrades, 2004

During late 2003 and early 2004, a permanent and improved version of ET was installed at the 2.1m (van Eyken et al. 2004b), with the following improvements:

- The spectrograph was completely replaced with a custom designed and built collimator and camera, allowing for a faster focal ratio ($f/5$ collimator and $f/2$ camera) with less aberration. The grating was replaced with an optimised volume phase holographic (VPH) transmission grating with substantially higher throughput, operating at around R $\sim$ 5000 and designed to peak in transmission at 5300 Å (Mahadevan 2006).

- The second interferometer output beam was also utilised. A Michelson interferometer actually has two output beams: the standard output, as shown previously in the ET schematic (figure 2-1), and also a second which in a fully on-axis design is actually reflected back down the input path. After the new upgrades, both output beams from the interferometer were utilised by picking off the second beam where it is slightly displaced from the input beam owing to the tilted mirror in the interferometer. The second output beam was then fed along a path parallel to the primary beam, imaging a second identical (though phase-inverted) spectrum onto the detector. An approximately $1/\sqrt{2}$ improvement in precision is obtained by averaging the results from the two spectra, equivalent to doubling the throughput.

- A new cryotiger-cooled 4kx4k back-illuminated Fairchild CCD detector was installed, allowing ~ 600 Å wavelength coverage, centred around roughly 5510 Å.

- A third beam was also passed through the instrument for calibration sources, allowing for simultaneous parallel ThAr fiducial measurements.

- The phase locking software was enhanced to lock fringe period and angle, in addition to the phase which was all that was measured previously.

- The instrument was installed in its current full thermal enclosure.

3.2.3 The current KPNO ET

The current set up was used to detect ET-1 (section 6.2), and is detailed fully in Mahadevan (2006). To summarise, a 200 μm fibre (2.5″ on sky) feeds the $f/8$ beam from either the 2.1 m or the 0.9 m coudé feed telescope into the instrument at $f/6$. The spectrograph grating is a VPH Dickson-type grating, with 92% efficiency over a 600 Å bandwidth, designed to create a spectrum with resolution $R \sim 5100$ on the detector with a coverage of 5000–5640 Å. The spectrum has a 6.7 pixel resolution element (FWHM) and covers 58 pixels in the slit direction for typically ∼4–5 fringe periods. The image is now recorded on a cryotiger-cooled back-illuminated 4k×4k CCD detector manufactured by Spectral Instruments Inc., with 15 μm pixels, 90% quantum efficiency, and linearity better than 1% up to 80% of full well capacity (where full well is at ∼99,000 e^-). The continuous-cycle cryo-cooling provided by the cryotiger allowed for a significant improvement in image stability. (Previously nitrogen cooled, the detector's position would drift by several pixels in the slit direction as the nitrogen evaporated and the dewar's

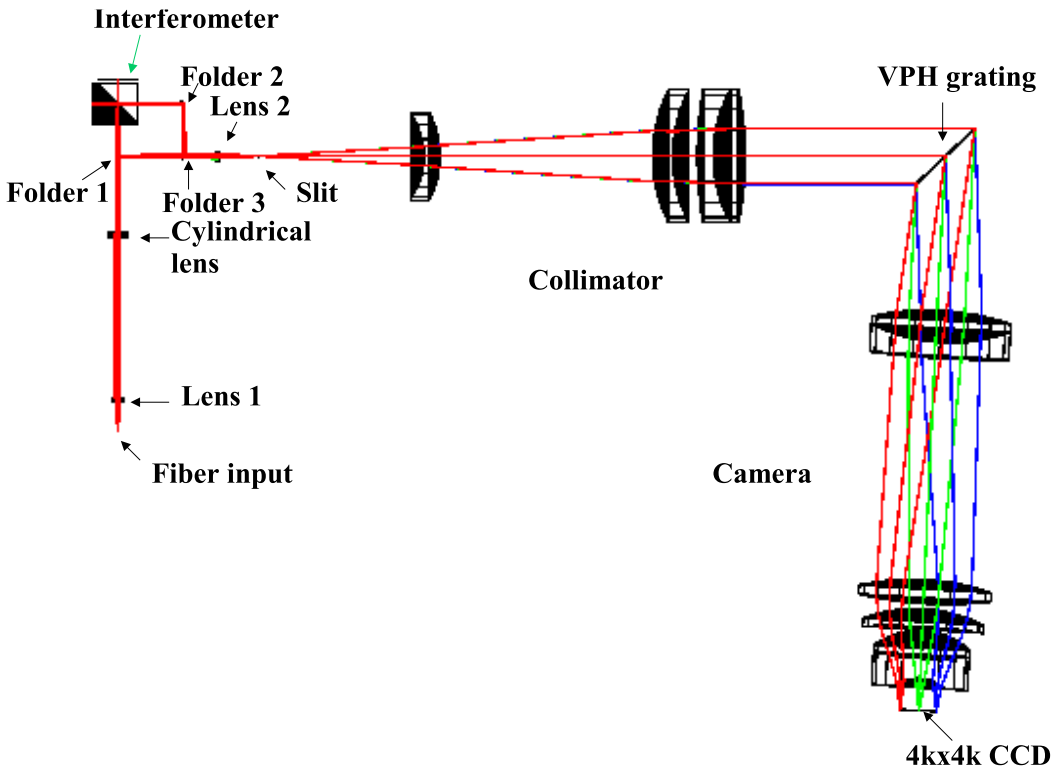

Figure 3-1. Optical design of ET at the KPNO 2.1m telescope, by Jian Ge and Deqing Ren. Designed to operate at resolution $R = 5100$, with a wavelength coverage of 640Å.

moment changed; manual refilling would again nudge it by several pixels.) A Melles-Griot HeNe laser at 0.6238 μm is used for the interferometer phase locking system. The glass iodine cell is cylindrical with a length of 150 mm and diameter of 50 mm, and is stabilised to a temperature of $60 \pm 0.1\,^{\circ}$C. The main specifications are summarised and compared with the multi-object Keck ET in table 3-1, and the layout of the optical design is shown in figure 3-1.

3.3 The Multi-Object Keck ET at APO

The Keck ET has also been through several engineering revisions. The original prototype was temporarily installed in the telescope pier in the basement of the Sloan Digital Sky Survey 2.5 m telescope in March and April 2005, designed for 20 simultaneous-object capability.

The current version, built in August 2005–February 2006, was fed by much longer fibres so that it could be installed in a separate nearby building where the Sloan plate-cartridges are plugged with fibres and stored, allowing for much better thermal control and isolation from telescope vibration. The design takes 60 200 μm target-fibres from the telescope to the instrument input (Wan et al. 2006). At the telescope end the fibres are plugged into a coupler in the plate cartridge, which contains a metal plate that fits at the focal plane of the telescope. Short 180 μm fibres inside the cartridge run from the coupler to holes drilled in the plate at positions corresponding to those of the target stars in the field that the particular plate is drilled for. The diameter of these short fibres corresponds to $3''$ on the sky, feeding at $f/5$. At the other end, the long fibres are fed to the instrument at $f/4$ in three vertical groups of 20 fibres, running alongside and slightly offset relative to each other. After passing through the interferometer these are imaged onto three slits in the instrument, so that the spectra cycle through the three slits in sequence as as we step spectrum by spectrum down the detector (with some exceptions in the ordering towards the first and last spectra for practical reasons.)

The detector is again a cryotiger cooled 4k×4k Spectral Instruments CCD, with 15 μm pixels, 92% quantum efficiency, and linearity better than 1% up to 80% of full well capacity (where full well occurs at $\sim$94,000 e^-). The VPH grating in the spectrograph (non-Dickson type) (Zhao & Ge 2006) operates at $R \sim 5100$, with spectra covering a larger range than the KPNO ET at 5000–5900 Å. 59 of the 60 spectra fall on the detector in the current alignment (the last does not quite fit), and each covers about 4096×50 pixels, again with typically on the order of $\sim$4–5 fringes in the slit direction. This information is also summarised in table 3-1. The optical layout for the interferometer is shown in figure 3-2, and that for the spectrograph in figure 3-3. A full detailed instrument description can be found in Ge et al. (2006b); Wan et al. (2006); Zhao & Ge (2006).

Table 3-1. Keck ET instrument specifications

	KPNO ET	APO Keck ET
# targets	1 + calibration	59
# interferometer outputs	2	1
Fibre diameter	200 μm / 2.5″	180 μm / 3″
Telescope $f/\#$	$f/8$	$f/5$
Instrument input $f/\#$	$f/6$	$f/4$
Resolution R	5100	5100
Resolution element	6.7 pix	5.0 pix
Detector	S.I. 4k×4k	S.I. 4k×4k
Spectral coverage	5000–5640 Å	5000–5900 Å
Size of spectrum	4096×58 pix	4096×50 pix

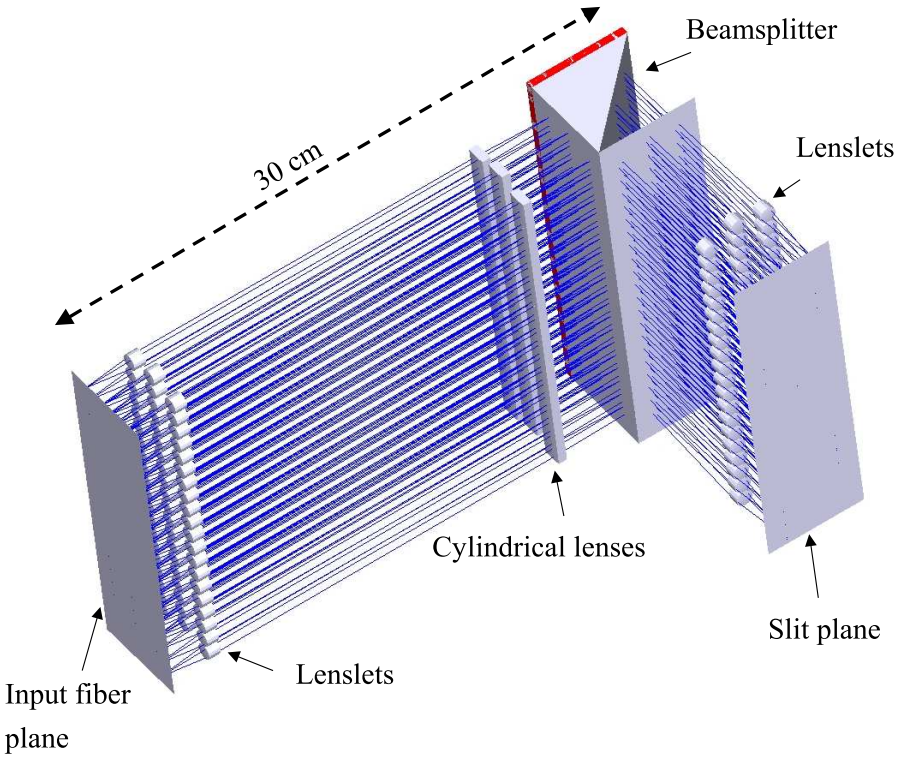

Figure 3-2. Optical design of the Keck ET interferometer, by Bo Zhao and Jian Ge, showing all sixty beams.

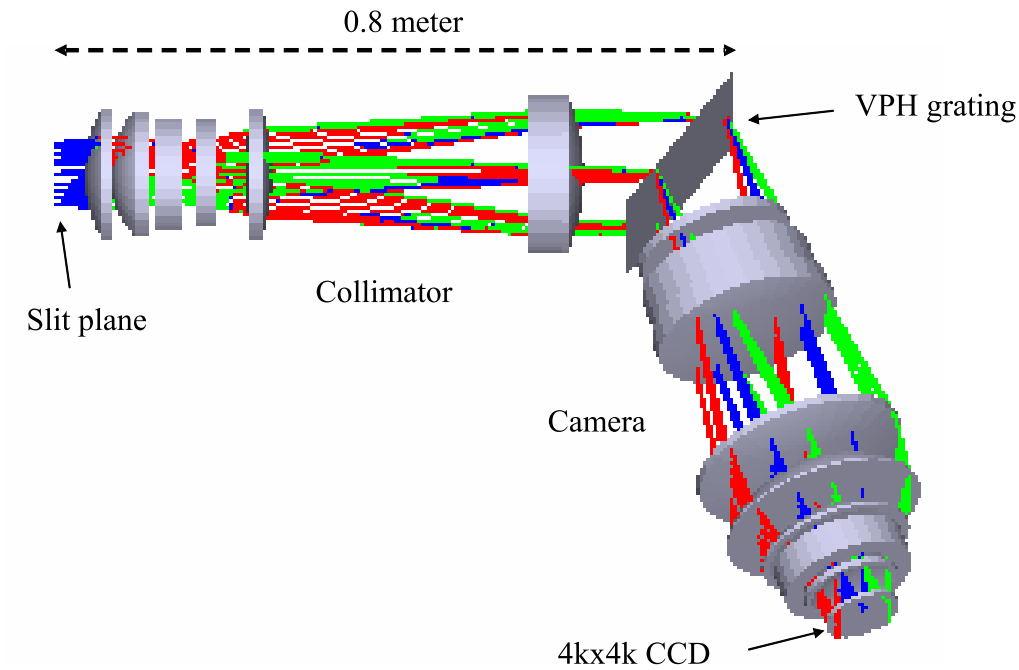

Figure 3-3. Optical design of the spectrograph for the Keck ET, which follows immediately after the interferometer shown in figure 3-2, matching at the slit plane. (By Jian Ge.)

CHAPTER 4
THE DATA REDUCTION PIPELINE

The data reduction software is a key element in the ET project. Real instrumental data are never perfect, and obtaining the high precisions called for in the ET project requires robust software that is able to account for outliers, systematic errors, and unexpected data problems, yet still extract the maximum possible information content from the data. The large quantities of data produced, especially by the multi-object instrument, render continual human interaction and oversight impossible: the software must be very reliable, and the process automated as far as possible. Furthermore, since the number of users is expected to increase, the software needs a streamlined and user-friendly front end that protects the user from unnecessary complication and is robust against any unanticipated input from a less experienced operator.

We have now almost completed the process of consolidating the data reduction procedure into a single stand-alone piece of software which will automate the entire process, taking raw data and electronic observation log files through to final radial velocity plots. As it currently stands, the software is able to handle full multi-object data and produce radial velocity measurements in a reasonably streamlined fashion, although some work remains to be done.

The pipeline and the specific algorithms used have been developed over time alongside the development of the ET instruments themselves, but the essential components have remained largely the same. Here we describe the most current version of the pipeline that takes us from raw data to RV measurements. In the future it will also be necessary to automate the process of analysing the RV data and identifying candidates for further follow-up, and to provide preliminary fits to the measured RV to estimate orbital parameters. (Preliminary work on such RV analysis has been undertaken by Stephen Kane.)

4.1 Observing Procedure and Calibrations

Over time, a standard observing procedure has been developed to ensure that the full set of necessary calibrations are taken. Calibration data taken as follows.

With the 'jitter' mode of interferometer mirror PZT switched on, washing out fringes as if there were no interferometer (section 3.1), we take:

- **Tungsten-illuminated iodine, no fringes:** several taken at the beginning and end of an observing run. Provides a low resolution iodine spectrum for reference, and in some circumstances these can, if necessary, be used as flatfields. (Although dividing out an iodine spectrum will remove the intensity variations due to the iodine absorption lines, it will not actually change fringe visibility or phase, and it has the advantage that the illumination function is likely to better match that of data with the iodine cell inserted, giving a better head start on illumination correction. The downside is that any hint of residual fringing in this calibration will introduce systematic phase errors in the data.)

- **Pure tungsten, no fringes:** several taken at the beginning and end of each observing run, providing a simple continuum illumination. An average over these frames is used as a master flatfield image.

- **Thorium-Argon, no fringes:** Several taken at the beginning and end of each run. Used to create an averaged template for slant correction of spectral lines to match the CCD axes.

- **Thorium-Argon with iodine, no fringes:** These can be used for slant correction of spectra where the iodine cell is switched in if it is suspected that insertion of the iodine cell is affecting the slant of the slit image in the spectra.

These calibrations are taken at the beginning and end of each run. Once jitter is switched off, the interferometer phase locker is restarted, and the phase restored to the same value, modulo 2π. Before and after jittering, a single tungsten-illuminated iodine frame is always taken. An on-the-fly rough reduction of these frames is performed, and the fringe phases differenced between the two frames across the length of the spectra: any phase wraps that may have occurred are easily identified because of the inverse dependency of phase shift on wavelength when the interferometer delay changes (equation 2–11). If any have occurred, the phase is stepped back to its original value using the phase-locker software, and another iodine frame is taken to verify that a near-zero radian

phase shift has been established across all wavelengths. In this way, the interferometer delay can in principle be restored to precisely the same value at any time, within the precision of the phase locking software and provided that the PZT has enough dynamic range.

After the beginning-of-run jitter calibration is complete, the following data are taken with the normal fringes:

- **Pure tungsten:** taken once or twice. Used to assess the presence (or otherwise) of interferometer comb, and any possible comb-aliasing. If the un-aliased comb is clear enough, it can be used for fringe counting to estimate the size of the interferometer delay (section 2.5).

- **Tungsten-illuminated iodine:** several taken at the beginning, middle and end of the night, used for determining a calibration of the fringe period in the slit direction as a function of wavelength, averaged over the night. This calibration is then applied to all fringe fitting for the whole night's data. An additional iodine frame is taken consecutively with any pure star template spectrum to provide a corresponding reference template spectrum. These frames can also be used periodically to check that the phase locker has not skipped any laser fringes, following the procedure outlined above for phase-wrap restoration.

- **Pure star templates:** taken at least once for each target or target field to be observed, to provide the template for separating out star and reference spectrum information in the combined star/reference data.

- **Star with iodine:** the actual data frames used to make RV measurements.

- **Bias frames:** several zero-length exposures taken each night and averaged for standard detector bias subtraction.

- **Dark frames:** usually several taken each night, matching the length of the longest exposure of the night, with no light input into the instrument. Averaged together over the whole run and used for standard subtraction of any stray background light sources that may be present in the instrument.

4.2 User Front-End

The pipeline is designed to be a single stand-alone piece of software which can handle both single- and multi-object data from our ET instruments, and can easily accommodate changes in instrument configuration. Pipeline options, such as which processing steps

to perform, whether to expect single- or multi-object data, whether to expect single or dual interferometer output beams, and the various options for each processing step, are controlled by means of a single master parameter file, which can be edited either with a simple text editor or by using a graphical user interface (GUI) written by Craig Warner (figure 4-1). Several subsidiary simple-text parameter files contain the remaining information needed for the pipeline configuration. The pipeline is written in the IDL data analysis language by Research Systems Inc., and is designed to be launched either from the IDL command line, or from the GUI. The software is designed so that different default parameter files can be chosen for different instruments and configurations, whilst easily allowing for fine tuning for particular unanticipated data situations. For every particular data reduction run, it should be possible to save the full set of parameters as a record of the particular processing that was done. In addition, records of all processing performed on an image are stored in the corresponding header information for that image.

The early parts of the data reduction procedure that followed standard astronomical data reduction methods were originally performed using the IRAF[1] (Tody 1986, 1993) software package. To enable incorporation of the entire procedure into a single pipeline, the initial IRAF parts of the data processing were later replaced with IDL code written by Craig Warner, along with the parameter-editor/launch GUI. At this time, the GUI only incorporates those parameters which affect the parts of the data reduction procedure that were originally performed using IRAF, and the subsequent data reduction must be launched independently from the IDL command line. Adding in the requisite parameter options to the GUI and incorporating launch of the full pipeline from there are relatively simple tasks, however, which are expected to be completed soon.

The raw data are stored as FITS images (Hanisch et al. 2001). Automation of the pipeline is facilitated in part by reading in what information is available from the FITS

[1] 'Image Reduction and Analysis Facility' – see also http://iraf.noao.edu/

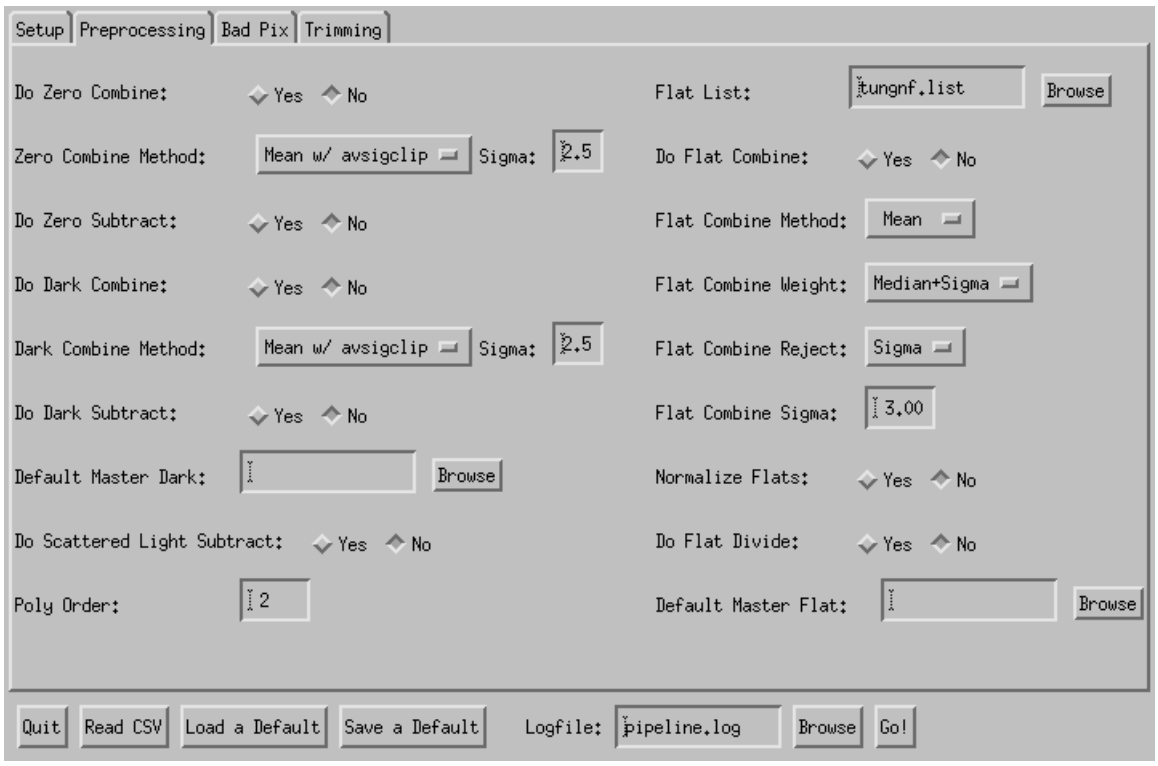

Figure 4-1. Example screenshot from the ET pipeline graphical user interface, written by Craig Warner, showing options for some of the preliminary pre-processing steps. The interface is used to edit the parameter file for the pipeline.

header information (epoch of observation, exposure length, etc.), and in part by reading in information from electronic log files created by the observer, which contain information on the image type (i.e., target or calibration frame and what kind of calibration, whether the iodine cell is switched in, whether the interferometer PZT is being jittered to wash out fringes, etc.), and the name of the target or target field. The electronic logs are created during observation using graphical software written in the Java language (Gosling et al. 1996)[2] by Craig Warner specifically designed for the ET instruments (figure 4-2), based on an original Microsoft Excel spreadsheet template by Andrew Vanden Heuvel. It is anticipated that in the future, more information will automatically be included in the FITS headers during observation, including telescope pointing information, observatory

[2] See also http://java.sun.com

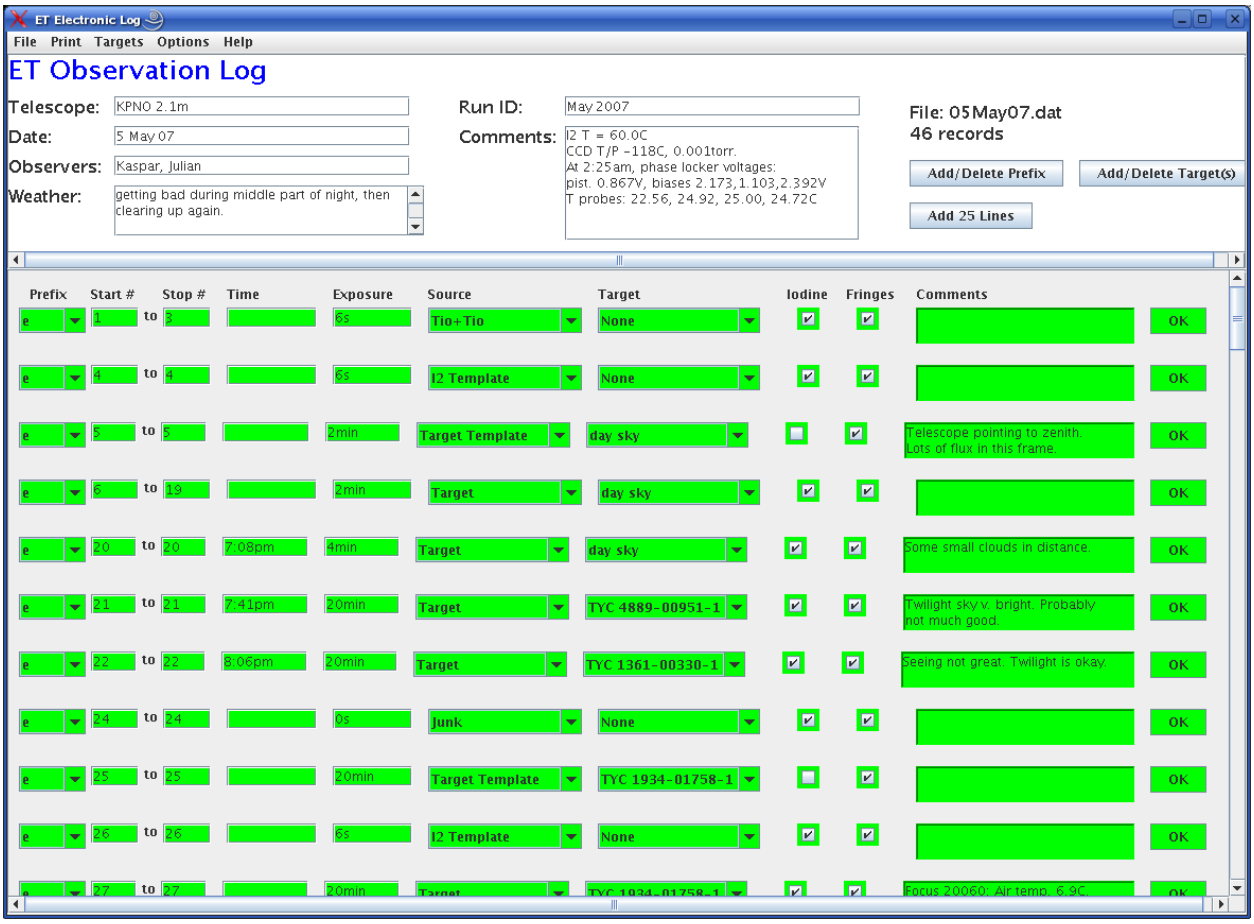

Figure 4-2. Example screenshot from the ET electronic observing log, written by Craig Warner. The log produces data files that are automatically readable by the data reduction pipeline.

information, instrument status (iodine cell in/out etc.), so that there can be less room for human error between sky and RV results.

4.3 Fibre Mapping

For the multi-object Keck ET, correctly identifying the right spectrum on the detector with the right target is of crucial importance. Each field to be observed with the Sloan telescope is associated with a metal plate which has holes drilled in it to match the positions of targets to be observed in the field. The plate is installed in a cartridge, and short colour-coded fibres running from a standard connector in the side of the cartridge are plugged into the various holes in the plate. When a field is to be observed, the

respective cartridge is installed in the base of the telescope so that the plug-plate is positioned at the focal plane.

During field selection and fibre plug-plate design, led by Suvrath Mahadevan and Roger Cohen, each hole on the plate is assigned an index number, and a data file created which contains the index number and information for each target in the field, including identifier, exact coordinates and visual magnitude. This file is used for the drilling of the plates. As the fibres are plugged, a careful record is kept of which fibre colour-code corresponds to which hole number, from which a standardised machine-readable text file is created for each plate plugged.

The pipeline takes all these files, along with a colour-code table which maps colour-codes to fibre index numbers, and a fibre-spectrum table which maps fibre index numbers to the corresponding spectrum index number on the detector, and collates all the information into a single 'master map' for each plate. In this way, every spectrum for a given plate observation can be associated with the correct target name, visual magnitude and celestial coordinates. Since the plug-plate target coordinates necessarily already include precession and proper motion effects for the epoch of observation, these coordinates can be directly used by the pipeline for Solar System barycentric correction (section 4.4.4).

In the electronic logs for the multi-object instrument, the 'target name' recorded is the plate identifier for the field. Using the master maps, the identifier is expanded out into a list of stars with their associated spectrum index numbers, and then the data can be processed by the standard single-object pipeline.

4.4 Data Reduction

The data reduction procedure itself can be divided into three main parts: image pre-processing, which prepares the raw spectra for extracting the RV information; fringe fitting, where fringe phase and visibility are measured along each column (or wavelength channel) on the detector, creating a file containing the 'whirl' information for each

spectrum; and RV measurement, where star template, reference template and data frames are combined to extract a final differential radial velocity.

4.4.1 Preprocessing

4.4.1.1 Bias/dark subtraction and flatfielding

Bias and dark-subtraction and flatfielding were originally performed in using standard IRAF routines (Tody 1986, 1993) in early versions of the data reduction procedure. These routines were later replaced with Warner's IDL code, and are performed as for standard astronomical data processing. A master bias frame is created by averaging (by mean or median) all the bias frames taken over a run, with outlier rejection, and subtracted from all data frames. A similarly averaged bias-subtracted master dark is also created for an entire run, using outlier rejection across the frames for each pixel to reliably remove cosmic rays. The master dark is scaled by appropriate exposure times and also then subtracted from all data frames.

The initial flatfielding is also performed following standard astronomical image processing procedures. A master flatfield is created from averaged jittered pure-tungsten continuum calibration frames without fringes, again with outlier rejection to remove cosmic rays. The average (either a mean or a median) can be weighted by the exposure time or by various measures of the average flux in the flatfield frames (mean, median, mean with outlier clipping, etc.), to account for fluctuation in the brightness of the tungsten lamp.

The flatfielding takes care of pixel-pixel gain variations in the CCD chip which can otherwise appear as noise similar to photon shot noise. It also usually goes a long way – although not all the way – towards flattening the illumination profile of the spectra.

4.4.1.2 Bad pixel masks and cosmic ray rejection

Options exist in Craig Warner's IRAF replacement software to apply bad-pixel masks, for replacing known bad pixels, columns or rows on the CCD chip; and to apply a cosmic ray rejection algorithm to all data frames. In practice we have had difficulty obtaining the

necessary calibration frames to create bad-pixel masks for the ET instruments, but in fact there appear to be very few bad pixels in any case in either of the CCDs. The cosmic ray rejection is not yet fully tested, and currently runs at a very large processor time overhead. Since there are various levels of outlier rejection in the code later on, cosmic ray rejection is usually left switched off, though this will likely be revisited.

4.4.1.3 Trimming

Trimming out the individual spectra is the final step that was originally performed in IRAF, and has now been replaced with Warner's IDL code. A simple-text parameter file lists the trim margins for each spectrum in a full CCD image, and is used to chop up the image into a single image file for each spectrum. The parameter file also serves to indicate to the pipeline the number of spectra to expect, whether many, as in the case of the multi-object instrument, or just three, as in the case of the single-object instrument (two output interferometer beams, plus an optional parallel calibration beam).

The data reduction steps to this point all represent more-or-less standard spectrographic data reduction. From here, the steps become more specialised to the DFDI technique and the ET instruments.

4.4.1.4 Slant correction

Owing to the fast focal-ratio of the ET instruments, there is unavoidably some asymmetry in the projection of the spectra onto the detector (see figure 4-3). This results in both slant and some degree of curvature of the slit image (referred to collectively as 'slant' for convenience), and hence of the spectral lines. In order to allow for the asymmetry, we apply an algorithm to correct the effect and align the spectral lines exactly with the CCD columns, so that we can easily fit sinusoids to each column during the fringe fitting later on.

The slant is measured using the non-fringing ThAr calibration frames, which provide sharp emission lines that help to provide an accurate determination. The spectrum is chopped into typically 25 segments in the dispersion direction. For each

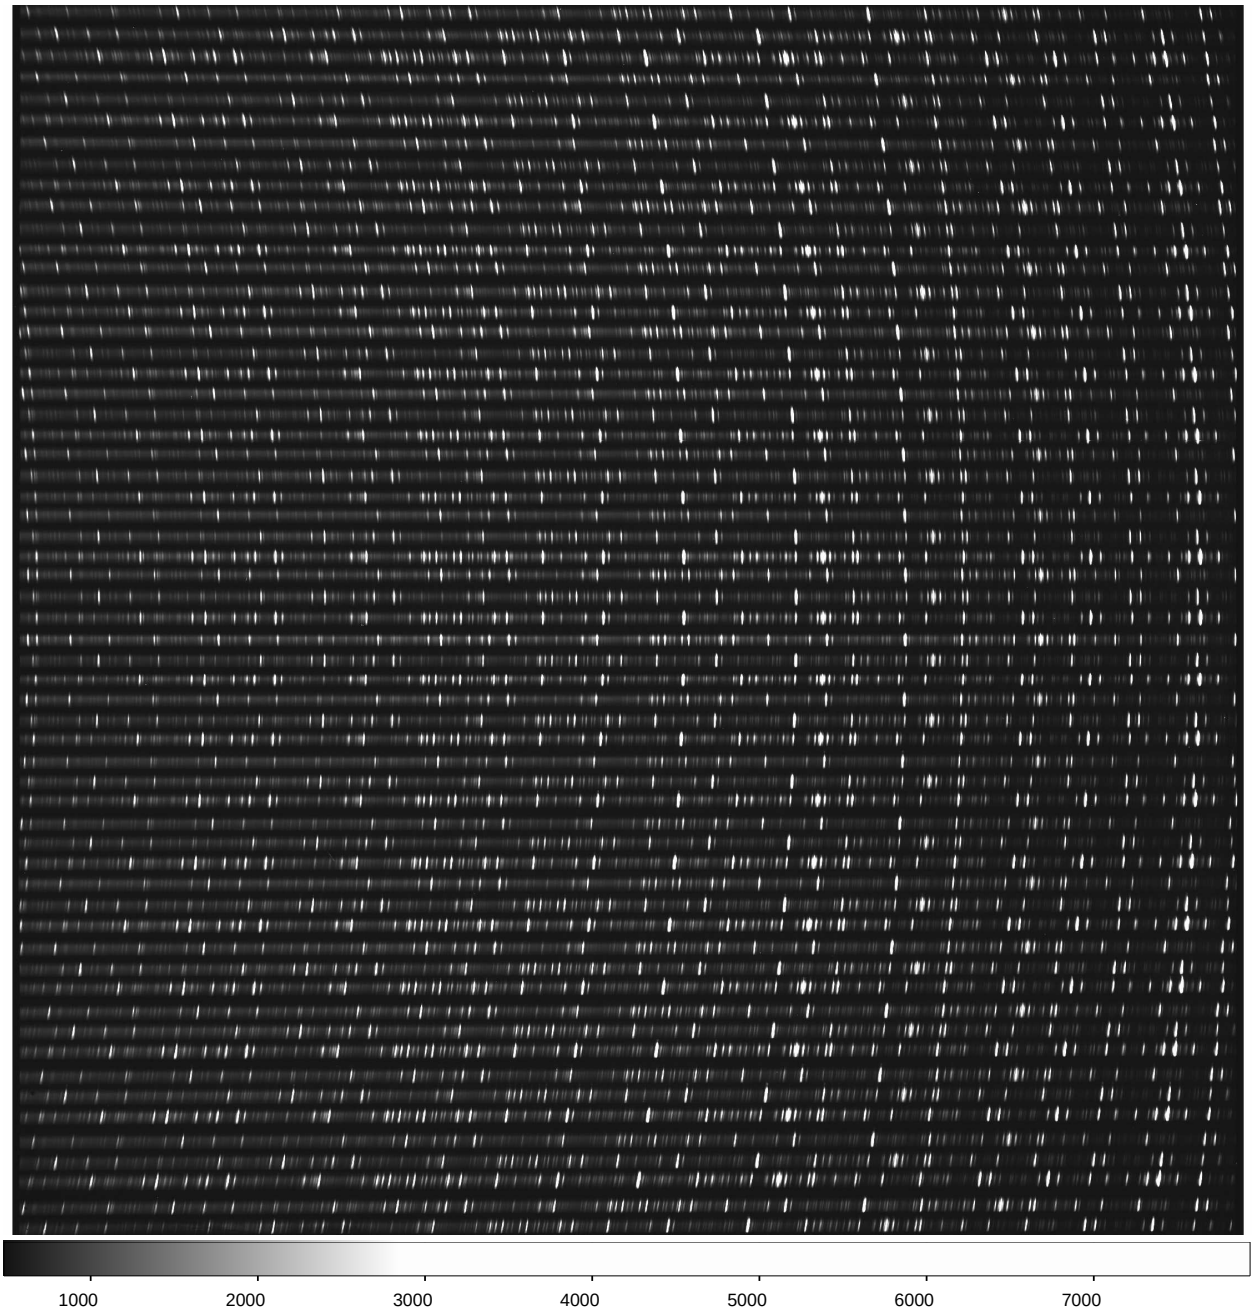

Figure 4-3. Example non-fringing ThAr frame from the multi-object Keck ET, showing all 59 spectra. Distortion of the field that leads to slant in the slit image is clear. Wavelength scales are offset in cycles of three spectra due to the arrangement of the fibre images on the spectrograph slits. Grey scale indicates ADU per pixel.

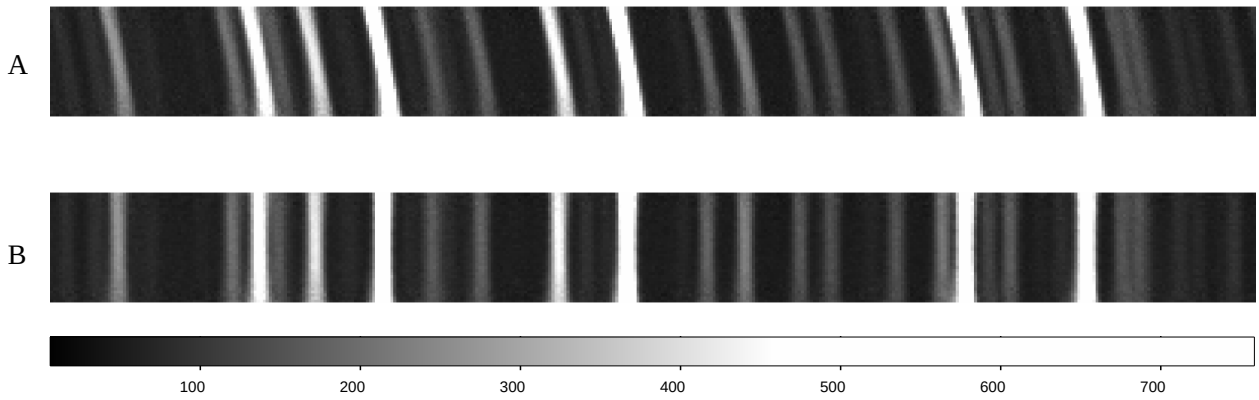

Figure 4-4. Slant correction of a non-fringing ThAr spectrum. A) Before correction. B) After correction. Image is from a portion of the fifth spectrum on the detector, taken from the same image as figure 4-3. Grey scale indicates ADU per pixel.

segment, the relative offset between each pixel row and the central row is determined by cross-correlation after applying a low-pass filter (Verschueren & David 1999). The cross correlation function is evaluated at each integer lag value, and an implementation of the interpolation scheme proposed by David & Verschueren (1995) used to measure the peak position to sub-pixel accuracy. This gives us a measurement of the mean slant for each segment.

Before applying any corrections, the measured slants for each segment are median filtered and then smoothed by boxcar averaging over the 25 segments, to reduce noise and to remove any outliers that may be caused by bad segments (e.g. any with very low flux, few or no spectral lines, or stray cosmic rays). Corrections are then applied to all the data spectra by interpolating the measured slants from segment to segment for each pixel, and then remapping each pixel of the spectrum to its new position, also by interpolation.

Tests of the internal precision of the algorithm are made by measuring the slant using the ThAr calibration frames, applying that correction to the calibration frames themselves, and then re-measuring the slant on the corrected calibration frames. Doing this gives typical RMS deviations from zero along the slit typically on the order of 0.01 pixels. An example of the effects of the slant correction can be seen in figure 4-4.

4.4.1.5 Illumination correction

Good correction of the illumination function in the data to create flat sinusoidal fringes has proven to be essential to avoid significant systematic errors in the fringe fitting. We have found that the illumination can vary significantly not only according to illumination source (i.e. tungsten lamp, ThAr lamp, or star) and presence of the iodine cell, but also with time (on the time-scale of $\sim 1\,\mathrm{hr}$), perhaps due to changes resulting from position on the sky, small changes in alignment of the instrument, etc. Thus, simply taking flatfields to account for the illumination, even for all the different combinations of source and iodine cell, is not sufficient unless it is done very regularly. Doing so is generally not practical, since the instrument needs to be interfered with[3] as little as possible during an observation run, and jittering the PZT mounted interferometer mirror to wash out the fringes is kept to a minimum.

The problem is overcome by taking advantage of the fact that the illumination varies only slowly in the dispersion direction, on a scale much larger (tens of pixels) than the period of the comb (typically $\lesssim$ 2–4 pixels). Since the comb is periodic, it is possible to boxcar average in the dispersion direction to obtain a reasonable approximation to the illumination function. The 'self illumination correction' algorithm first goes through and renormalises each wavelength channel in the slit direction, to largely eliminate absorption line features. It then repeatedly boxcar averages in the dispersion direction with an averaging window of $\sim$ 50 pixels until a smooth illumination-function image is obtained. The original image is divided by this function, and a surprisingly well flatfielded image is usually obtained.

Figure 4-5 shows the effect of this flatfielding. A raw data file and the flattened and pixel-pixel corrected version are shown along with the extracted illumination function, and

[3] No pun intended.

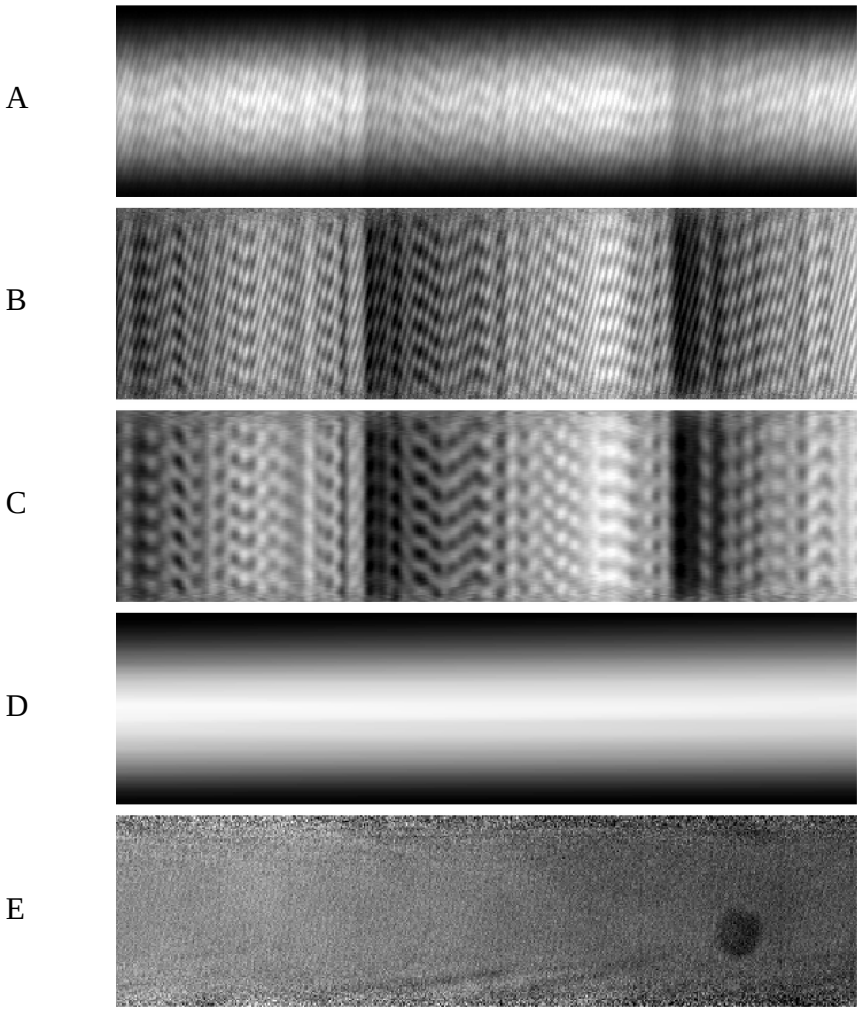

Figure 4-5. Pre-processing steps. A) Raw iodine+bromine fringing spectrum from the early prototype single-object ET. Interferometer comb is clearly resolved. B) Image after pixel-pixel flatfielding and illumination correction, and C) after low pass filtering to remove the interferometer comb. D) Illumination function extracted by self-illumination correction algorithm. E) Pixel-pixel gain variation seen by self-illumination correction of a flatfield frame. Note that the spectral lines are close-to vertical in the raw data, so slant-correction is difficult to see.

the pixel-pixel flat obtained by removing the illumination function from one of the flat files, where imperfections in the CCD response can clearly be seen.

4.4.1.6 Filtering

As seen in section 2.6.1, it is important that there is no interferometer comb presence in the data. This is ensured by a simple one dimensional low-pass spatial Fourier filter of

all the data in the dispersion direction, at a cutoff frequency just below that of the comb. The filter used is the IDL language 'digital filter' function, which generates a smoothly shaped filter to avoid the ringing effects that would be seen with a simple step-shaped filter. Edge effects at the ends of the spectrum caused by the filtering mean that the two ends of the spectrum must be trimmed where the data becomes meaningless (typically about 10 pixels at each end).

The effect of the filtering is shown in figure 4-5C. After filtering, the image preprocessing is complete and the data are ready for fringe measurement.

4.4.2 Fringe Fitting

Fringe fitting entails fitting sinusoids to measure the phase and visibility of each wavelength channel for a spectrum: in other words, measuring the complete whirl. This is achieved by simply fitting a sine wave to each vertical column in the image, using a standard χ^2 minimisation algorithm (IDL's 'curvefit' routine). The fit is weighted according to the number of counts in the original dark- and bias-subtracted non-flatfielded data, on the assumption of photon counting statistics combined with the read-noise to prevent over-weighting of zero- or very low-flux pixels. For the algorithm to converge it is necessary to first obtain good initial estimates of each of the four parameters of the sine wave: phase, spatial frequency, amplitude and offset. The amplitude can be estimated using the standard deviation across the channel, given that the amplitude of a sine wave is $\sqrt{2}$ times its standard deviation; the offset can be estimated from the mean values across the channel. The frequency is estimated by finding the peak frequency in the Fourier transform of the fringe, and the phase by finding the position of one of the peaks and relating it to the estimated fringe frequency (this latter is less important, but is used to speed up the convergence). The visibility is then given by the ratio of the amplitude to the offset.

The algorithm may sometimes find a local minimum in χ^2, rather than the true minimum. A common example is to obtain a very low frequency, high amplitude fit

which is an almost straight line through the centre of the data points. The fit is therefore iterated with different initial guesses for the phase until either parameter values are returned within sensible bounds and with an acceptable value of χ^2, or the fit is deemed to be impossible. In general local minima usually represent obviously wrong fits, so rejection does not normally pose a problem. An iterative sigma-clipping algorithm is used to ignore any outlying pixels in the fit, so that it is not affected unduly by cosmic rays or bad pixels.

An accurate determination of the fringe frequency is essential to determine the phase accurately. If the frequency is slightly wrong, then the phase measurement quickly becomes meaningless. Fortunately there is some redundancy in the data which helps: the fringe frequency varies only slowly with wavelength. For the fringing iodine frames taken at the beginning, middle and end of each night of observation, two passes of the curve-fitting are performed. The first determines the frequencies of each channel. A polynomial is then fit to the fringe frequencies as a function of wavelength, weighted according to the returned standard curve-fit frequency errors, and a second fringe-fitting pass is performed with the frequencies now fixed to match the polynomial. (In principle the fringe frequency should vary as the reciprocal of the wavelength. Distortion in the image means that the wavelength scale on the detector is not linear, however, and distortion in the slit direction further complicates things. Since the wavelength scale does not otherwise need to be calibrated, the polynomial fit leads to better results.)

The iodine frames provide the best data for determining the fringe frequency function, since the spectra have good flux and the fringe coverage is more or less continuous over the entire spectrum. Once all the iodine frames for the night are measured, the fringe frequency functions are averaged together (with outlier rejection in case of any bad frames), and the resulting function is used as a master frequency function for fitting all the other data from the night. This elimination of one of the free parameters in the sinusoidal fits yields a significant reduction in the fitting errors, and is crucial for obtaining photon

limited performance (section 5.2, figure 5-1). The results from the fringe fitting are saved as 'whirl' files for each spectrum.

4.4.3 Radial Velocity Measurement

4.4.3.1 Reference extraction

In the simplified case of a perfectly stable instrument and data with no overlaid reference spectrum, calculating the velocity shift is relatively simple: the phase shift in a given wavelength channel is directly proportional to the differential RV shift of the target, scaled by the phase-velocity factor Γ (section 2.4). For small shifts, it is adequate to perform a weighted mean of the phase shifts over all the channels to obtain a high precision Doppler shift that uses the information across the whole spectrum. For larger shifts, the phase shift can no longer be assumed to be independent of wavelength, λ, and a 'twist' is introduced into the phase shift as a function of wavelength. Because the wavelength scale is not linear on the detector, the functional form of this shift (as a function of detector coordinate) is not simple, but in principle could be established by deliberately stepping the interferometer cavity, and empirically fitting a polynomial function to the twisted function $\Delta\phi(\lambda)$. Once established,the function could be scaled and fit to any $\Delta\phi(\lambda)$: at a chosen reference point in the spectrum (usually the centre), the phase shift of the function would contain the information of an averaged value, and could easily be translated to a high precision velocity shift. This is the approach that would be used for a parallel simultaneous reference spectrum (iodine or ThAr), or for bracketing pure-star data exposures with pure reference frames: any apparent drift in the reference would simply be subtracted, leaving us with the true stellar Doppler shift.

In practice, however, a superposed iodine reference spectrum has always been used to account for instrument drift, due to various instrument instabilities that render other approaches impractical. In order to extract the stellar and iodine whirls from the combined whirl, $\mathbf{M}_j$, two templates are taken at each observation, giving a pure iodine whirl $\mathbf{I}_j$ (taken with a tungsten lamp as a light source), and a pure star whirl, $\mathbf{S}_j$. In

general, $\mathbf{M}_j$ will be some linear vector combination of the two template whirls, where the template whirls are rotated in phase and scaled in visibility by some amount to be determined. The basic mathematics for extracting the stellar and iodine components of the combined data were outlined in section 2.6.1. Essentially, we are aiming to solve equation 2–26, which we restate as:

$$\mathbf{M}_j \approx K_\mathrm{s}\mathbf{S}_j + K_\mathrm{i}\mathbf{I}_j, \tag{4–1}$$

where we have added the subscript j as a reminder that the whirls $\mathbf{M}$, $\mathbf{S}$ and $\mathbf{I}$ are functions that in practice have discrete values at each integer value of j, where j represents the pixel column index in the spectrum. $\mathbf{S}$ and $\mathbf{I}$ must be allowed to rotate in phase and scale in visibility, giving us the free parameters that we need to solve for. We note here that for the current analysis, we assume that the visibility scaling constants K_s and K_i are constant across all channels j. They are allowed to float to match their best-fit value, but their proper channel-dependent form given in equation 2–27 has only recently been appreciated, and has not been incorporated into the pipeline yet.

We express a general vector for the jth wavelength channel of a whirl in polar coordinates using the notation $\mathbf{M}_j = [M_j^\gamma, M_j^\phi]$ with M_j^γ representing the vector length (i.e. the visibility) and M_j^ϕ the phase. Using similar notation for $\mathbf{S}_j$ and $\mathbf{I}_j$, we can then write:

$$\mathbf{M}_j = [M_j^\gamma, M_j^\phi] = [K_\mathrm{s}^\gamma S_j^\gamma, S_j^\phi + K_\mathrm{s}^\phi] + [K_\mathrm{i}^\gamma I_j^\gamma, I_j^\phi + K_\mathrm{i}^\phi], \tag{4–2}$$

where K_s^γ and K_i^γ are the stellar and iodine template scaling factors respectively. K_s^ϕ gives the phase shift of the star component relative to the star template, and K_i^ϕ likewise for the iodine. These are the four unknowns in the equation, and for notational convenience are thought of as the components of two vectors, $\mathbf{K}_\mathrm{s}$ and $\mathbf{K}_\mathrm{i}$ (although technically they do not necessarily transform as vectors under vector operations). These are approximately

constant over all channels j if the wavelength range is small and/or the phase shifts are small. The corrected phase change[4] is then given by $\Delta\phi = K_{\rm s}^{\phi} - K_{\rm i}^{\phi}$.

To solve for our four unknowns, we begin by linearising the equations to simplify numerical solution. Converting our whirls to Cartesian coordinates, we can write for the x and y components of $\mathbf{M}_j$:

$$M_j^x = K_{\rm s}^{\gamma} S_j^{\gamma} \cos(S_j^{\phi} + K_{\rm s}^{\phi}) + K_{\rm i}^{\gamma} I_j^{\gamma} \cos(I_j^{\phi} + K_{\rm i}^{\phi}), \tag{4–3}$$

$$M_j^y = K_{\rm s}^{\gamma} S_j^{\gamma} \sin(S_j^{\phi} + K_{\rm s}^{\phi}) + K_{\rm i}^{\gamma} I_j^{\gamma} \sin(I_j^{\phi} + K_{\rm i}^{\phi}). \tag{4–4}$$

Looking at the right hand side of the x component and using trigonometric identities, we can rewrite the first term:

$$K_{\rm s}^{\gamma} S_j^{\gamma} \cos(S_j^{\phi} + K_{\rm s}^{\phi}) = K_{\rm s}^{\gamma} S_j^{\gamma} (\cos S_j^{\phi} \cos K_{\rm s}^{\phi} - \sin S_j^{\phi} \sin K_{\rm s}^{\phi}) = K_{\rm s}^{x} S_j^{x} - K_{\rm s}^{y} S_j^{y}, \tag{4–5}$$

where we have broken down $\mathbf{K}_{\rm s}$ and $\mathbf{K}_{\rm i}$ into their Cartesian components. The second term can be rewritten the same way, and a similar treatment can be applied to M_j^y in equation (4–4). Hence we end up with two linear equations in four unknowns:

$$M_j^x = K_{\rm s}^{x} S_j^{x} - K_{\rm s}^{y} S_j^{y} + K_{\rm i}^{x} I_j^{x} - K_{\rm i}^{y} I_j^{y} \tag{4–6}$$

$$M_j^y = K_{\rm s}^{x} S_j^{y} + K_{\rm s}^{y} S_j^{x} + K_{\rm i}^{x} I_j^{y} + K_{\rm i}^{y} I_j^{x} \tag{4–7}$$

One channel is thus insufficient to determine $\mathbf{K}_{\rm s}$ and $\mathbf{K}_{\rm i}$ completely. However, by combining all the available wavelength channels j, we effectively have ~ 2000 equations across a 4096 pixel-wide detector. The system is then *overdetermined*: in an ideal world only four equations would be required and all other sets of four would yield the same result, but in the real world, all sets of four will give slightly different results. The system

[4] $\Delta\phi$ in fact includes an arbitrary additive constant, since the instrument drift between the times the iodine template and the star template are taken is unknown. Since we are only interested in differential RV measurements, this constant is not of any consequence.

of equations can still, however, be written as a matrix equation, and then solved using the technique of singular value decomposition (SVD) (Press et al. 1992, section 2.6). This yields what is essentially a 'best fit' result. Once $\mathbf{K}_\mathrm{s}$ and $\mathbf{K}_\mathrm{i}$ are determined, they can be converted back to their polar equivalents, and the values K_s^ϕ and K_i^ϕ extracted, hence yielding $\Delta\phi$ and finally the relative velocity shift via equation 2–11.

This new approach yields results numerically extremely close to the previous approach employed by Erskine & Ge (2000), but has the advantage that it makes no assumptions about the values of the scaling constants K_s^γ and K_i^γ beyond their wavelength independence. It also allows for weighting of the equations to account for different measurement accuracies in the individual channels (although in practice we generally assume uniform weighting – see section 5.1).

The SVD solution provides the innermost layer of nested iteration in the radial velocity determination. We have not yet accounted for shift of the spectral lines themselves in the dispersion direction. This can be caused both by natural Doppler shift (by a small but significant amount) and by image shift on the detector because of slight offset of the beam path when the iodine absorption cell is inserted and because of possible instrument instability. It is important to compare the phases and visibilities of exactly the same part of the spectrum from data frame to data frame for a meaningful shift comparison. To accomplish this, the template whirls are shifted by sub-pixel amounts using interpolation (either linear or spline, depending on speed and accuracy required). The shifts are allowed to float as a free parameters, and the χ^2 value for the residuals between the best-fit solution and the combined star/iodine data frame is minimised using the IDL 'AMOEBA' algorithm (based on the routine of the same name in Press et al. 1992).

Finally, where required, a third layer of iteration can be used to iteratively reject outlier channels in the best-fit solution, to eliminate the effects of any remaining cosmic

rays or bad channels; usually this last iteration only makes a very small difference, however.

4.4.3.2 Bulk shift compensation

In the pipeline, the data are treated as if phase shifts were independent of wavelength, on the assumption that the wavelength dependency will average out to give a meaningful mean phase shift. Although it is averaged out, this phase shift still needs to vary consistently with a constant proportionality to Doppler shift and interferometer delay change. Relying on this broad assumption alone, however, leads to clear systematic errors. One approach would be to chop the spectrum up into small segments over which wavelength dependency is negligible, and treat each segment independently. This has never met with much success though, probably because of the addition approximation error dominating when small sections of the spectrum are used.

Early experiments showed that purely following the wavelength-independent approximation could give non-zero results when the interferometer delay, d_0 was changed by small amounts (equivalent to a Doppler shift), even when using a stationary tungsten-illuminated bromine absorption cell as a simulated target, with iodine as a reference. It was found that this error could be dramatically improved by using what we have termed 'bulk-shift compensation'. This technique attempts to mitigate some of the effects of wavelength dependency by subtracting from the data frame phases a polynomial fit to any apparent twist in the phase residuals between the best fit solution and the combined data frame. This correction is also applied iteratively as another layer in the iterations towards the best fit RV solution in the iodine/star extraction algorithm.

The reason the subtraction can be done is because we are interested in only differential phase shifts. In the simple case of no wavelength dependency, one can imagine that we could subtract any uniform phase shift we wanted from the data frame whirl, since the same phase would be subtracted from both star and iodine components: the *difference* between the star and iodine shifts compared to the templates would remain the same.

Subtracting a polynomial does the same thing, but allows for wavelength dependency (again, if the wavelength scale were truly linear, we could simply subtract a scaled $1/\lambda$ relationship instead). Without this bulk-shift compensation technique we were unable to obtain useful results.

4.4.3.3 An improved approach

Bulk shift compensation should work well in the regime where both target and reference shift by similar amounts, so that the difference between the two shifts is small and a uniform phase shift can be subtracted off both components in the combined data frame. This will be the case when instrument drift dominates and target Doppler shifts are small. As we push to longer time baselines where the Earth's barycentric motion becomes large, and as instrument stability improves, this regime no longer holds, and it is reasonable to expect systematic errors to creep in.

A more robust approach is simply to perform a full non-linear χ^2 minimisation over the full parameter space, allowing extra parameters to account for wavelength dependency. A version of such a code has been written, using Craig Markwardt's IDL fitting routine, 'mpfitfun'.[5] Instead of linearising the equations as in section 4.4.3.1, we simply add up to second-order polynomials to both the star and iodine templates, with floating parameters. Two visibility scaling parameters are also allowed for, as before (the full channel dependency according to equation 2–27 has not yet been incorporated, though this can be easily done). Free parameters are also allowed to account for the dispersion direction shift as well: in this case, again up to a second-order polynomial wavelength-dependent shift can be allowed for, instead of simply assuming a uniform shift. A polynomial shift allows for both uniform image shifts caused by image shift on the detector, and wavelength-dependent Doppler shift. We now have a total of up to 14 free

[5] http://cow.physics.wisc.edu/~craigm/idl/idl.html, (last accessed Oct 2007).

parameters. Any of these parameters can be held fixed or restricted in range, and any pair can be constrained to have equal values.

The new code is in fact computationally considerably faster, and has the advantage of being much simpler and easier to read. Issues with convergence failure at this point have hampered implementation, however, although initial tests show that it often gives comparable or better results than the earlier code. More work is needed to provide improved initial guesses for the parameters, to improve the constraints on their values, and to perform a more reliable search of χ^2 space that does not get trapped in local minima.

One way of reducing the number of free parameters – which should help the convergence issues – is to fully calibrate the wavelength dependency of the phase shifts. This can be done simply by stepping the interferometer delay by small amounts and taking tungsten-iodine exposures, and empirically fitting the phase shifts between the frames. Once the function is known, only a single scaling parameter per template is needed for is needed to phase-shift the star and iodine templates, instead of the three parameters per template for a polynomial fit. Such data have been obtained, but we have not yet incorporated them into the pipeline.

4.4.4 Barycentric Correction

Once the intrinsic differential RV of a target has been measured, the final step is to correct for the Earth's motion by translating RV's to the Solar System barycentre frame of reference. Over the course of a year, the RV of a stable target can vary by as much as $\sim 60\,\mathrm{km\,s^{-1}}$ if the line of sight is close to the plane of the Earth's orbit, simply from the orbital velocity of the Earth. A secondary effect is the Earth's diurnal rotation, which can cause variations of up to $\sim 1\,\mathrm{km\,s^{-1}}$ for targets in the equatorial plane. Many other effects need to be considered to obtain corrections needed for precision at the $\sim 1\,\mathrm{m\,s^{-1}}$ level (see McCarthy 1995, for a detailed discussion).

Our barycentric correction code is based on a slightly modified routine originally written by Suvrath Mahadevan. If necessary, the target stellar coordinates are first

precessed to the epoch of observation using the 'precess' routine from the IDL Astronomy User's Library (hereafter IAUL) (Landsman 1993).[6] This is important because at the $1\,\mathrm{m\,s^{-1}}$ precision level, position on sky needs to be known to within a few arc-seconds (McCarthy 1995). In the case of the Keck ET instrument the precession is not needed since the coordinates are obtained directly from the plate drilling files, which already include precession and proper motion effects. For the single-object KPNO ET, coordinates are automatically obtained from the online SIMBAD database[7] by target name (proper motion is not yet accounted for, though this should be a relatively simple matter to correct).

The barycentric velocity due to the Earth's orbital motion is determined using the IAUL 'baryvel' routine, based on the FORTRAN program of Stumpff (1980), which takes into account Earth-moon motion and accounts for the perturbations of the Solar System planets, claiming a $1\,\mathrm{m\,s^{-1}}$ precision level. The coordinates and altitude of the observatory are converted from geodetic to geocentric coordinates, and used to calculate the observatory velocity due to the diurnal rotation of the Earth, following the calculations outlined in the documentation for the IRAF routine 'rvcorrect'. The orbital and diurnal velocities are summed to give the full barycentric correction.

Relativistic effects (gravitational and transverse Doppler shifts) are neglected, as is parallactic motion. Overall precisions are found to be good to approximately the $1\,\mathrm{m\,s^{-1}}$ level (see section 5.6).

[6] See also http://idlastro.gsfc.nasa.gov/ (last accessed Oct 2007).

[7] http://simbad.u-strasbg.fr/simbad/ (last accessed Nov 2007).

CHAPTER 5
SOURCES OF ERROR

Understanding sources of error in the measurements is a crucial part of developing a new instrument. Establishing and, where possible, quantifying the significance of each source is an important factor in guiding the process of instrument and data reduction pipeline development. Characterising the errors for the ET instruments is an ongoing process, and the errors have to be tackled in order of importance: it is difficult to tackle small sources of error when RV signals are drowned out by other larger ones. Some, such as photon shot noise, are relatively simple random errors that are relatively easy to quantify and handle statistically; others are more complex systematic or correlated errors, that can be harder to assess. An exhaustive error budget is beyond the scope of this work, but some of the salient error sources are discussed here.

5.1 Error Bars

The error bars presented in all the ET work are intended to represent only short-term, random, uncorrelated noise. This representation helps to give a baseline against which to assess the instrument and pipeline performance, which is of great importance during the development of a new instrument where the software and instrument configuration are in a continual state of flux. The calculation of errors due to photon shot noise was already discussed in 2.7. For the most part, we have settled on using the photon limiting noise for our error bars, calculated individually for each RV data point. Short-term (few-hour) experiments have shown a good match with these results, and more internally-consistent error estimates have usually yielded fairly similar-sized error bars.

A second, and probably more inclusive estimate for the error bars could in principle be obtained by looking at the scatter in the residuals in phase and visibility between the data frame and the best-solution superposition of the transformed star and iodine templates. Since such a residual is measured at each CCD column in the spectrum, it is easy to calculate an RMS value for the residuals, and it should be possible to translate

this to a final estimate of the internal error of the RV measurement. This would be useful in that it would include errors in the pipeline and data artifacts that are not captured by the photon limit estimate. At this stage, however, it would be dominated by the addition approximation error (section 5.10), in non-trivial ways. Since the addition approximation causes systematic errors that appear only over large velocity shifts, using this method to represent error bars would currently not be helpful on the scales where we are trying to make measurements: within small velocity shift regimes (i.e. over the short term), the error bars could appear much larger than the actual RMS scatter in the RV results.

An alternative way to estimate error bars is to propagate the standard errors from the fringe fitting (section 5.2) through the iodine extraction to the final result. This was the approach used early in the ET project, and to produce error bars for the 51 Peg results in section 6.1. To convert the standard errors in the visibility-phase vector for a given wavelength channel, $\Delta\gamma$ and $\Delta\phi$, to errors in Cartesian coordinates, Δx and Δy, as necessary for the linearisation in equations 4–6 and 4–7, we use the standard formalism for combining independent errors:[1]

$$\Delta x^2 = \left(\frac{\partial x}{\partial \gamma}\right)^2 \Delta\gamma^2 + \left(\frac{\partial x}{\partial \phi}\right)^2 \Delta\phi^2 \tag{5–1}$$

$$\Delta y^2 = \left(\frac{\partial y}{\partial \gamma}\right)^2 \Delta\gamma^2 + \left(\frac{\partial y}{\partial \phi}\right)^2 \Delta\phi^2, \tag{5–2}$$

where the Δ represents the respective measurement errors. Given $x = \gamma\cos\phi$ and $y = \gamma\sin\phi$, we find:

$$\Delta x^2 = \Delta\gamma^2\cos^2\phi + \Delta\phi^2\gamma^2\sin^2\phi \tag{5–3}$$

$$\Delta y^2 = \Delta\gamma^2\sin^2\phi + \Delta\phi^2\gamma^2\cos^2\phi. \tag{5–4}$$

[1] Note that x and y here are not related to the physical x, y positions on the CCD detector.

The system of equations for all the channels in the iodine extraction algorithm (section 4.4.3.1) can be weighted according to these errors, and error bars for the finally determined velocity shift can be extracted from the results using standard statistics for singular value decomposition (Press et al. 1992).

Equations 5–3 and 5–4 will in general be useful in that they may partially catch errors from the fringe fitting due to unexpected artifacts in the data reduction, such as non-uniform fringes and so forth. However, it is interesting to note that in the case of well behaved, evenly-distributed independent random errors, assuming $\Delta\gamma \ll \gamma$, it can be shown that:

$$\Delta\phi = \Delta\gamma/\gamma. \tag{5–5}$$

This relation follows simply from the dimensionality of ϕ and γ. If we represent the pair as a vector, $\boldsymbol{\gamma}$, it is reasonable to expect the error in this vector measurement to be an error circle in two dimensional space. The radius of this circle is equal to $\Delta\gamma$, so in the direction parallel to $\boldsymbol{\gamma}$ the error is simply $\Delta\gamma$. In the perpendicular direction, the error is $\gamma \tan \Delta\phi \approx \gamma\Delta\phi$, but since the error is circular, it must also equal $\Delta\gamma$ and hence we end up with equation 5–5. In Cartesians, the error in both axes is likewise given by the radius of the error circle, so that

$$\Delta x = \Delta y = \Delta\gamma. \tag{5–6}$$

Following a treatment similar to that in Ge (2002) for the measurement error in phase, it can also easily be shown that for an ideal sinusoidal fringe the error in the *visibility* measurement is given simply by

$$\Delta\gamma = \sqrt{\frac{2}{\overline{F}}}, \tag{5–7}$$

where $\overline{F}$ is the mean flux in the fringe, and is independent of the fringe visibility itself. For star and iodine data, the flux does not vary dramatically from fringe to fringe: it is therefore not unreasonable to make the approximation that $\Delta\gamma$, and hence Δx and Δy, are roughly constant across all channels, and solve the ensemble of equations with no

weighting. Owing to the continuum shapes of the stellar targets and the tungsten lamp used for obtaining the iodine spectra, along with the blaze function of the spectrum, the approximation could be improved, but the non-weighted approach is generally used for simplicity at the moment and has usually given good results.

5.2 Fringe Fitting

To assess the intrinsic accuracy of the curve-fitting routines used for measuring the phase and visibility of fringes, we ran Monte-Carlo simulations, simulating sinusoidal fringes with uniformly randomly distributed phases and various fixed values of visibility. Gaussian noise was added to simulate a signal/noise ratio (S/N) of 100 (a reasonable approximation to Poisson noise at this level of flux).

The results of such a simulation are shown in figure 5-1. 100 simulations were run for each point on the plot. The dotted line with diamond points represents the mean of the standard errors returned by the curve fitting routine, with the error bars representing the standard deviation in that mean over the 100 iterations; the thin solid line with the crosses represents the RMS difference between the input phase in the simulation and the phase measured by the curve-fit, with the error bars representing the expected statistical deviation in that RMS due to the finite number of simulations; and the thick solid line represents the theoretical photon limiting curve predicted by equation 2–31. It can clearly be seen that all three curves match very well. It is therefore not thought that there is any significant error due to the curve-fitting routines themselves.

5.3 Illumination Correction

Good illumination correction is important to be able to make good sinusoidal fits to the fringes. Normal flatfielding takes care of pixel-pixel gain variations, but only partially flattens the illumination function of a spectrum. The rest of the flattening depends on the self-illumination correction algorithm (section 4.4.1.5).

To the curve-fit routine, any illumination correction deficiency looks like additional noise in the fringe. We might therefore expect the flatfielding errors to become significant

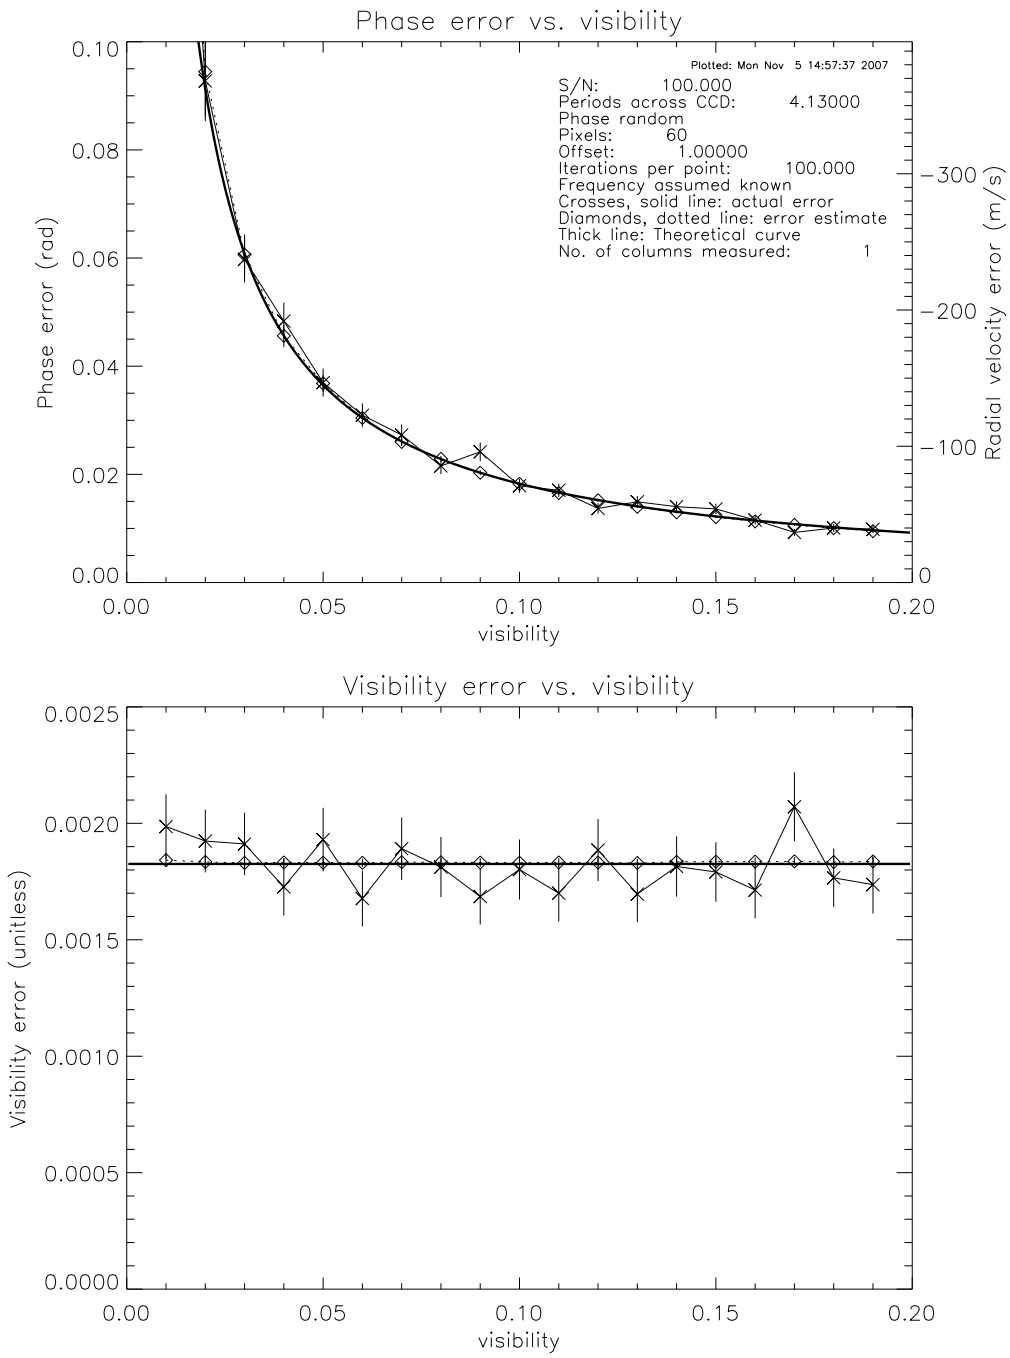

Figure 5-1. Phase and visibility errors due to curve-fitting alone, for a single wavelength channel. Thin solid line with crosses: RMS error from Monte-Carlo simulations; dotted line with diamond points: mean standard error estimate from curve-fitting; thick solid line: theoretical curve due to photon limit.

when they become comparable to the photon noise S/N ratio. A similar plot to figure 5-1, this time against S/N, (figure 5-2) shows that the errors begin to diverge from the predicted photon errors at S/N ratios above about 20, just as we would expect for a 5% (i.e. S/N=20) illumination correction error. (It is also interesting to note that the standard curve-fit returned errors apparently do not capture this effect at all.) It therefore seems reasonable to suggest that illumination correction needs to be good to better than the photon S/N for the spectrum; otherwise precisions are likely to reach a floor at the equivalent S/N of the illumination correction. In practice, the pipeline can allow extra parameters in fitting the sinusoids, however, to at least partially compensate for poor flatfielding, and so the errors are actually likely to be less severe than this estimate.

There are two particular effects which can trip up the illumination correction algorithm, requiring that some care be taken. The first is where, by pure coincidence, a large portion of a spectrum contains fringes that all happen to have similar phases, which we term 'fringe conspiracy'. In this case, the illumination correction algorithm will tend to interpret the aligned fringes as background illumination features, and divide them out, subtly affecting the phases of those and neighbouring fringes. Given the sheer number of spectra we obtain, this inevitably happens on occasion, and can cause significant systematic errors in the same way that a contaminating spectrum would (see section 5.7). The option exists in the pipeline software to try and overcome fringe conspiracy by performing additional smoothing of the extracted illumination function in the slit direction, in addition to that in the dispersion direction, to remove any residual fringes. Unfortunately this transverse smoothing also has the effect of removing any genuine features in the illumination function that have similar spatial scales to the fringes, which can themselves contribute systematic errors if not successfully removed.

Both these effects are very hard for the illumination correction algorithm to distinguish from genuine fringes. They can best be considered as contaminating spectra, rather like residual comb (section 5.9): for features at the 0.5% fractional level in

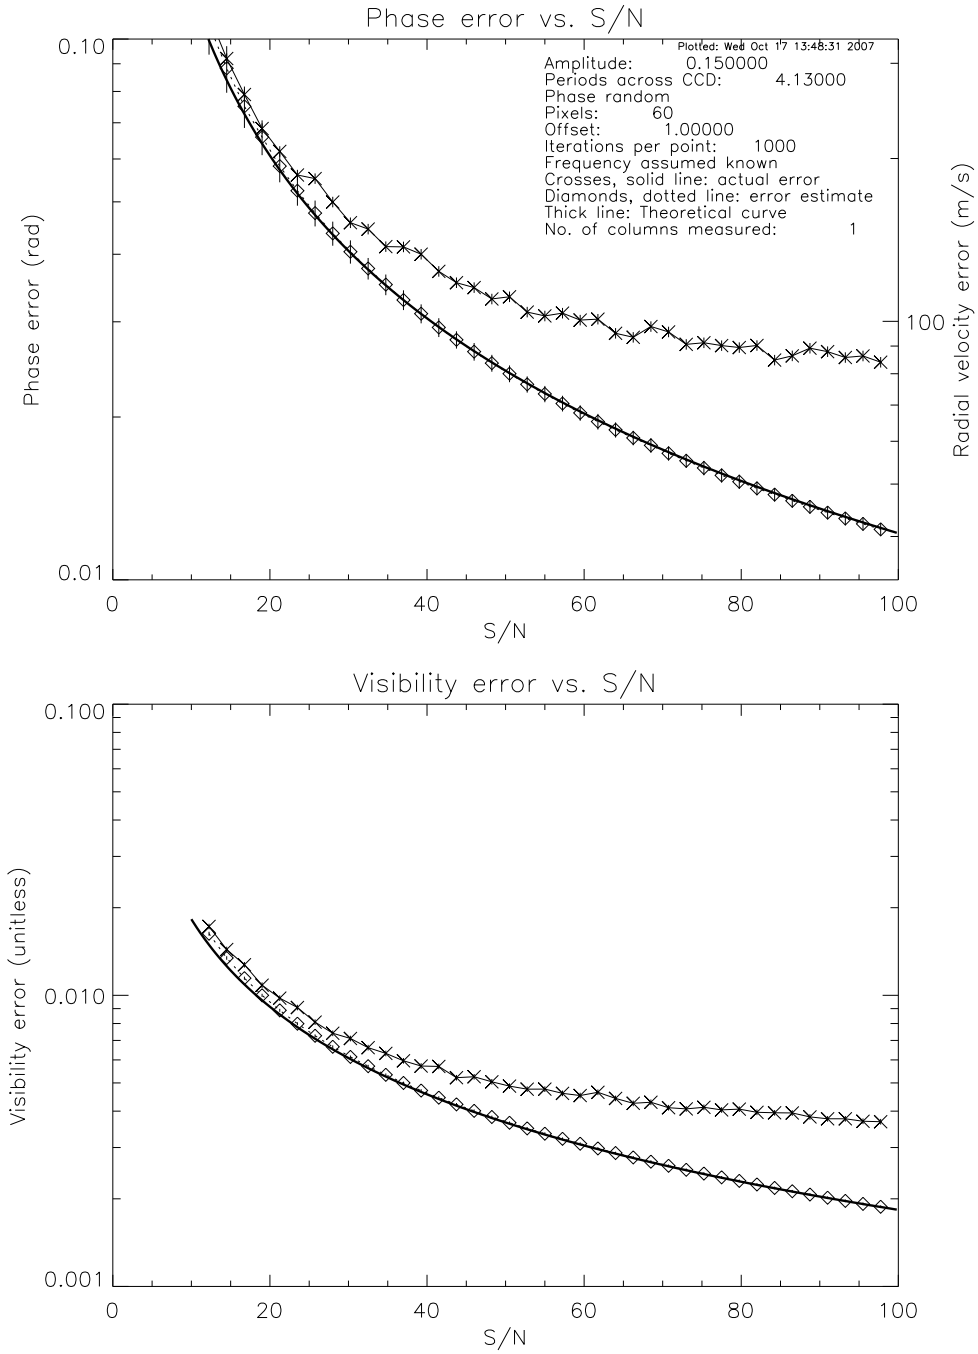

Figure 5-2. Fringe fitting errors for poor illumination correction. Thin solid line with crosses: actual RMS error from Monte-Carlo simulations; dotted line with diamond points: mean standard error estimate from curve-fitting; thick solid line: theoretical curve due to photon limit. 1000 Monte Carlo simulations are performed per point, with a 5% modulation of the fringes along the slit length. The actual RMS errors begin to deviate from the photon limit at $S/N > 100\%/5\% = 20$.

illumination, we might again expect errors of around $9\,\mathrm{m\,s^{-1}}$ just as for residual comb – though in this case the error will probably be rather smaller because the features are more likely to be appear as single small bumps modulating the normal fringes, rather than full superposed fringes.

Ultimately, if there are illumination features which have similar spatial scales and amplitudes to the actual fringes being measured, it is extremely hard, if not impossible, to separate them out. The instrument therefore needs to be built to provide as smooth an illumination function as possible.

For all these types of illumination problem, it is difficult to establish to what degree these errors will be systematic and correlated from data point to data point (though it is unlikely that they will be purely random). Occasional problems have certainly been seen in the past due to the fringe-conspiracy effect. It is hoped that carefully prepared simulations run through the pipeline will enable us to gain a more clear picture of how illumination correction affects RV results.

5.4 Slant Correction

Correction of slant in the spectral lines to match the CCD columns (section 4.4.1.4) is an important step for achieving the best possible precision. Spectra are corrected for slant by shifting each row of the spectrum in the dispersion direction until all the rows are aligned, with the slit image following exactly along the CCD column direction. We can gain a crude handle on the required precision of the correction by considering the effects of shifting an entire spectrum in the dispersion direction by sub-pixel amounts. We take a fully processed data whirl (from real data that have been slant corrected as well as possible), shift it by small amounts in the dispersion direction using interpolation, and then measure the apparent ensemble phase shift across all wavelength channels between the shifted and unshifted frames. The results of such an experiment from an early 51 Peg fringing spectrum taken with the single-object ET are shown in figure 5-3.

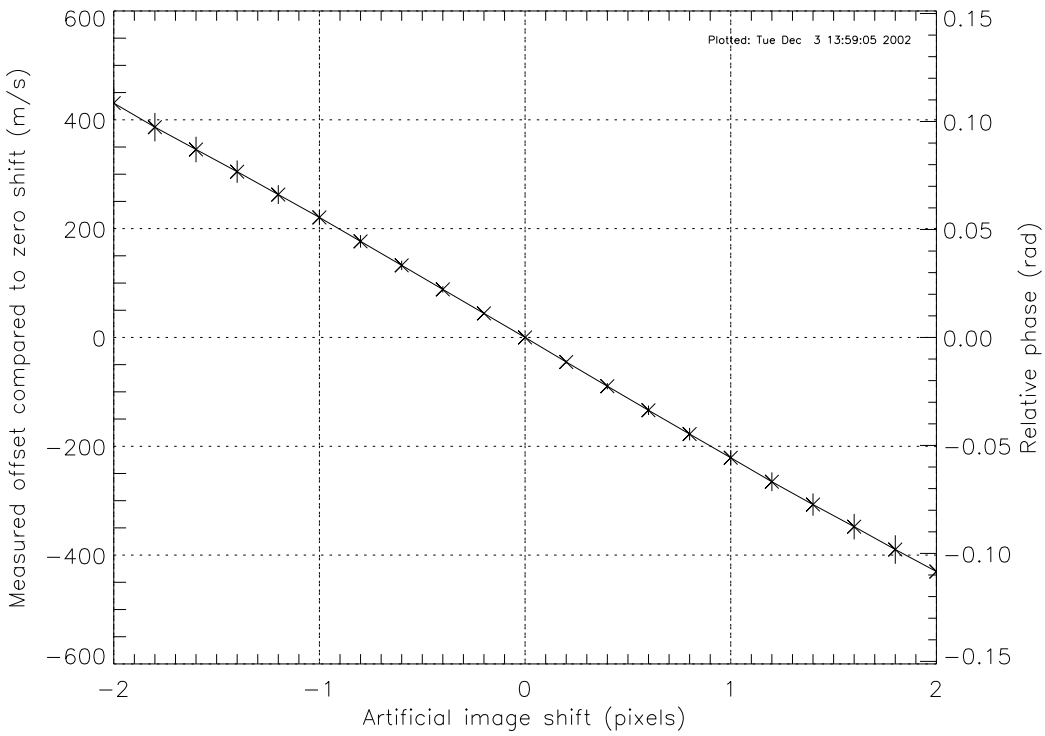

Figure 5-3. Effect of uniform artificial spectrum shift in the dispersion direction.

The magnitude of the effect depends on the particular spectrum used, and the way the fringes for that spectrum happen to align over the full wavelength coverage of the spectrum. Figure 5-3 shows one of the more extreme cases, showing an error of $\sim 200\,\mathrm{m\,s^{-1}}$ per pixel of shift. For an acceptably small error of, say, $1\,\mathrm{m\,s^{-1}}$ (smaller than the photon noise in almost all cases), we would therefore require an image shift of less than 0.005 pixels. We can think of this as a worst-case limit for the precision of the dispersion shift to be applied to each single row cut along the spectrum. Since in reality the slant correction is not applied uniformly to each row, but rather with a roughly monotonic variation from positive to negative along the slit direction, the actual precision requirement is likely to be somewhat less stringent. In practice the slant correction shows an RMS precision of the order of ~ 0.01 pixels, and this is deemed sufficiently accurate for the time being. (This figure is determined by taking slant-corrected images, re-measuring the slant as a function of row number, and calculating the RMS about zero.)

5.5 Flux Centroiding

Flux centroiding has not yet been fully implemented in the data reduction pipeline. Flux centroiding involves finding a flux-weighted centre time for the exposure, so that the time of exposure remains meaningful even if there is temporary cloud cover or a loss of telescope guiding, for example. The diurnal motion of the Earth can change by as much as $\sim 1\,\mathrm{m\,s^{-1}}$ per minute (simply the centripetal acceleration at the equator). Over the course of a one hour exposure, therefore, this can amount to a very significant change. We end up with essentially an averaged velocity over the length of the exposure, and it is important to use a correct and consistent definition for the time of the exposure.

At the moment, we default to simply using the exact mid-time of the exposure to calculate the appropriate barycentric correction. In practice, the flux levels that we obtain are coincidentally such that the errors due to this approximation are generally sufficiently small compared to the photon errors that the approximation is not a major concern. Tests using the PMT records (section 3.1) show that, except in cases of extreme varying cloud cover (where data is likely to be rejected anyway), flux weighting usually modifies the time-of-exposure by ∼1–10% of the exposure length. Even taking the extreme example of complete cloud coverage for just the second half of an exposure would move the time-of-exposure forward by a only a quarter of the exposure time. For a one hour exposure, we might therefore expect, even in such a particularly bad case, an error of 15 min, or $\sim 15\,\mathrm{m\,s^{-1}}$ in barycentric correction. For the multi-object Keck ET, targets that require such long exposures usually have photon errors at the $\sim 30\,\mathrm{m\,s^{-1}}$ level or more, so the flux centroiding error is not dominant. For short exposures at Kitt Peak on bright targets, we might take, say, 5 min exposures at the shortest. In a very bad case, this might lead to a $\sim 1.3\,\mathrm{m\,s^{-1}}$ error: again, this is sufficiently below the photon limit (usually 2–3 $\mathrm{m\,s^{-1}}$ at best) that flux centroiding, while significant, is not a primary concern.

These errors are mentioned here since they are relevant to the results reported so far with ET. Nonetheless, flux-centroiding is easy to implement, the hardware is already in

place, and it is simply a matter of dropping in the appropriate mechanisms for including the data into the pipeline.

5.6 Barycentric Correction

Correction of RV to the Solar System barycentric frame of reference (section 4.4.4) can be a very significant source of systematic error if not done accurately. Fortunately the problem has already been fairly thoroughly addressed in the field of astronomy. A comparison of our barycentric correction routine to *helio*centric velocities given by the IRAF routine, 'rvcorrect' (Tody 1986, 1993),[2] spanning 100 years from July 2004, gives an RMS difference of $3.2\,\mathrm{m\,s^{-1}}$, within the quoted precision of $5\,\mathrm{m\,s^{-1}}$ stated in the IRAF documentation.

The rvcorrect routine, however, does not handle the necessary full correction to the Solar System barycentre, so that perturbations due to the Solar System planets are not accounted for. A comparison made by Scott Fleming in early 2006 of the full barycentric correction against independent software used by William Cochran of the Texas Planet Search team showed agreement to within $1\,\mathrm{m\,s^{-1}}$ (neglecting gravitational redshift and transverse Doppler shift) (W. Cochran, 2006, private communication).

5.7 Contaminating Spectra

In order to estimate the kinds of errors that we expect from various other sources, it is useful to calculate a rough estimate of the errors due to contaminating background spectra. This formalism will enable us to calculate the effect of background moonlight contamination or double-lined spectroscopic binaries, for example. In addition, by treating any residual unfiltered comb presence and the iodine/star cross-talk term in equation 2–24 as background spectra, we can also try to assess their relative significance.

Figure 5-4 shows a fringe along one detector column due to the target source alone, with fringe amplitude a_s, mean flux F_s, and phase ϕ_s. For simplicity we assume no

[2] Also see http://iraf.noao.edu/, last accessed Oct 2007.

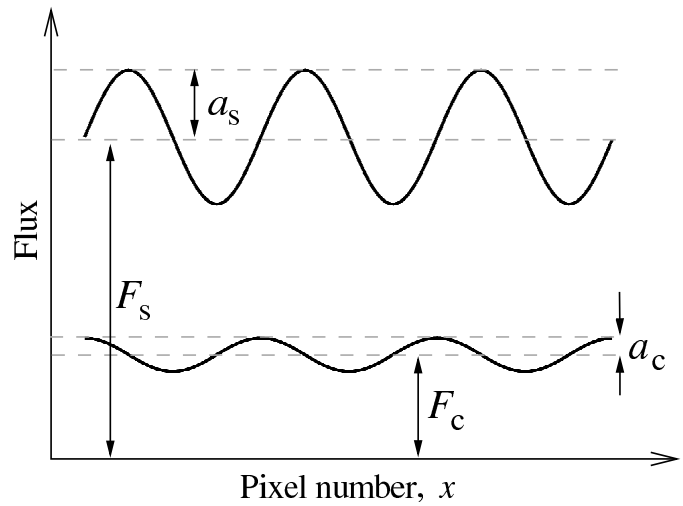

Figure 5-4. Fringe along one channel due to target (upper curve), and contaminating low flux fringe (lower curve). Measured fringe is a summation of these two fringes.

iodine fiducial reference, since we are only aiming for an order-of-magnitude estimate. A second contaminating fringe of lower amplitude a_c and mean flux F_c due to background contamination is also shown, with phase ϕ_c. If the spatial frequency of the fringes is f, then the summation of these two fringes will give the total measured fringe:

$$\begin{aligned} & F_s + \Re\{a_s e^{i(fx+\phi_s)}\} + F_c + \Re\{a_c e^{i(fx+\phi_c)}\} \\ = \quad & F_s + F_c + \Re\{a_s e^{i(fx+\phi_s)} + a_c e^{i(fx+\phi_c)}\}, \end{aligned} \tag{5–8}$$

where x identifies position along the slit. $F_s + F_c$ represents the mean value of the measured flux. The last term represents the varying sinusoidal fringe.

We are interested in the phase error, ε_ϕ, introduced into the measured fringe by the contaminating spectrum. Since we are only interested in the phase information, we ignore the offset term $F_s + F_c$, and represent the varying term as a vector summation, as shown in figure 5-5, where a_s and a_c represent the source and contaminant fringe amplitudes as before. The angle $\Delta\phi$ is the difference between the source and contaminant fringe phases, $\Delta\phi = \phi_c - \phi_s$. Using the sin and cosine rules for triangles we can show:

$$\frac{\sin \varepsilon_\phi}{a_c} = \frac{\sin \Delta\phi}{\sqrt{a_s^2 + a_c^2 + 2a_s a_c \cos \Delta\phi}}. \tag{5–9}$$

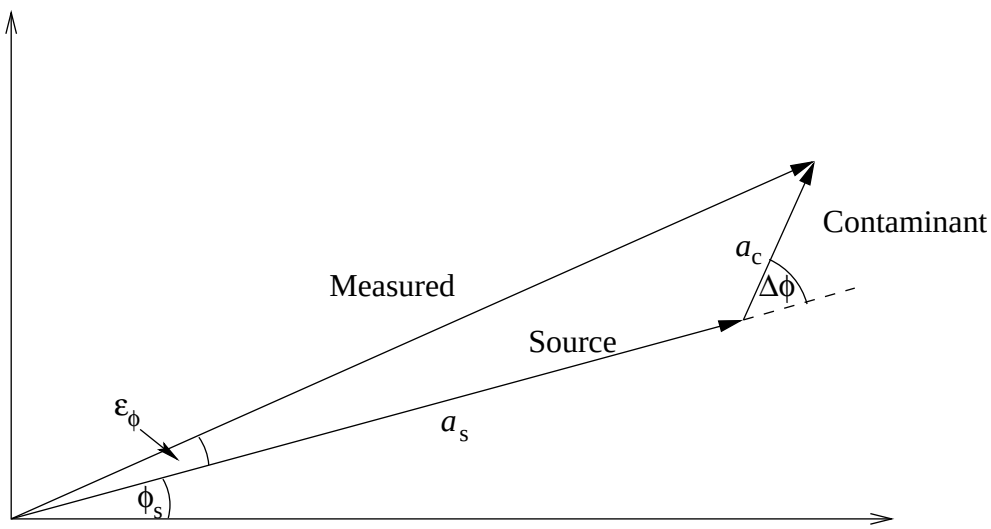

Figure 5-5. Vector representation of the summation of the fringes due to the target source and background contamination.

In the limit that the two spectra are of similar form and very close in velocity, so that $\Delta\phi$ is *small*, then

$$\sin\varepsilon_\phi \approx \frac{a_\mathrm{c}\Delta\phi}{a_\mathrm{s}+a_\mathrm{c}}. \tag{5–10}$$

If we assume the source and contaminant fringe visibilities are approximately equal, so that $a_\mathrm{s}/F_\mathrm{s} \approx a_\mathrm{c}/F_\mathrm{c}$, then $a_\mathrm{c}/a_\mathrm{s} \approx F_\mathrm{c}/F_\mathrm{s}$, which is equal to the flux ratio of the two fringes. If the contaminating fringe is much fainter than the source, so that $F_\mathrm{s} \gg F_\mathrm{c}$ and $a_\mathrm{s} \gg a_\mathrm{c}$, and ε_ϕ is therefore small, then

$$\sin\varepsilon_\phi \approx \varepsilon_\phi \approx \frac{a_\mathrm{c}}{a_\mathrm{s}}\Delta\phi = \frac{F_\mathrm{c}}{F_\mathrm{s}}\Delta\phi. \tag{5–11}$$

Still assuming that the two spectra are of similar type and velocity, then all wavelength channels will see approximately the same phase difference between source and contaminant fringes, and this error will be systematically close to the same across all wavelength channels. Therefore the same result will be expected finally even after averaging over all channels. Since $\Delta\phi$ is proportional to the difference in velocity, Δv, between source and

contaminant, then we can write the final error in the measured velocity, ε_v, simply as:

$$\boxed{\varepsilon_v \approx \frac{F_\mathrm{c}}{F_\mathrm{s}} \Delta v.} \tag{5–12}$$

In other words, the systematic velocity error in a single fringe is simply the velocity difference scaled by the flux ratio of the contaminant to the source fringe.

This relation does not hold well to arbitrarily large velocities, however. From the geometry of figure 5-5 it can be seen that a worst case scenario is where the contaminant in all wavelength channels is systematically offset by an amount such that the background contaminant vector is perpendicular to the measured vector (or, approximately, where $\Delta\phi = \pi/2$). In this case, $\varepsilon_\phi \approx a_\mathrm{c}/a_\mathrm{s} \approx F_\mathrm{c}/F_\mathrm{s}$, so that:

$$\boxed{\varepsilon_v \approx \Gamma \frac{F_\mathrm{c}}{F_\mathrm{s}}}, \tag{5–13}$$

where Γ is the phase/velocity scaling factor.

In the limit that the spectra are completely uncorrelated, or are sufficiently separated in velocity space that overlapping features are in no way correlated, then the phase errors will be randomly distributed across all wavelength channels. Following again from equation 5–9, again assuming $F_\mathrm{s} \gg F_\mathrm{c}$ and $a_\mathrm{s}/F_\mathrm{s} \approx a_\mathrm{c}/F_\mathrm{c} \Rightarrow a_\mathrm{c}/a_\mathrm{s} \approx F_\mathrm{c}/F_\mathrm{s}$, but now taking $\Delta\phi$ as uniformly randomly distributed, we can find the RMS value for the phase error in one channel as:

$$\mathrm{rms}(\sin \varepsilon_\phi) \approx \mathrm{rms}(\varepsilon_\phi) \approx \mathrm{rms}(\frac{a_\mathrm{c}}{a_\mathrm{s}} \sin \Delta\phi) = \frac{F_\mathrm{c}}{F_\mathrm{s}} \frac{1}{\sqrt{2}} \tag{5–14}$$

Assuming an average over n independent wavelength channels gives a $1/\sqrt{n}$ reduction in the final error, so that for uncorrelated spectra, we can expect a final velocity error of:

$$\boxed{\varepsilon_v \approx \frac{\Gamma}{\sqrt{2n}} \cdot \frac{F_\mathrm{c}}{F_\mathrm{s}}}, \tag{5–15}$$

where Γ is the phase-velocity scale factor for the instrument. The phase error is now independent of differential velocity between source and contaminating spectra, since the two spectra no longer bear any relation to each other.

5.8 Moonlight Contamination

Equations 5–12 to 5–15 can be applied directly to estimate the magnitude of the errors introduced by background scattered moonlight contamination. A $3''$ fibre with a bright-time sky background of $19\,\mathrm{mag\,arcsec^{-2}}$ due to scattered moonlight from the atmosphere gives a total of 16.9 mag of sky background. For a magnitude 12 star, this gives a source-to-contamination flux ratio of about 90. In the worst case scenario, from equation 5–13, and assuming $\Gamma \sim 3300\mathrm{m\,s^{-1}\,rad^{-1}}$ (roughly correct for the KPNO ET), we find $\varepsilon_v \approx 37\,\mathrm{m\,s^{-1}}$. This will apply where the stellar spectrum is similar to the moonlight spectrum (not uncommon, since most targets are sun-like), and in the case where the velocity difference between star and moonlight, Δv, is coincidentally around $\sim 5\,\mathrm{km\,s^{-1}}$. At smaller velocities, the error will scale roughly linearly as $\varepsilon_v = \Delta v/90$ up to this point. After that, it will improve again as $\Delta\phi$ increases to π, where the phase error once again approaches zero. As $\Delta\phi$ increases, the behaviour is likely to be somewhat oscillatory, until Δv is large enough that the two spectra are completely uncorrelated. For $n = 1000$ *independent* wavelength channels (ie. 4000 pixel channels with an LSF ~ 4 pixels wide), the error should then approach equation 5–15, with a value of around $\varepsilon_v \approx 0.8\,\mathrm{m\,s^{-1}}$. This should also be the typical error size when the star and moon spectra are very different in form. Simulations of the effect of moonlight contamination by Mahadevan show reasonable agreement. Figure 5-6 shows the RV deviation caused by synthetic moonlight contamination added to a synthetic stellar spectrum, and then multiplied by the interferometer response function and degraded in resolution to simulate real instrument spectra. The resulting spectra are put through the standard pipeline (ignoring the iodine reference) to assess the effects of the contamination.

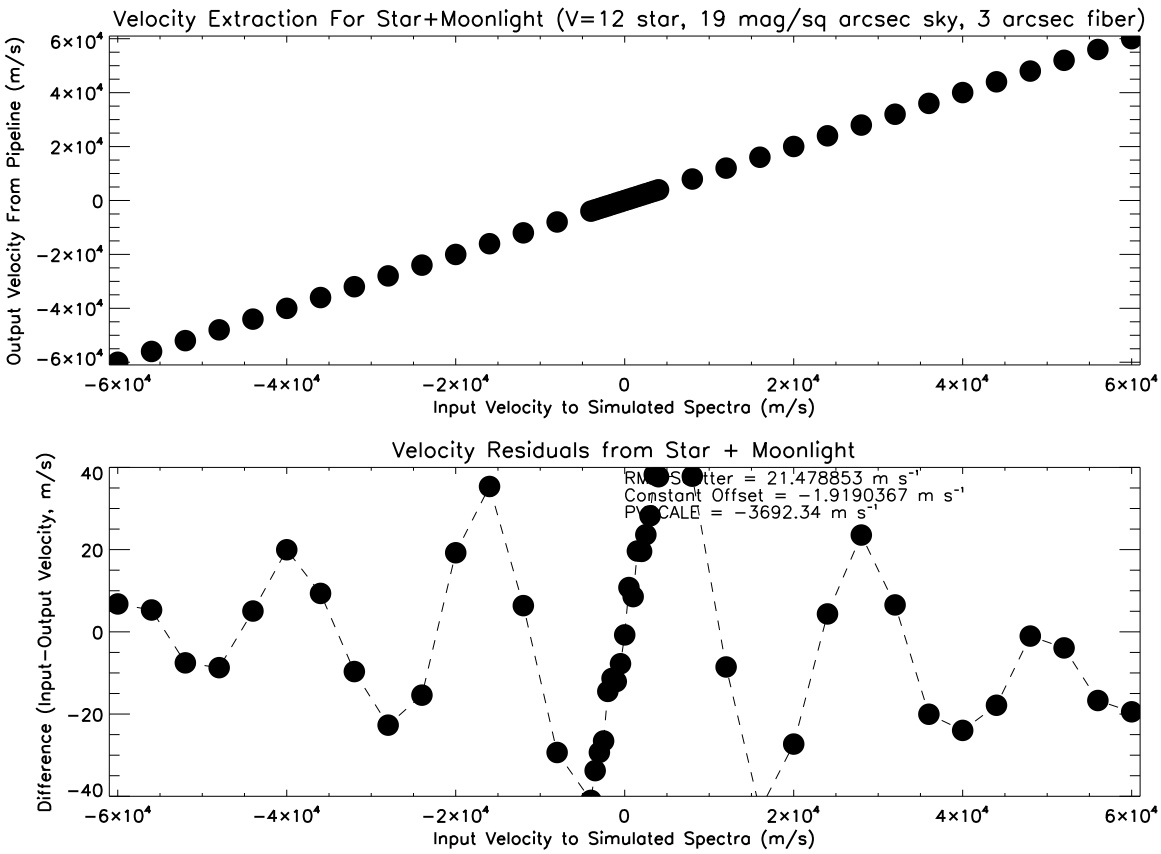

Figure 5-6. Simulations of moonlight contamination, showing the systematic error introduced by contaminating moonlight at 19 mag arcsec^{-2} for a V=12 F9V star on a 3″ fibre, as a function of velocity difference between target star and moon spectrum. [*Courtesy of Suvrath Mahadevan, private communication.*]

5.9 Residual Interferometer Comb

If the interferometer comb is not completely removed – either by careful tuning of interferometer delay and slit width or by Fourier filtering in the data processing – then it acts like a contaminating spectrum. This has in the past been a problem with the Keck ET data in particular: in some cases the sampling was such that the comb was aliased in places, creating a low frequency pattern in the dispersion direction which was impossible to filter out without losing Doppler information.

We can follow the same formalism as for the moonlight contamination. In this case, however, instead of the source and contaminant fringe visibilities being similar, the fluxes are similar, so that $F_{\rm s} \approx F_{\rm c}$. Dividing the denominators of equation 5–9 through by

$F_\mathrm{s} \approx F_\mathrm{c}$, we can replace the fringe amplitudes a_s and a_c with their respective visibilities $V_\mathrm{s} \equiv a_\mathrm{s}/F_\mathrm{s}$ and $V_\mathrm{c} \equiv a_\mathrm{c}/F_\mathrm{c}$. The source spectrum and interferometer comb are completely unrelated in form. Assuming $V_\mathrm{c} \ll V_\mathrm{s}$ we can follow the same reasoning as for equation 5–15 and write:

$$\boxed{\varepsilon_v \approx \frac{\Gamma}{\sqrt{2n}} \cdot \frac{V_\mathrm{c}}{V_\mathrm{s}}}, \tag{5–16}$$

so the error is now proportional to the ratio of visibilities. For a residual comb visibility of 0.5% on top of a spectrum of typical mean visibility of say, 4%, and again taking $\Gamma \sim 3300\mathrm{m\,s^{-1}\,rad^{-1}}$ and $n = 1000$ independent wavelength channels, this gives an expected error of $\varepsilon_v \approx 9\,\mathrm{m\,s^{-1}}$.

5.10 The Addition Approximation

In order to estimate the errors introduced by the addition approximation discussed in section 2.6.1, we can also follow a similar approach, treating the cross term from equation 2–24 which is ignored in the approximation (or rather, treating the *lack* of cross term) as if it were a contaminating spectrum. First, we consider the simplified case of two discrete overlapping Gaussian absorption lines, from template spectra labelled A and B (e.g. an iodine and a stellar line), combined by multiplication to give the measured spectrum, labelled M. Both line centres are exactly coincident. The fractional line depths are represented by D ($0 \leq D \leq 1$), with corresponding subscripts a, b and m. From Ge (2002), we have that:

$$\gamma = De^{-3.56d^2/l_\mathrm{c}^2} = KD, \tag{5–17}$$

where γ is the absolute fringe visibility (i.e. $\mathbf{V} = \gamma e^{\mathrm{i}\phi}$), d is the interferometer delay, $l_\mathrm{c} = \lambda^2/\Delta\lambda$ is the coherence length of the interferometer beam with line width $\Delta\lambda$ at wavelength λ, and $K = \exp(-3.56d^2/l_\mathrm{c}^2)$. Although not very realistic, we assume both lines A and B and the resulting line M are of similar width, and that the measured line, which is the product of the two lines, is also approximately Gaussian. K is then approximately

the same for all three lines. We can then write:

$$\begin{aligned}\gamma_{\mathrm{m}} &\approx D_{\mathrm{m}}K = [1-(1-D_{\mathrm{a}})(1-D_{\mathrm{b}})]K \\ &= [D_{\mathrm{a}}+D_{\mathrm{b}}-D_{\mathrm{a}}D_{\mathrm{b}}]K \\ &= \gamma_{\mathrm{a}}+\gamma_{\mathrm{b}}-D_{\mathrm{a}}D_{\mathrm{b}}K. \end{aligned} \tag{5–18}$$

In the addition approximation, the complex visibilities of the template spectra are added together. In this simple case, the two lines are centred at the same wavelength and both are Gaussian, so that one line is simply a scaled version of the other. By the linearity of Fourier transforms, this means that the phases of the two complex visibilities must be identical, so that in the addition approximation, the two absolute visibilities add to give $\gamma_{\mathrm{m}} \approx \gamma_{\mathrm{a}} + \gamma_{\mathrm{b}}$. The remaining term in equation 5–18, $D_{\mathrm{a}}D_{\mathrm{b}}K = \gamma_{\varepsilon}$, is therefore approximately the error, the difference between the added templates and the actual measured visibility.

In the more general case that the two line centres are not exactly coincident or the same shape, so that the respective template fringes are not in phase, the error term will also include a phase difference, becoming a two dimensional vector, $\gamma_{\varepsilon}e^{\mathrm{i}\phi_{\varepsilon}}$. Taking the error term above as a reasonable estimate of the length of this vector and assuming ϕ_{ε} is uniformly randomly distributed, we can calculate a corresponding representative error in phase of the summation approximation. Figure 5-7 shows the addition of the "true" (measured) complex visibility and the error term to give the solution according to the summation approximation (compare to figure 5-5). ϕ represents the phase of the true complex visibility, and ε_{ϕ} represents the error in the measurement of that phase. If we assume the resulting measured visibility vectors and the error terms are uncorrelated from channel to channel, and if we take D_{a} and D_{b} to be some kind of representative average line depth for the two spectra across all j, we can derive the typical expected velocity

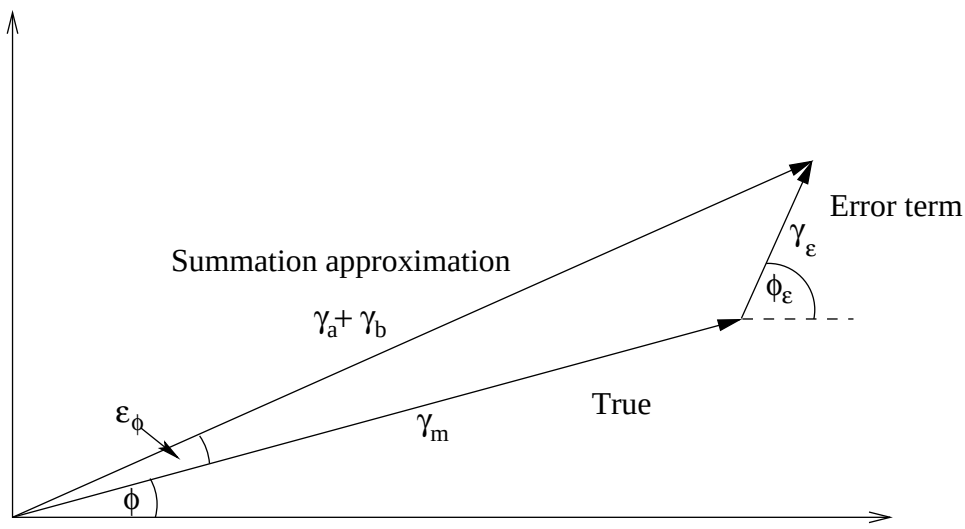

Figure 5-7. Vector representation of the summation of the true complex visibility and the error term due to the addition approximation.

error following the same reasoning as for equation 5–16 and write:[3]

$$\begin{aligned} \varepsilon_v &= \frac{\Gamma}{\sqrt{2n}} \cdot \frac{\gamma_\varepsilon}{\gamma_\mathrm{m}} \\ &= \frac{\Gamma}{\sqrt{2n}} \cdot \frac{D_\mathrm{a} D_\mathrm{b}}{D_\mathrm{a} + D_\mathrm{b} - D_\mathrm{a} D_\mathrm{b}}, \end{aligned} \tag{5–19}$$

where n is again the number of independent wavelength channels.

Figure 5-8 shows the expected typical error as a function of average line depth for the simplified case where the typical depths of the two spectra are equal. For average line depths of, say, 80% for both star and iodine, and again taking $\Gamma \sim 3300\,\mathrm{m\,s^{-1}\,rad^{-1}}$ and $n = 1000$ this gives a typical error due to the addition approximation of $\sim 50\,\mathrm{m\,s^{-1}}$, which is clearly very significant. The error will manifest as a systematic error in the velocity response of the instrument, essentially adding noise which varies as a function of the specific overlapping of the lines between target star and reference spectrum. It

[3] Although angles ϕ_ε in figure 5-7 and $\Delta\phi$ in figure 5-5 are measured from different origins, they are in both cases taken to be uniformly randomly distributed between 0 and 2π, so that the same reasoning applies for both.

will therefore vary with stellar spectral type, class and line width, and will also vary as a function of the intrinsic absolute Doppler shift of the stellar spectrum. Since the stellar lines are generally considerably broader than the iodine lines, if the stellar lines slowly shift relative to the iodine lines, the noise term will slowly change until the point where a shift of more than a stellar line width has been reached. At this point, the stellar lines are overlapping completely new iodine features, and the noise term will take on a new value that is completely uncorrelated with its previous value. Hence, we expect a non-linearity in the velocity response of the instrument, with a standard deviation somewhere on the order of $50\,\mathrm{m\,s^{-1}}$ and that varies with Doppler shift on a scale of approximately the line width of the star. For solar-type stars observable with ET, this variation will be over scales typically on the order of 5–$10\,\mathrm{km\,s^{-1}}$.

Figure 5-9 shows the results of simulated fringing spectra run through the pipeline by Mahadevan to see the effect of non-linearity due to the addition approximation, and shows broad agreement with these expectations. We note here that, along the same lines, using the correct, flux-dependent, and hence channel-dependent, visibility scaling factors, K_s and K_i (equation 2–27) is also of significance. Since these have only recently been fully understood, and had been previously been allowed to be free parameters that are constant over all channels, the full form has not yet been included in the pipeline. This omission can itself be expected to contribute some systematic error. Mahadevan has performed tests with simulations to model the cross-talk term in the addition approximation, and found that including the full flux dependency along with the cross-talk term is indeed essential to avoid systematic errors. However, including the flux-dependency alone without including the cross-talk term left remaining systematic errors on a similar scale to those found without including either.

Clearly the addition approximation is a very significant source of systematic error, and it will be essential to solve the problem. It can, however, be mitigated in the meantime by judicious selection of observation times and positions of targets on the

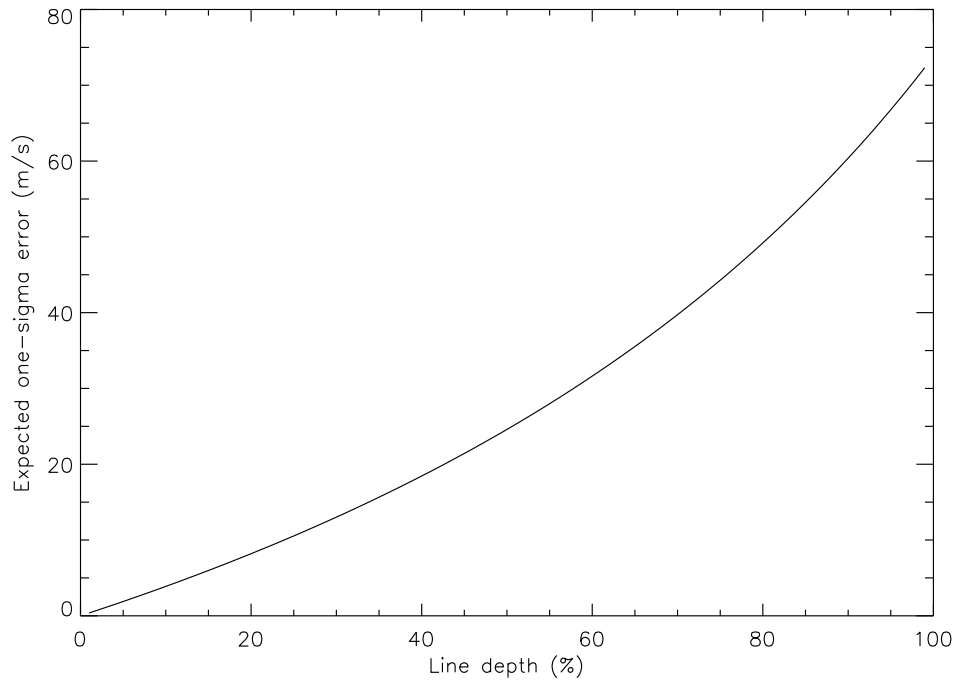

Figure 5-8. Analytically calculated expected error due to the addition approximation, assuming approximately equal line depths for both star and reference spectra.

sky, so that the barycentric motion of the Earth – usually the dominant effect that causes the non-linearity to become significant – is minimised. Such observations are often not hard to achieve at least over periods of a few days, and because of this fact, we have been still been able to make many useful detections of genuine astrophysical RV signals.

5.11 Other Sources of Error

The list of error sources addressed here is by no means exhaustive, and there are many other sources still to be addressed. Examples include how much the assumption of uniform phase-shift across all wavelengths introduces systematic errors, how much the 'bulk-shift compensation' algorithm (4.4.3.2) corrects this effect, and how much the same algorithm introduces errors of its own. (Although future versions of the software should render this problem irrelevant by no longer making such an approximation.) The effect of LSF variation along the length of the slit, caused by non-uniform slit illumination, also has yet to be quantified, as has the effect of distortion along the length of the slit giving rise to non-uniform spatial frequency of the fringes. Likewise, it would be good to establish

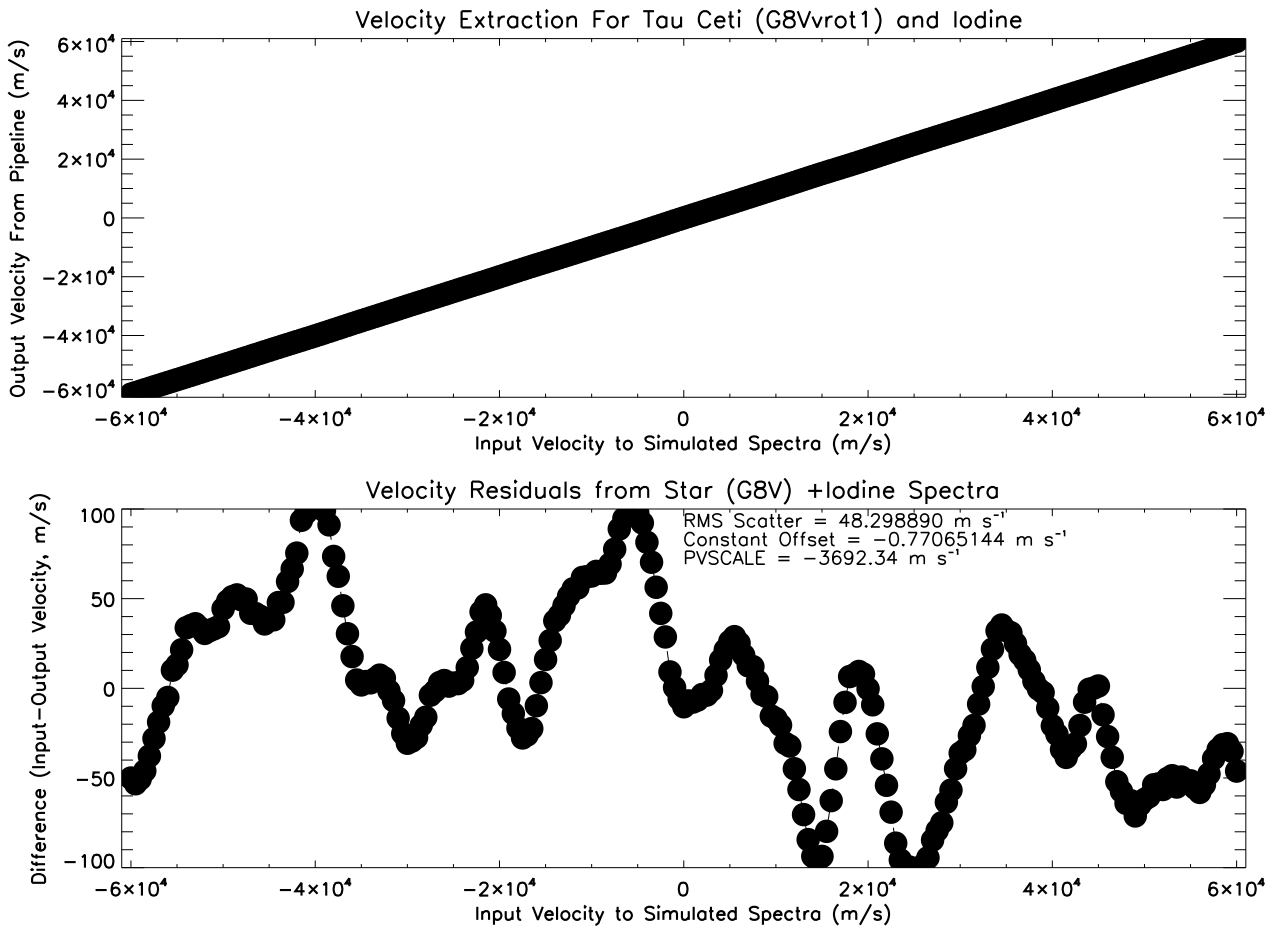

Figure 5-9. Simulations showing the addition approximation error, by Mahadevan. The non-linearity in the RV response has the same order of magnitude and occurs on the same input velocity scale as expected from theoretical predictions. [*Courtesy of Suvrath Mahadevan, private communication.*]

at what lower flux limit the data reduction algorithms begin to break down so that the results no longer follow simple photon noise predictions, and to better establish the noise floor at which point there is no gain to be had by increasing the photon flux.

Many of these issues will be easier to address once Mahadevan's code for producing raw data simulations is fully operational: it should then in principle be possible to simulate various effects one at a time in a controlled manner, run the simulations through the pipeline, and establish semi-empirically the size of the error caused by each one.

5.12 Planet Detection Limits

Figure 5-10 shows the velocity semi-amplitude, K, expected in the reflex motion of a star due to a planetary companion as a function of orbital period, at different minimum planet masses ($m \sin i$). This can be used as a rough guide to the masses of detectable planet companions at a given error level, at least in the case of random uncorrelated errors. K scales relatively weakly as the $-2/3$ power of stellar mass, where masses are generally in the range ~ 0.7–$1.6\,M_\odot$ for the ET surveys. At the current best ET precision of $\sim\ 3\,\mathrm{m\,s^{-1}}$ short-term for bright Sun-like stars, we find a detection limit of about $m \sin i\ >\ 0.1\,M_\mathrm{J}$ for hot-Jupiters (period $<\ 10\,\mathrm{d}$), at the 3-sigma confidence level. At typical Keck ET precisions of 30–40 $\mathrm{m\,s^{-1}}$, we might expect something more like a $1.0\,M_\mathrm{J}$ detection limit at periods less than 10 d.

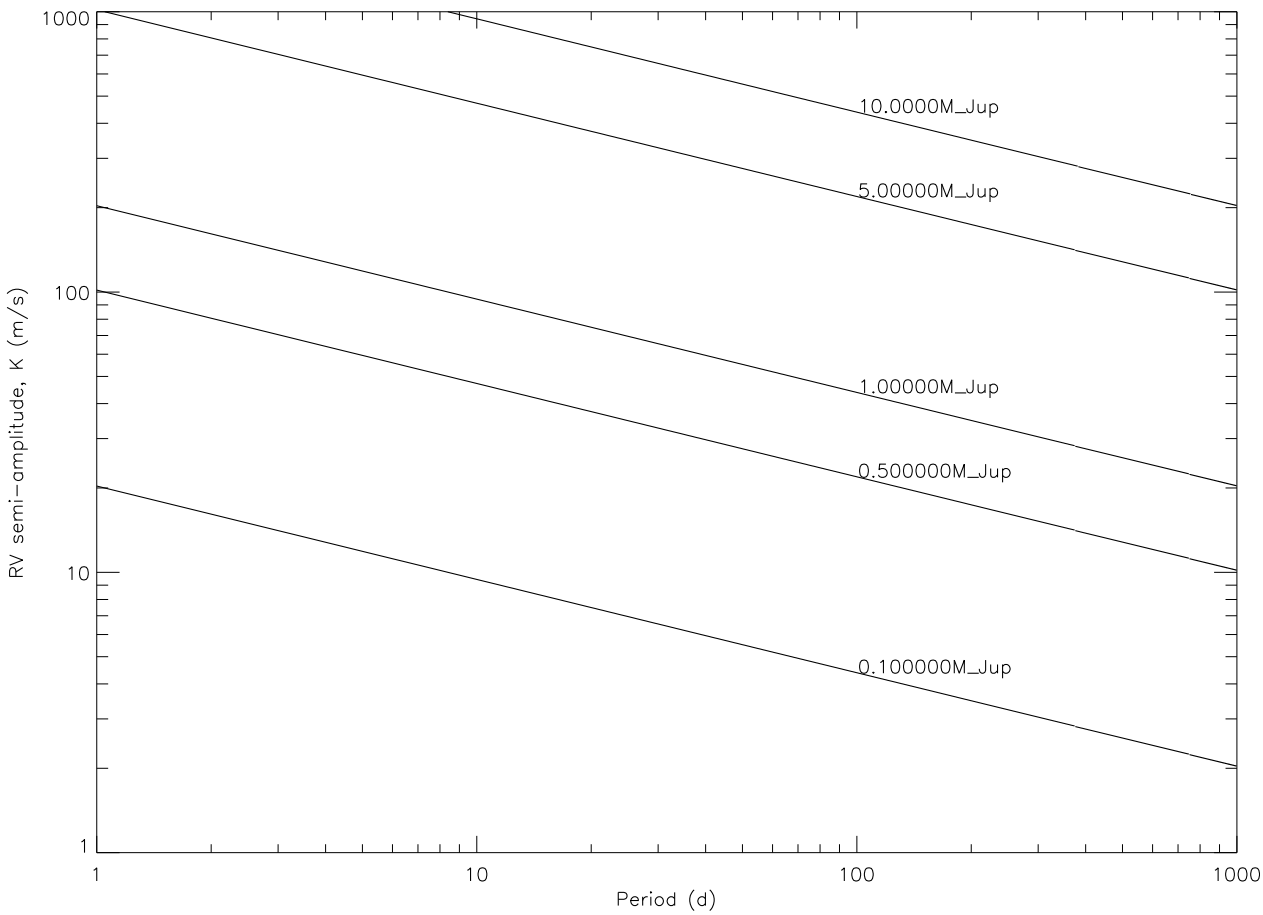

Figure 5-10. Radial velocity semi-amplitude, K, for different minimum planet masses ($m \sin i$), assuming a $1\,M_\odot$ primary and zero-eccentricity orbits.

CHAPTER 6
RESULTS FROM THE KPNO ET

The single-object ET instrument installed at the KPNO 2.1 m telescope is the precursor to the multi-object Keck ET. It was used for the confirmation of the well-established planetary companion to 51 Pegasi, marking the first ever planet detection using the DFDI technique (van Eyken et al. 2003, 2004a). It was also used to make the first actual discovery of an exoplanet using the DFDI technique (Ge et al. 2006a), the find being designated ET-1 (HD 102195b). An overview of the instrument itself, along with references to more detailed descriptions, is given in chapter 3.

6.1 Confirmation of 51 Peg b

The detection of 51 Peg b was made using a temporary prototype during a test run in August 2002 (section 3.2.1), showing a root-mean-square (RMS) scatter of 11.5 $\mathrm{m\,s^{-1}}$ in the measurements over five days. The results were reported in van Eyken et al. (2004a), where we also showed comparison measurements of the RV-stable star, η Cas, which demonstrated an RMS scatter of 7.9 $\mathrm{m\,s^{-1}}$ over seven days, starting to approach the precision levels obtained with traditional RV techniques based on cross-dispersed echelles. These results are summarised here.

6.1.1 Observations

During the 2002 test run, we were able to obtain regular observations of a number of stars, including known planet-bearing stars 51 Peg, υ And and HD 209458; RV stable stars η Cas, τ Ceti and and 31 Aql; and a bright star, α Boo, over a period of about seven days. The weather was good for the most part, with typical seeing of around 1.7″.

6.1.2 Data Analysis

The raw data were reduced using an early version of the current pipeline, following roughly the same basic procedures as outlined in chapter 4. Raw spectra were first trimmed and dark subtracted using standard IRAF routines, with bias being subtracted along with the darks in one step. Pixel-pixel flatfielding was performed using non-fringing

quartz-lamp continuum spectra as flatfields (where the fringes were eliminated by rapidly oscillating the interferometer PZT mirror during the exposure, as described in section 3.1).

The rest of the data reduction was then performed using custom software written in the IDL data analysis language. The images were 'self illumination corrected' using an algorithm to extract the underlying continuum illumination function which is divided out from each image. This avoids problems with changes in the illumination over time. The spectra were then corrected for slant so that the slit direction was exactly aligned with the CCD pixel axes. They were then low-pass Fourier filtered in order to remove the interferometer comb.

After these pre-processing steps, the phase and visibility were determined for each wavelength channel by fitting a sinusoid to each column of the CCD image, each pixel being weighted according to the number of counts in the original non-flatfielded data, on the assumption of photon noise dominated error. Since fringe spatial frequency varies only slowly as a function of wavelength, we fit a smooth function to the frequencies obtained from the sinusoid fits, and then performed a second pass with the frequencies fixed to match this function, helping to reduce random errors.

Pure stellar and pure iodine template spectra were taken at the beginning of each observation, and these were used to mathematically extract the phase shifts of the star and the iodine individually, and hence calculate the intrinsic stellar velocity shift corrected for instrumental drifts (see section 4.4.3.1).

Finally, the RV due to the motion of the Earth was subtracted to leave an intrinsic stellar relative velocity curve. The exposure time was taken to be the centre of the exposure.

Error bars were based on the standard statistical curve-fitting errors determined during measurement of phase and visibility. The errors were translated to error bars through calculations appropriate to the algorithms used to extract the final intrinsic stellar

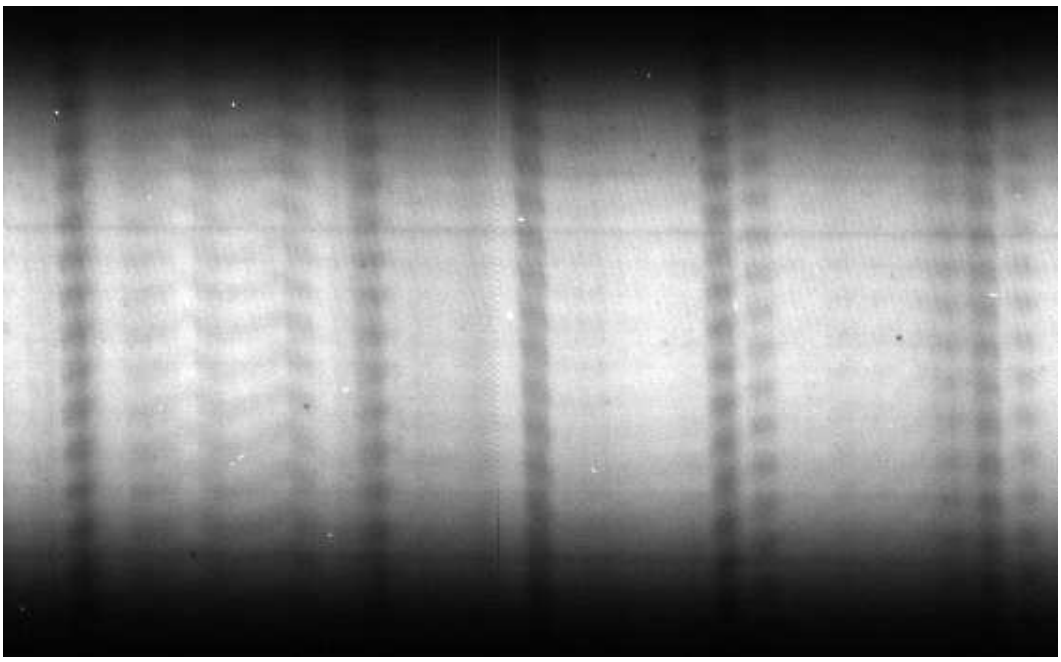

Figure 6-1. Small section of raw fringing spectrum of 51 Peg with iodine, obtained at KPNO on the night of August 14, 2002 ($V = 5.5$ mag, S/N$\sim$ 50 per pixel). (The interferometer comb can be very faintly seen near the top part of the image.) [*Reproduced from van Eyken et al. (2004a).*]

RV (section 5.1). They are expected to give a reasonable guide to the random scatter expected in the data, although they may not catch all systematic errors.

On closer inspection, the data from this trial run were found to show significant variation in the fringe phase and visibility along the length of the slit, perhaps due to non-uniform slit illumination, or to aberration and distortion in the optics. We therefore cut the spectra into three slices along the dispersion direction and treated each slice separately, in order to obtain sinusoidal fits less affected by this systematic error. A weighted average of the three results was used to give the final RV plot.

6.1.3 Results

Part of the raw fringing spectrum for 51 Peg with the overlaid iodine spectrum is shown in figure 6-1, obtained in 25 minutes at visual magnitude 5.5 with S/N per pixel in the central strip of around 50. Typical exposure times for η Cas (mag 3.5) were 30 min at an S/N of 80 (including iodine cell losses).

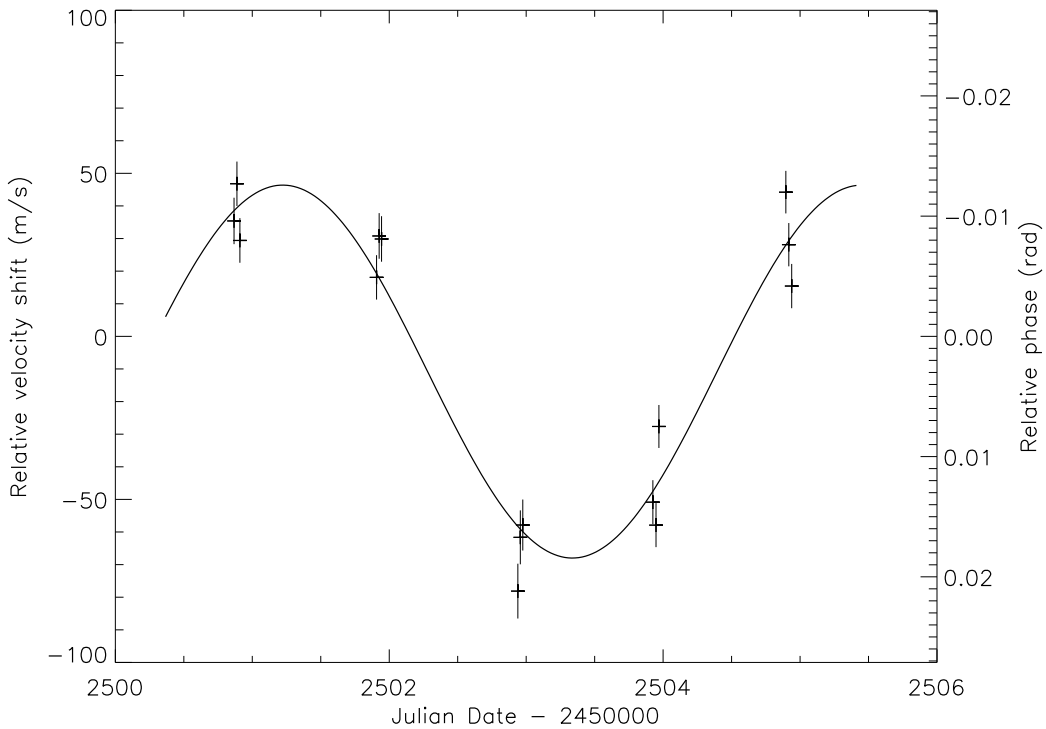

Figure 6-2. Plot of RV measurements for 51 Peg, with the predicted curve over-plotted. RMS residuals are $11.5\,\mathrm{m\,s^{-1}}$. [*Reproduced from van Eyken et al. (2004a).*]

Figure 6-2 shows the radial velocity variation measured for 51 Peg after diurnal motion was subtracted. The zero point was chosen arbitrarily. Over-plotted is the expected curve extrapolated from orbital parameters determined by Naef et al. (2004). The same data are listed in table 6-1. S/N per pixel ratios obtained were in the range ∼40–60 for star+iodine spectra. The templates used for the processing were from the night of August 16, 2002 (August 17 UT), and S/N for the iodine and star templates were approximately 300 and 70 per pixel respectively (for the central of the three strips). Averaging over the three detector strips gave an RMS deviation from the predicted curve of $11.5\,\mathrm{m\,s^{-1}}$. The value of the reduced χ^2 is 2.70.[1]

[1] These results represented a substantial improvement over earlier reported measurements (van Eyken et al. 2003), due in part to using all three detector strips and also to several improvements in the reduction software.

Table 6-1. RV measurements for 51 Peg [*Reproduced from van Eyken et al. (2004a).*]

Julian date −2450000	Radial velocity ($\mathrm{m\,s^{-1}}$)	Error ($\mathrm{m\,s^{-1}}$)
2500.8665	35.4	7.1
2500.8879	46.8	6.8
2500.9094	29.4	6.8
2501.9077	18.1	6.8
2501.9263	30.8	7.0
2501.9441	29.9	7.0
2502.9392	−78.1	8.4
2502.9573	−61.6	8.3
2502.9751	−57.9	7.8
2503.9250	−50.9	6.7
2503.9480	−57.9	6.8
2503.9692	−27.6	6.6
2504.8960	44.2	6.5
2504.9177	28.1	6.6
2504.9392	15.4	6.8

Residuals after diurnal correction for the star η Cas are shown in figure 6-3 and table 6-2, using templates from the night of August 15, 2002 (Aug 16 UT). η Cas is a known RV stable star (W. D. Cochran 2002, private communication) and is therefore expected to show zero shift at our current level of precision. The three image strips are averaged, weighted according to flux. The RMS scatter is $7.9\,\mathrm{m\,s^{-1}}$, with a reduced χ^2 of 2.03. Typical S/N per pixel in the central strip is around 70–90 for star+iodine spectra, 270 for the iodine template, and 100 for the star template.

Under 1.5 arc-sec seeing conditions, we obtained a total instrument throughput of $\sim 4\%$, from above the atmosphere to the detector, including sky, telescope transmission, fibre loss, instrument and iodine cell transmission and detector quantum efficiency. This throughput was obtained using only one interferometer output. Excluding slit loss, the transmission of the instrument itself from fibre to detector was 19%.

6.1.4 Discussion

It is possible to make an estimate of the photon limited error in these measurements following the formalism discussed in section 2.7.1, using equation 2–37. For the 51 Peg

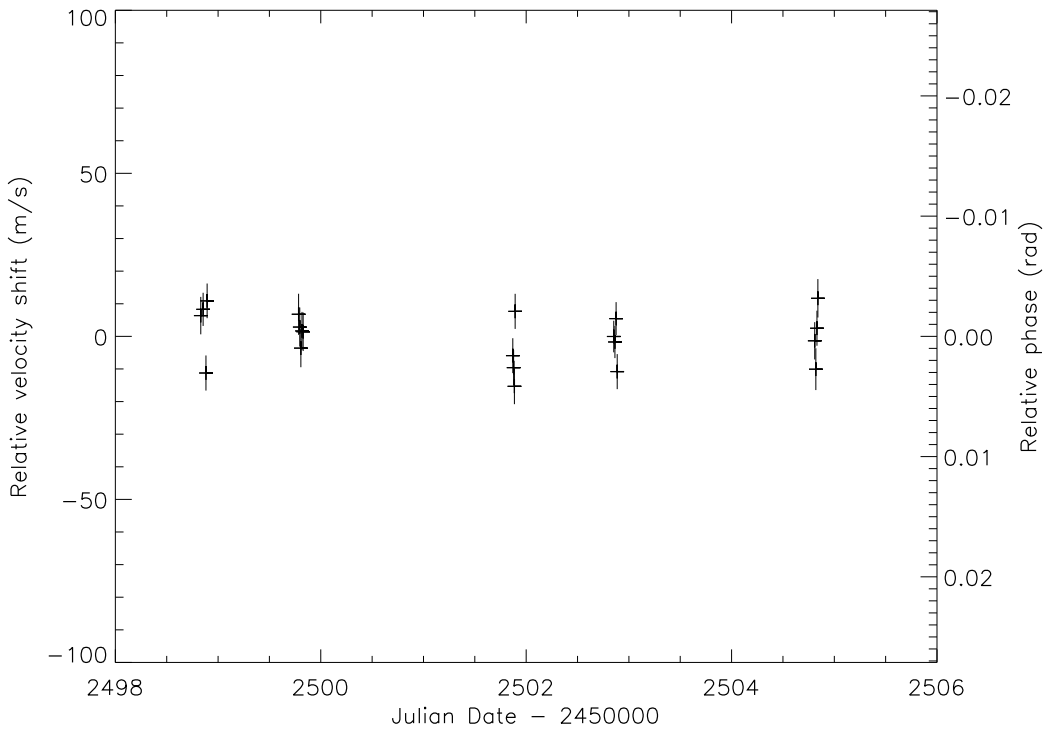

Figure 6-3. Plot of RV measurements for η Cas, an RV stable star, expected to show zero shift. RMS residuals are $7.9\,\mathrm{m\,s^{-1}}$. [*Reproduced from van Eyken et al. (2004a).*]

measurements, averaging over all the data points gives a mean photon error due to the star component of $6.7\,\mathrm{m\,s^{-1}}$, and a mean error due to the iodine component of $7.2\,\mathrm{m\,s^{-1}}$. Doing the same for the η Cas measurements yields mean errors of $5.3\,\mathrm{m\,s^{-1}}$ and $4.3\,\mathrm{m\,s^{-1}}$ for star and iodine components, respectively. Adding the component errors in quadrature gives a total mean photon error of $9.8\,\mathrm{m\,s^{-1}}$ for 51 Peg, and $6.8\,\mathrm{m\,s^{-1}}$ for η Cas, as summarised in table 6-3.[2]

Considering the uncertainty in the RMS residual values obtained for the data due to the small number of data points ($11.5 \pm 2.1\,\mathrm{m\,s^{-1}}$ for 51Peg and $7.9 \pm 1.2\,\mathrm{m\,s^{-1}}$ for η Cas),

[2] The photon errors listed here are a little smaller than those originally reported in van Eyken et al. (2004a), where the error terms due to the templates themselves should have been neglected, as discussed in section 2.7.1. The figures given here are more realistic, but the conclusions remain essentially the same.

Table 6-2. RV measurements for η Cas. [*Reproduced from van Eyken et al. (2004a).*]

Julian date −2450000	Radial velocity ($\mathrm{m\,s^{-1}}$)	Error ($\mathrm{m\,s^{-1}}$)
2498.8318	6.4	5.7
2498.8545	8.3	5.1
2498.8823	−11.2	5.4
2498.8936	10.9	5.3
2499.7839	6.8	6.3
2499.7952	2.9	6.0
2499.8064	−3.6	5.8
2499.8174	1.7	5.8
2499.8284	1.4	5.9
2501.8701	−5.9	5.3
2501.8777	−9.6	5.4
2501.8850	−15.3	5.5
2501.8926	7.7	5.4
2502.8528	−0.0	4.9
2502.8643	−1.7	4.9
2502.8752	5.4	5.1
2502.8860	−10.8	5.4
2504.8093	−1.4	5.7
2504.8201	−10.1	6.4
2504.8311	2.5	5.4
2504.8401	11.7	5.9

Table 6-3. Mean photon limiting error estimation for 51 Peg and η Cas observations, compared with RMS residuals from the data.

Target	Star component ($\mathrm{m\,s^{-1}}$)	Iodine component ($\mathrm{m\,s^{-1}}$)	Total photon err. ($\mathrm{m\,s^{-1}}$)	RMS residuals ($\mathrm{m\,s^{-1}}$)
51 Peg	6.7	7.2	9.8	11.5
η Cas	5.3	4.3	6.8	7.9

we can see that the instrument was already performing close to the photon limit: the reduction software was successfully extracting almost the maximum possible information from the data. For comparison, RMS scatters obtained previously for 51 Peg had been $13\,\mathrm{m\,s^{-1}}$(Mayor & Queloz 1995), $5.2\,\mathrm{m\,s^{-1}}$ (Marcy et al. 1997), and $11.8\,\mathrm{m\,s^{-1}}$ (Naef et al. 2004).

Given the photon limit, the error bars in the data appear to be somewhat underestimated (leading to the large values for the reduced χ^2). A possible cause of this was the low pass

Fourier filtering done to remove the interferometer comb. In addition to removing the comb, filtering has the effect of smoothing the photon noise in the data, reducing the residuals in the sinusoid fits to the fringes and thereby reducing the resulting error estimates for each fringe. This reduction may be an artificial effect, however, which in fact does not improve the precision of the fits. The extent to which artifical reduction of the standard errors occurs has still to be determined. With later data, however, it was usually found more instructive to actually use the calculated photon limit for each data point to create the error bars, and this has generally been the convention adopted since then.

The entire run was conducted over the period of about two weeks, including set-up. The RMS residuals we obtained, along with the first planet confirmation, clearly showed that the precision we were able to attain with the DFDI technique was becoming comparable with traditional echelle techniques.

6.2 Discovery of ET-1

The years 2003 and early 2004 saw the permanent installation of ET at the 2.1 m telescope, with major upgrades (see section 3.2.2). Along with continuing refinements to the software, the upgrades enabled us to make substantially improved measurements of 51 Peg in December 2004, showing RMS residuals now as low as $7.9\,\mathrm{m\,s^{-1}}$ over 6 data points and an average photon limiting error of $9.4\,\mathrm{m\,s^{-1}}$ using only the 0.9 m coudé telescope, rather than the 2.1 m (see figure 6-4). Earlier, in March 2004, we had been able to obtain a short-term RMS of $3.6\,\mathrm{m\,s^{-1}}$ over $\sim$ 2 hr on another known stable reference star, 36 UMa (visual magnitude 4.83), using the 2.1 m telescope, with mean photon limiting errors of $4.0\,\mathrm{m\,s^{-1}}$ (see figure 6-5). In both these cases, the RMS was again consistent with the photon limit, given the uncertainty in the RMS due to the small number of data points. The error bars represent the calculated photon limiting error in all of the following results in this section.

Beginning in the winter of 2004, we began a small-scale survey with the instrument to search for short period (< 10 d) planets, and we were eventually able to uncover our first

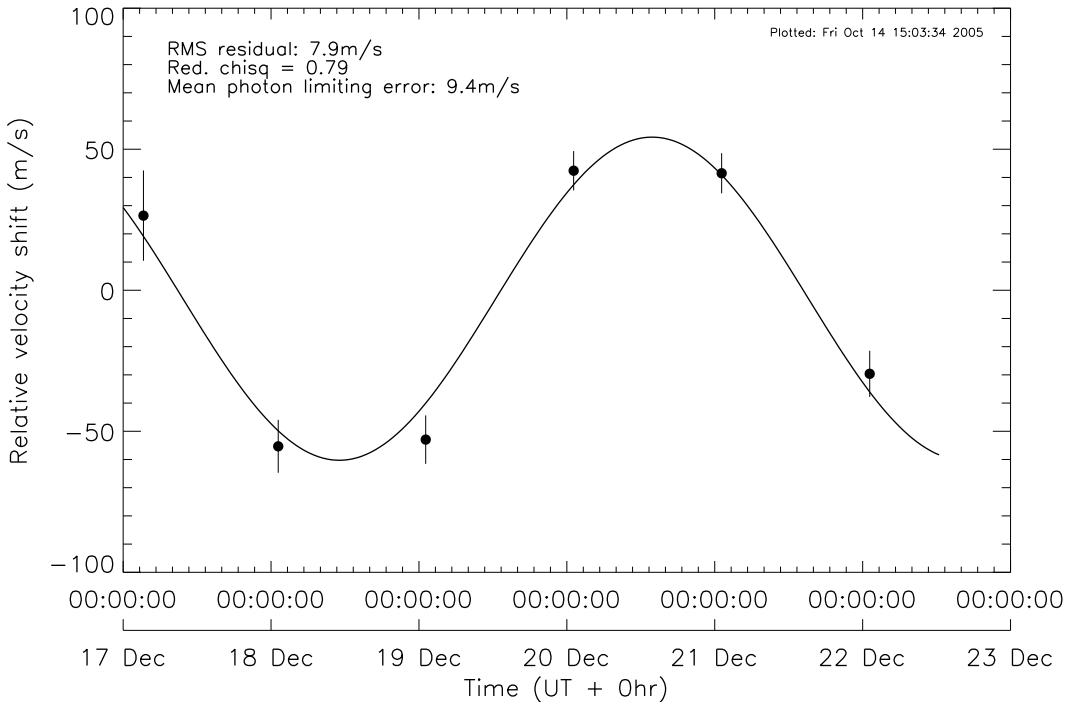

Figure 6-4. 51 Peg RV measurements in December 2004 with the upgraded ET, with predicted RV curve overplotted. 7.9 m s^{-1} RMS residuals were obtained with the 0.9 m coudé telescope – an improvement over previous the 11.5 m s^{-1} obtained with the 2.1 m in 2002 (figure 6-2).

new planet, ET-1, orbiting the star HD 102195. The initial tentative findings, discussed here, were reported in van Eyken et al. (2005); the official discovery was announced in Ge et al. (2006a), *representing the first time an exoplanet has been discovered by RV measurements around a star fainter than magnitude V=8 using a sub-meter sized telescope.*

6.2.1 Early Measurements

6.2.1.1 Observations and data analysis

Observations were conducted of around 90 previously unsearched stars during the period from December 2004–May 2005 using the KPNO 0.9m coudé feed telescope from December to March and the 2.1 m telescope in May. We chose most targets primarily from the N-star catalogue (Gray et al. 2003), selecting dwarf stars of type F–K with visual magnitude V=8–9, high metallicity, and where known, slow rotation and low activity

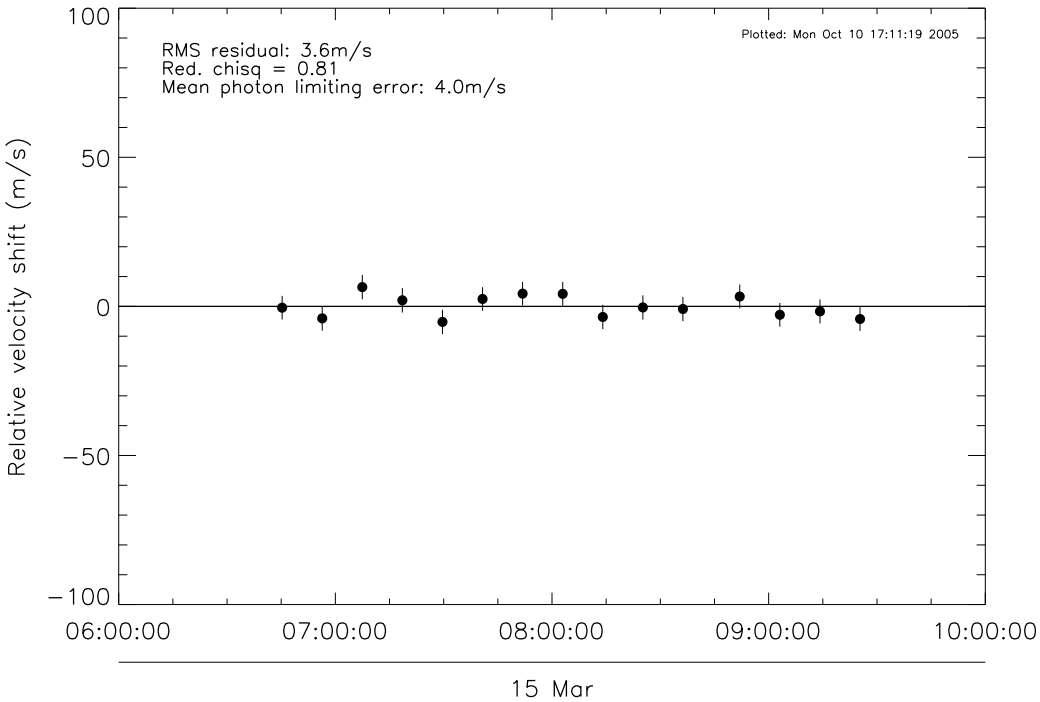

Figure 6-5. 36 UMa (known RV stable star) short term precision measurements with the upgraded ET in March 2004, showing photon-limiting performance at $3.6\,\mathrm{m\,s^{-1}}$ RMS residuals.

indicators (target selection is discussed in more detail in Mahadevan 2006). Any known visual doubles or variable stars were rejected from the list. Where possible, candidates showing significant RV variation early in the survey were followed up using the 2.1 m telescope in May 2005. The survey was divided into five observing runs of durations varying from 8–21 nights, the first four runs using the coudé feed telescope, and the last switching to the 2.1 m part way through the run. A total of 59 nights were allocated for the coudé feed, and 7 nights on the 2.1 m. Observations were made on 43 of the nights, the remainder being lost primarily due to bad weather.

Star and iodine templates were taken at the beginning and end of each observing run, and typically 5–6 RV measurements with the iodine cell in the beam path were acquired per survey star for each run. Typical exposure times for the coudé feed were ~ 25 min for stars at magnitude $V \sim 8$ and ~ 40 min for $V \sim 9$. On the 2.1 m, exposures were generally kept to 10 min. A typical raw spectrum from the current upgraded instrument is shown

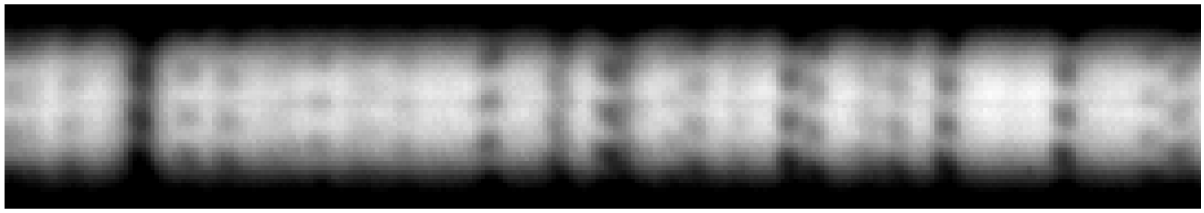

Figure 6-6. Example section of raw spectrum taken with upgraded ET (55 Cnc without iodine, May 2005).

in figure 6-6. The data were reduced using essentially the same pipeline as is now used, described in chapter 4, although the process was then not as streamlined, and lacked the fully automated multi-object multiplexing capability.

6.2.1.2 RV results

The initial report of tentative results was given in October 2005 (van Eyken et al. 2005). In order to establish an independent estimate of our survey precision, we defined stars which appeared RV stable at our precision level to be those showing a reduced $\chi^2 < 2$ over the period of observation. At the 0.9m coudé during one typical observing block, 7 out of 15 search stars satisfied this criterion. For those, we found a mean RMS scatter of $18.9\,\mathrm{m\,s^{-1}}$ compared to a mean photon limiting error of $21.2\,\mathrm{m\,s^{-1}}$, at a mean visual magnitude of 8.26 and typical exposure times of 20-30min. (The lower value of the RMS is comfortably within the range expected due to statistical variation). From the 2.1 m observations, using the same criterion, we found 12 stable stars out of 24 search stars, for which we found a mean RMS of $17.6\,\mathrm{m\,s^{-1}}$, compared to a mean photon limiting error of $17.4\,\mathrm{m\,s^{-1}}$. In this case, the mean V magnitude was 8.48, and typical exposures were 10 min. In both cases the errors were consistent with the expected photon limit.

Of the 90 stars surveyed, we selected those that showed an RMS deviation about a constant RV greater than 2.5 times the mean photon error as possible candidates. We ruled out any that showed variation greater than $1000\,\mathrm{m\,s^{-1}}$ on the assumption that these were likely to be binary stars. 10 candidates remained, and HD 102195 appeared to be the most promising.

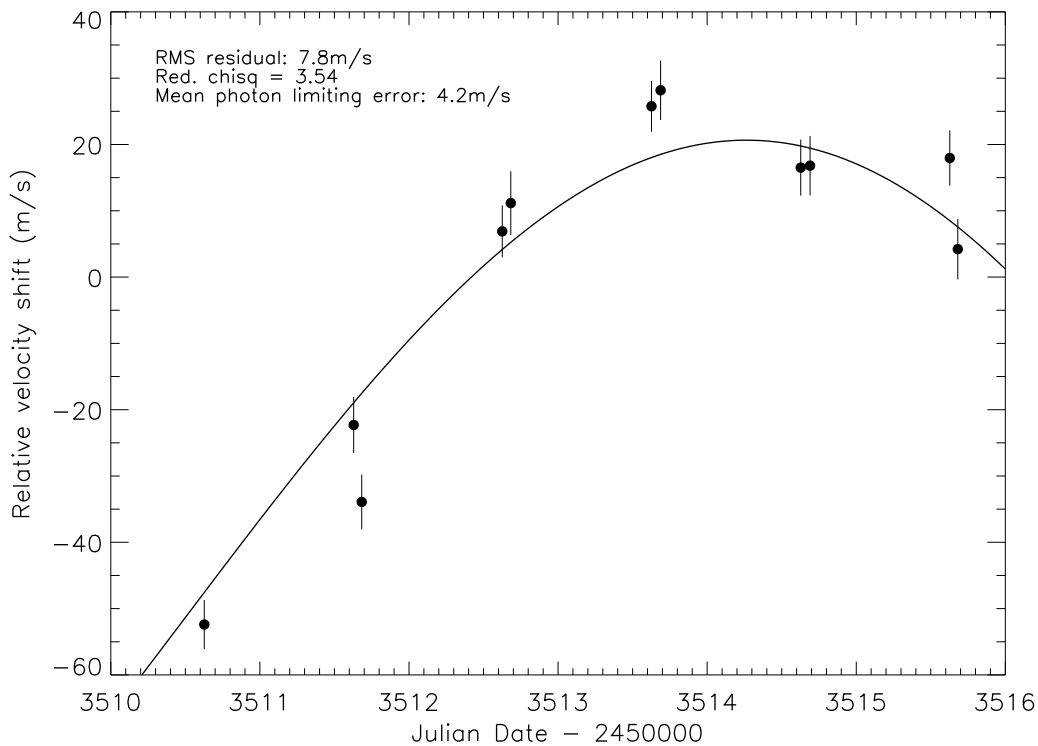

Figure 6-7. Known planet-bearing star, 55 Cnc, measured with current KPNO ET. Predicted RV curve due to 55 Cnc b is overplotted. Additional scatter may be due to a closer-in planet not included in model.

Figure 6-7 shows results then obtained at the 2.1 m for the known planet bearing star, 55 Cnc, used as a control, with the predicted RV curve due to its 14.7 d companion 55 Cnc b over-plotted. The other known companions (Butler et al. 2006) are neglected here, since they have much longer periods and would have very little effect over this timescale, with one exception: part of the reason for the somewhat large reduced χ^2 value may well be because of an inner planet, 55 Cnc e, reported by McArthur et al. (2004) to have a 2.8 d period and velocity semi-amplitude of 6.7 m s^{-1}. (Modelling this planet into the predicted curve is difficult because it is likely to be gravitationally interacting with 55 Cnc b.) RMS deviation about the predicted curve was 7.8 m s^{-1}, compared to a mean photon limiting error of 4.2 m s^{-1}. Figure 6-8 similarly shows the performance over the period of a few days again using the bright RV-stable reference star 36 UMa.

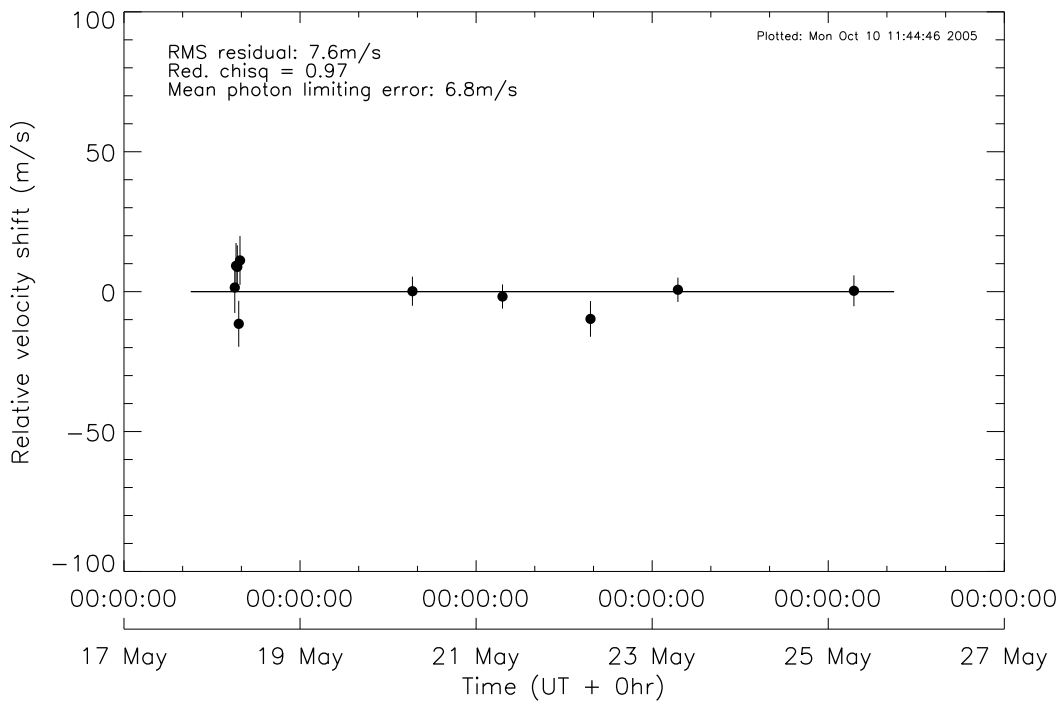

Figure 6-8. RV stable star, 36 UMa, measured over a few days with current KPNO ET in 2005. Smallest error bar is $\sim 4.1\,\mathrm{m\,s^{-1}}$.

Figures 6-9 and 6-10 show a preliminary best-fit Keplerian orbit curve to the first sets of data for HD 102195, showing the planetary companion RV signal. Data for this candidate were obtained in three sets of observations, each of about one week, in January, March (0.9m) and May 2005 (2.1 m). The offsets between the data sets were arbitrary, each observing run being reduced separately: the offsets were therefore chosen to give a good match to a single Keplerian signal with no long-term trend.

A zero-eccentricity RV fit at the most likely period given the data, 4.85 d, was consistent with a planetary companion of minimum mass $M \sin i = 0.41\,M_{\mathrm{J}}$ orbiting at 0.052 AU from the host star, with an RV semi-amplitude of $58.1\,\mathrm{m\,s^{-1}}$, RMS residuals of $20.2\,\mathrm{m\,s^{-1}}$, and reduced χ^2 of 1.49 (where the error bars shown are again the photon limiting errors). These values gave a mass limit of $M \sin i = 0.41\,M_{\mathrm{J}}$, consistent with other known hot-Jupiter exoplanets in the same period range.

Figure 6-11 shows a Lomb-Scargle periodogram (Lomb 1976; Scargle 1982) for the data then taken to that point, showing the power at various periods in the RV

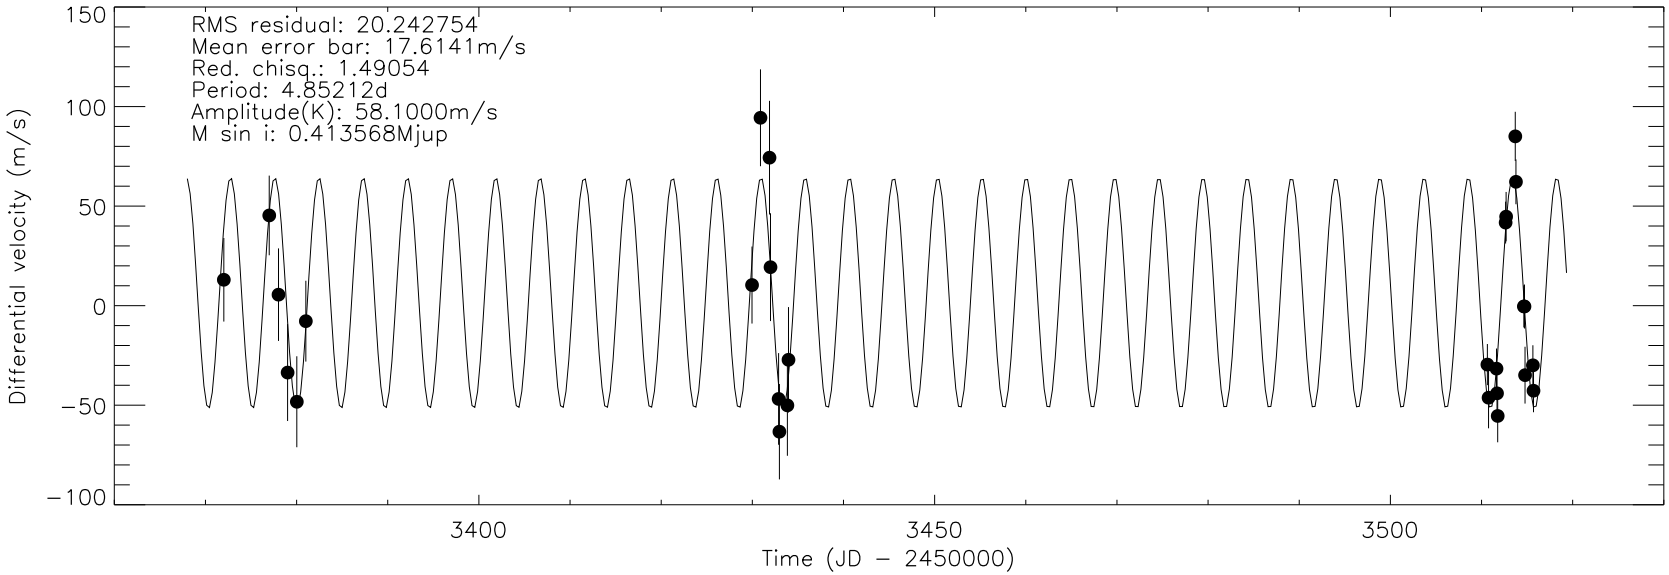

Figure 6-9. Best fit Keplerian orbit for early ET-1 RV measurements, assuming a zero eccentricity planetary companion. Offsets between the three observing runs were arbitrary, and chosen assuming no long term RV drift between the runs.

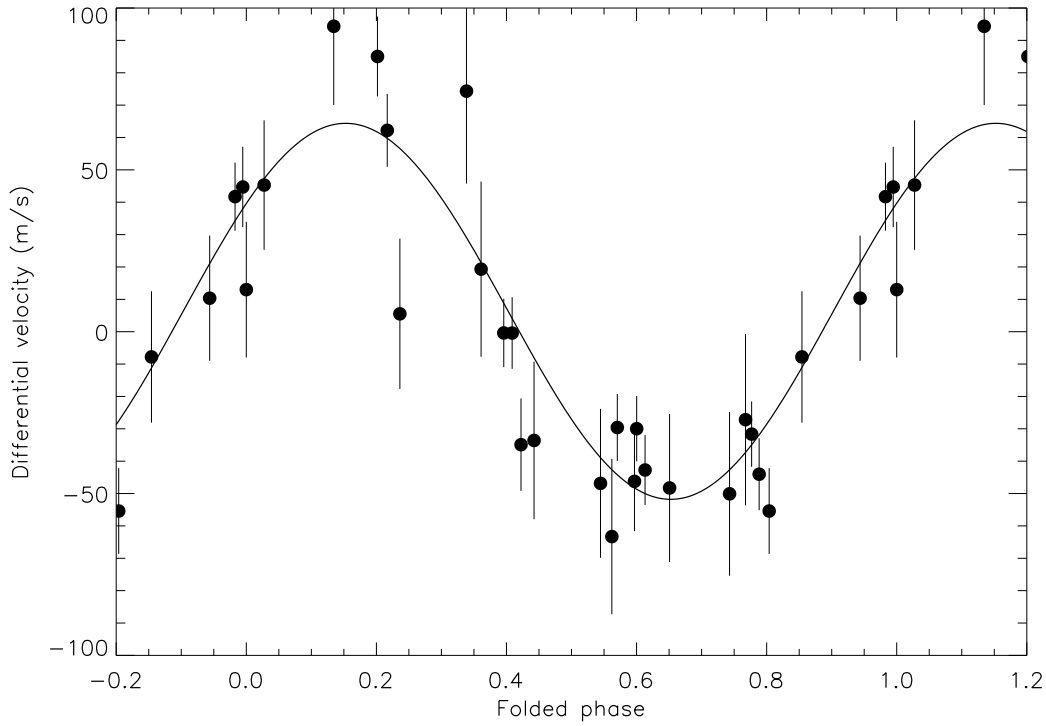

Figure 6-10. As for figure 6-9, but phase folded on the best-fit period of 4.85 d.

signal. The most significant periodicity of 4.85 d was used as an initial guess for the χ^2 minimisation in the curve fitting. It can be seen, however, that the sampling of the data gave a number of other periods of similar likelihood. More data were needed to narrow down the possibilities: later followup data (section 6.2.2) were to show that in fact the peak at around 4.1 d was the correct one.

6.2.1.3 Photometry

It is important to rule out possible non-planetary sources of RV variation when searching for planets. One effect that can be very significant is the presence of stellar variability – in particular, variability caused by star spots. Any time-varying inhomogeneity in the stellar disc can appear as an RV variation. This is because the RV measured is actually effectively an averaged RV over the surface of the disc. Since our target stars typically have projected rotational velocities on the order of several $\mathrm{km\,s^{-1}}$ and we are looking for signals at only the $\mathrm{m\,s^{-1}}$ level, a dark spot on the surface of the stellar disc can change the measured RV by a significant amount, potentially at least

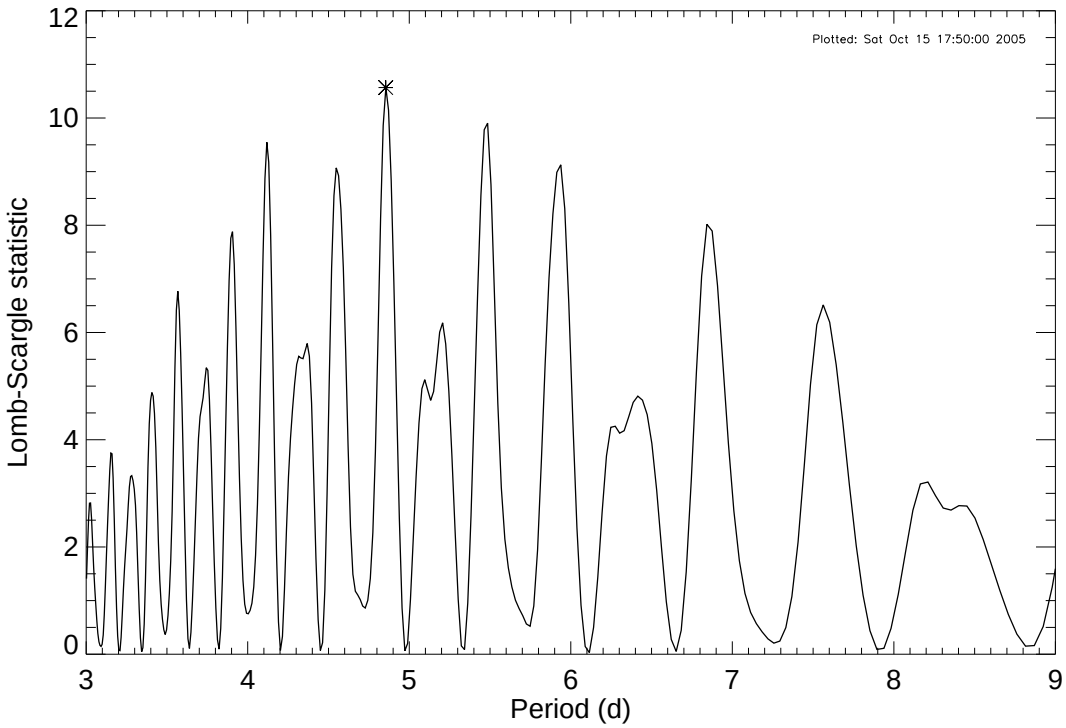

Figure 6-11. Lomb-Scargle periodogram for early ET-1 data. The most probable period is marked with an asterisk at 4.85 d, showing a signal above the background at the 99.7% significance level.

comparable to a planetary signature. As stellar rotation brings the spot in and out of view, covering different parts of the disc, the RV signal can be modulated. The effect can be distinguished from true Doppler modulation, however, because in the case of surface features, the apparent magnitude of the star will vary in synchrony with the RV signal. It is therefore fairly standard procedure to take photometric measurements of suspected planet bearing stars to rule out such problems.

Figure 6-12 shows preliminary photometric measurements of the same target taken a few weeks later by Gregory Henry at Tennessee State University, using the Automatic Photometry Telescope at Fairborn Observatory. Some variation was seen at the 0.01 mag level, with somewhat of a periodicity around roughly 12 days, showing that there was indeed some stellar variation. The 12-day period, however, would likely correspond to the rotational period of the star, and was far enough removed from the period of the RV curve not to rule out the possibility of a planetary companion.

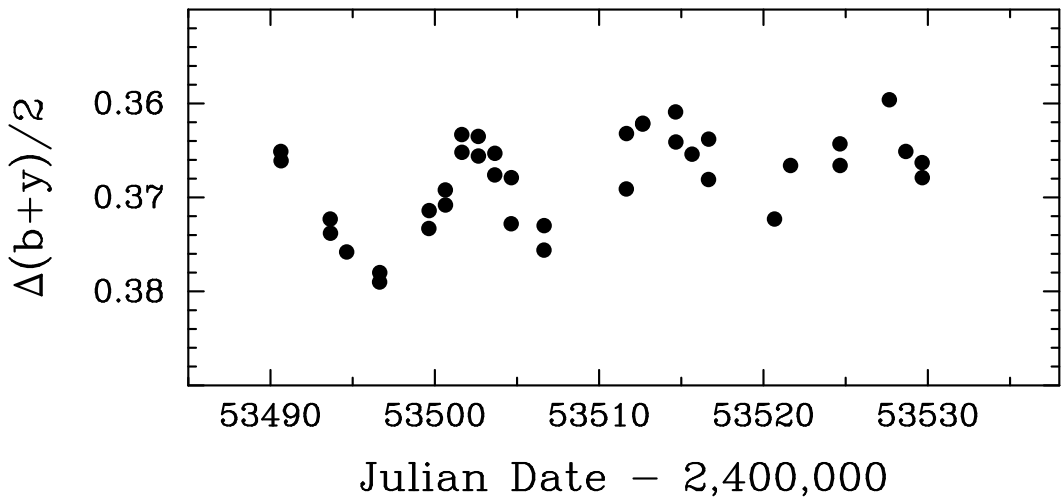

Figure 6-12. Early photometry of HD 102195, the parent star for ET-1 (combined Strömgren b and y differential magnitudes). Taken by Gregory Henry using the APT at Fairborn Observatory. A $\sim 12\,\mathrm{d}$ periodicity is evident, at the $0.01\,\mathrm{mag}$ level. Error bars are $\sim 0.001\,\mathrm{mag}$. [*Reproduced from Ge et al. (2006a).*]

6.2.2 Followup and Confirmation

Later followup was to confirm the planetary nature of the RV signals detected, and the planet, HD 102195b, became referred to as ET-1. In addition to further RV measurements made using the 2.1 m telescope, much of the followup work was undertaken by other collaborators. The results are reported in detail in Ge et al. (2006a), but are briefly summarised here for completeness.

In December 2005, a further 21 data points were made using ET and the 2.1 m telescope to try and weed out the harmonics in the periodogram and pin down the correct period. From November 2005 – January 2006, a further 10 RV measurements were independently made with conventional RV methods using the McDonald Observatory 9 m Hobby-Eberly telescope (HET), in an effort led by Donald Schneider at Penn State University. Using the High Resolution Spectrograph (HRS), the much larger telescope aperture allowed for internal errors of only $\sim 2\,\mathrm{m\,s^{-1}}$, and the measurements convincingly confirmed the RV signal detected with ET. Combining the HET measurements with the ET measurements allowed for a much more precise constraint on the orbital parameters

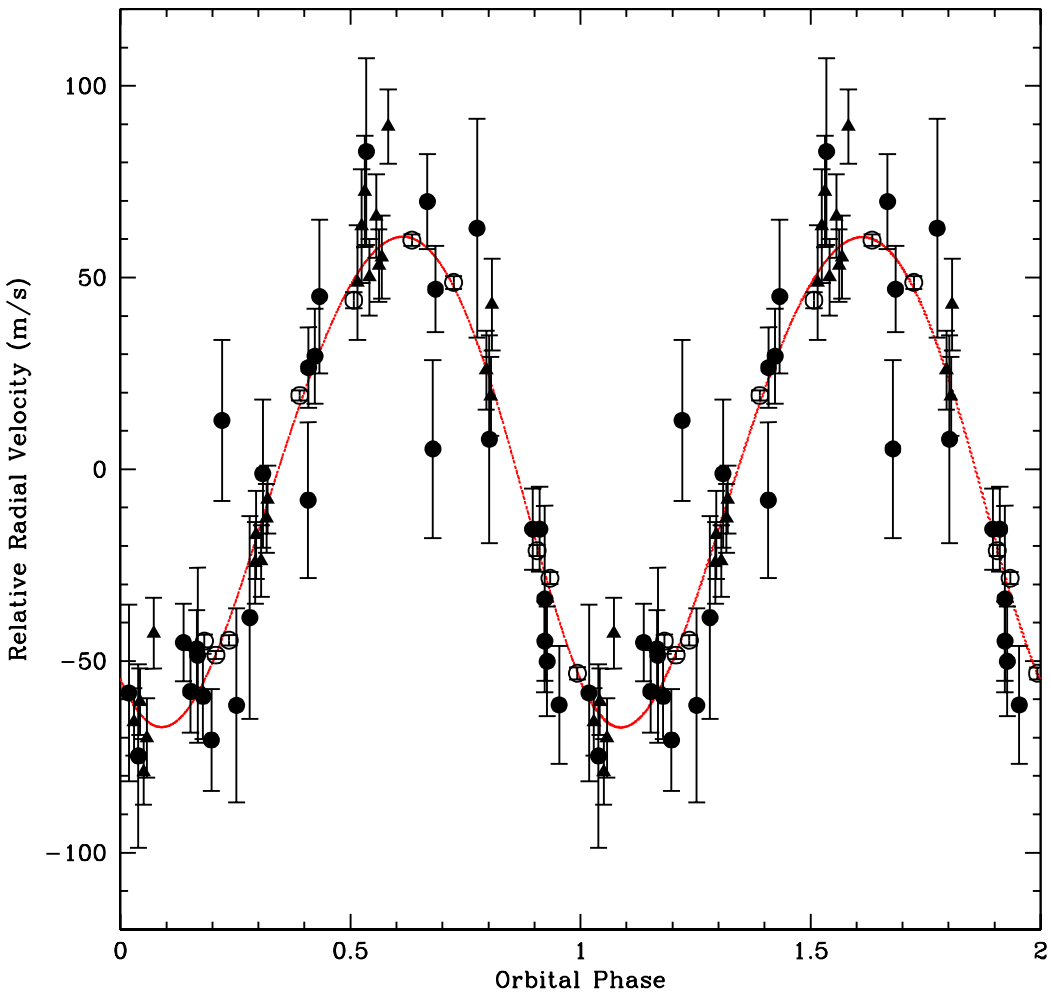

Figure 6-13. Folded combined radial velocities for HD 102195, combining ET measurements from the KPNO Coudé Feed telescope (filled circles), 2.1 m (filled triangles), and HRS measurements from the HET (open circles). Two orbital cycles are shown and each observation is plotted as two points, folded on the best-fit orbital period. The best Keplerian fit to the data (overplotted) yields an orbital period of 4.11434 d. [*Plot produced by Robert Wittenmyer, University of Texas; reproduced from Ge et al. (2006a).*]

of ET-1. The combined RV data sets are plotted in figure 6-13, along with the best-fit Keplerian orbit solution calculated by Robert Wittenmyer at the university of Texas.[3]

[3] Further HET data were also later proposed for and analysed by Suvrath Mahadevan: those results are reported in Mahadevan (2006).

Table 6-4. Orbital parameters for HD 102195.

Parameter	Value
P	4.11434±0.00089 days
T_p	2453732.7±0.5
e	<0.14
a	0.0491 AU
ω (deg)	143.4±15.4
K	63.4±2.0 m s^{-1}
$m \sin i$	$0.488 \pm 0.015\ M_J\,(\frac{M_*}{0.93\,M_\odot})^{2/3}$
σ_j	5.8±1.8 m s^{-1}
RMS	16.0 m s^{-1}

Symbols represent standard orbital parameters; σ_j represents instrinsic stellar jitter. [*Reproduced from Ge et al. (2006a); an error in the dependence of the minimum planet mass on the stellar mass has been corrected here.*]

The periodogram was updated by Stephen Kane using the full data set, and is shown in figure 6-14 for comparison with figure 6-11. The 4.11 d peak is now clearly the strongest. Eric Ford, then at Berkeley, performed a full Bayesian analysis of the data to obtain the best possible determination of the orbital parameters, and found fully consistent results, attributing the previously used 4.8 d period to aliasing by the lunar cycle (since observations were generally taken only during bright time, around the time of full-moon). As before, the offsets between the different runs and instruments were arbitrary, and so were allowed to float as free parameters in the analysis. The orbital parameters for the best solution are listed in table 6-4. The full set of radial velocity measurements from all the observing runs is listed in table 6-5.

A measurement of the stellar parameters for HD 102195, listed in table 6-6, was lead by Eduardo Martín at the Instituto de Astrofisica de Canarias, using the high resolution SARG spectrograph on the 3.5 m Telescopio Nazionale Galileo at La Palma in June 2005. This measurement was also used to check for the presence of any faint stellar spectroscopic binary companion which might be offsetting the line centroid positions and masquerading as a Doppler shift; none was found.

The SARG spectra were also used by Martín's team to perform a line bisector analysis. A 'line bisector' traces the centre of an absorption line as a function of depth

Table 6-5. Complete radial velocities for HD 102195

JD −2450000	Instrument[a]	RV m s^{-1}	Errors m s^{-1}	JD −2450000	Instrument[a]	RV m s^{-1}	Errors m s^{-1}
3372.014	CF, BK2	133.0	21.0	3696.034	HET	−60.6	1.3
3376.999	CF, BK2	165.3	20.0	3697.034	HET	−52.0	1.1
3378.011	CF, BK2	125.5	23.2	3701.029	HET	−55.9	0.8
3379.013	CF, BK2	86.4	24.3	3704.016	HET	−35.8	1.3
3380.024	CF, BK2	71.7	22.8	3718.966	2.1 m	56.0	10.9
3381.011	CF, BK2	112.2	20.3	3719.023	2.1 m	90.0	9.7
3429.966	CF, BK4	−99.6	19.3	3720.978	2.1 m	−69.5	10.3
3430.892	CF, BK4	−15.6	24.3	3721.039	2.1 m	−42.1	9.2
3431.882	CF, BK4	−35.7	28.5	3721.956	2.1 m	−16.5	11.5
3431.991	CF, BK4	−90.7	27.1	3722.001	2.1 m	−23.3	9.3
3432.881	CF, BK4	−156.9	23.0	3722.048	2.1 m	−12.2	9.0
3432.966	CF, BK4	−173.3	24.0	3722.059	2.1 m	−7.2	8.9
3433.844	CF, BK4	−160.1	25.3	3722.967	2.1 m	50.7	10.0
3433.963	CF, BK4	−137.2	26.5	3723.029	2.1 m	66.6	10.9
3510.641	2.1 m, BK5	−29.6	10.3	3723.053	2.1 m	53.7	9.5
3510.769	2.1 m, BK5	−46.2	15.4	3724.012	2.1 m	26.5	10.3
3511.643	2.1 m, BK5	−31.6	10.1	3724.053	2.1 m	19.6	10.3
3511.699	2.1 m, BK5	−44.0	11.1	3724.064	2.1 m	43.6	11.9
3511.774	2.1 m, BK5	−55.4	13.2	3724.975	2.1 m	−65.3	8.8
3512.641	2.1 m, BK5	41.7	10.5	3725.029	2.1 m	−60.0	8.7
3512.699	2.1 m, BK5	44.7	12.4	3725.062	2.1 m	−78.3	8.6
3513.704	2.1 m, BK5	85.0	12.3	3726.062	2.1 m	−23.8	10.6
3513.777	2.1 m, BK5	62.2	11.2	3726.975	2.1 m	49.3	14.9
3514.646	2.1 m, BK5	−0.4	10.6	3727.012	2.1 m	64.1	14.8
3514.708	2.1 m, BK5	−0.4	11.0	3727.040	2.1 m	73.1	14.5
3514.775	2.1 m, BK5	−34.9	14.3	3731.949	HET	41.3	1.4
3515.638	2.1 m, BK5	−30.0	10.1	3737.946	HET	−52.2	1.3
3515.700	2.1 m, BK5	−42.7	10.8	3740.922	HET	−28.6	1.1
3694.035	HET	36.7	1.8	3742.912	HET	11.9	1.1
				3743.915	HET	52.3	1.1

[a] 'CF' = 0.9 m coudé feed; 'BK' indicates which block of observations the measurement is from. [*Reproduced from Ge et al. (2006a)*]

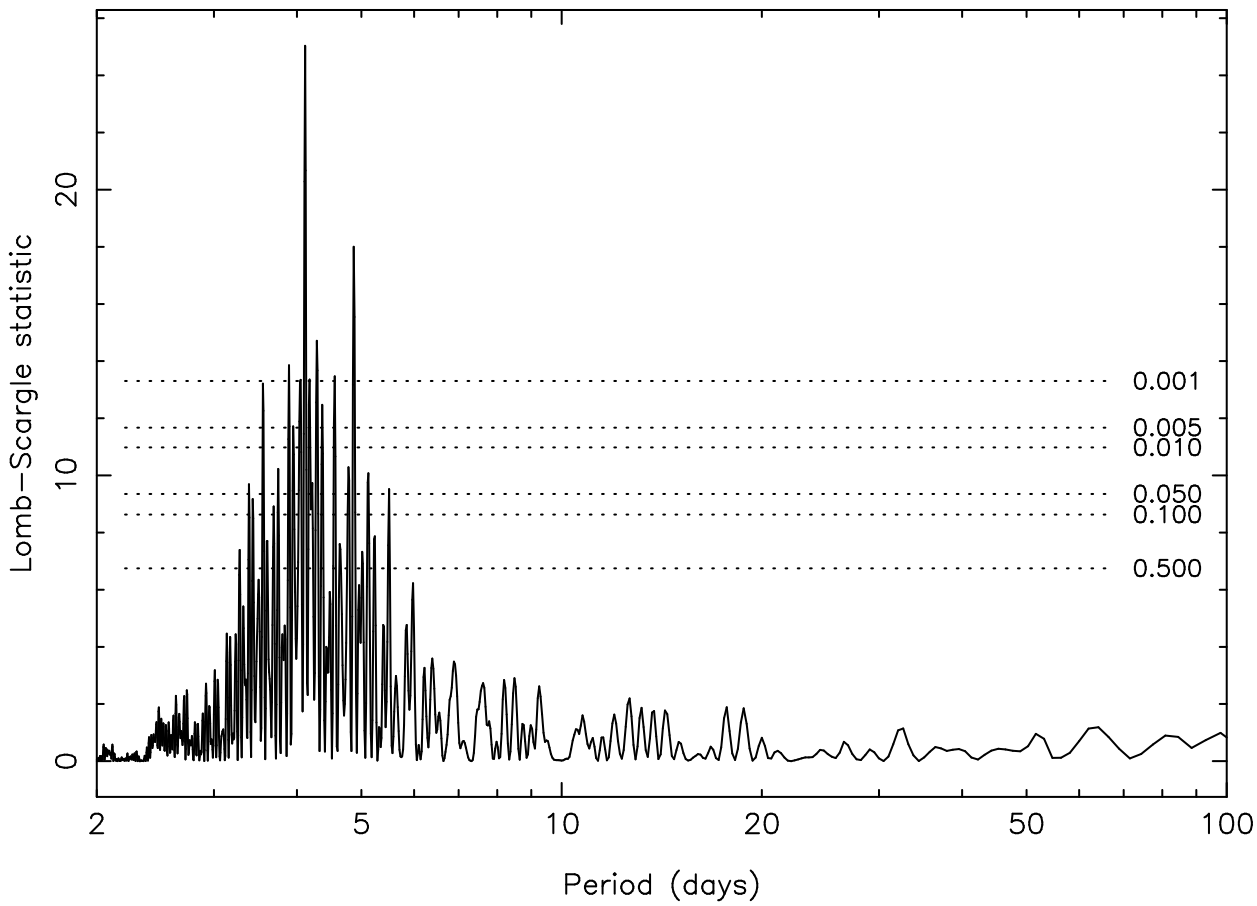

Figure 6-14. Updated periodogram for HD 102195, by Stephen Kane, using the full combined data sets. Different false alarm probabilities are indicated by the dashed lines. The peak power is now very clear at 4.11 d, with a false alarm probability of $\sim 10^{-6}$. [*Reproduced from Ge et al. (2006a).*]

in the line, giving an indication of variation in the symmetry of the line. Analysis of the variation of the line bisector can give information about stellar variability, particularly the presence of star spots or changes in the granulation pattern, which can create line asymmetries that look like RV variations (Martínez Fiorenzano et al. 2005). Over the 3 days of SARG observations, no such bisector variation was found.

Further high-precision photometric analysis by Gregory Henry up until February 2006 showed the same periodicity as before at 12.3 ± 0.3 d, but now at much lower amplitude. This behaviour is consistent with slow starspot or plage variation on a star rotating at a 12.3 day period. The period is also consistent with the projected rotational velocity

Table 6-6. Stellar parameters for HD 102195

Parameter	Value
V	8.05
M_V	5.73
$B-V$	0.84
Spectral type	G8V
Distance	29 pc
[Fe/H]	0.096 ± 0.032
$T_{\rm eff}$	5330 ± 28 K
$v \sin i$	3.23±0.07 km s^{-1}
$\log g$	4.368±0.038 [log(cm s^{-2})]
BC	−0.177
$L_{\rm star}$	0.463±0.034 L$_\odot$
$M_{\rm star}$	0.926±0.016 $M_\odot$
$R_{\rm star}$	0.835±0.016 $R_\odot$
log R'_{HK}	−4.30
$P_{\rm rot}$	12.3±0.2 days
Age	0.6–4.2 Gyr

[*Reproduced from Ge et al. (2006a)*]

($v \sin i$) with the SARG measurements if the planet's orbit is close to edge-on, although the photometry also ruled out any transit signatures to a high degree of confidence.

Finally, measurements and analysis of emission line cores at the bottom of the CaII H and K lines, created in the stellar chromosphere, were undertaken by Suvrath Mahadevan (Mahadevan 2006), using the KPNO 0.9 m coudé spectrograph. These can be used as an indicator of stellar activity, and the results were consistent with the interpretation of HD 102195 as a mildly active star. Similar measurements were made using a spectrum obtained with the FOCES high resolution echelle spectrograph at the German-Spanish Astronomical Observatory (CAHA, Almería, Spain), led by David Montes at the Universidad Complutense de Madrid, leading to similar conclusions. Variation in the chromospheric activity measured by Mahadevan showed no signal at the orbital period of the planet that would cause concern for the planetary interpretation, although some variation was seen that could have been correlated with the 12.3 d stellar rotation period.

6.2.3 Discussion

The full set of observations of HD 102195 taken together made a convincing case for a hot-Jupiter type planetary companion orbiting a mildly active star. Although the level of stellar activity could be expected to cause some RV variation, the activity measurements implied that the variation would be on the level of only $\sim$ 10–20 $\mathrm{m\,s^{-1}}$; even had the period of the variation matched the 4.11 d planetary orbital period, it would have been insufficient to explain the measured 63 $\mathrm{m\,s^{-1}}$ RV semi-amplitude.

The only concern was that the 12.3 d stellar rotational period appeared to be quite precisely 3 times the planetary orbital period. The only physical explanation we were able to think of for this was a 3-fold symmetric positioning of star spots on the stellar surface, at 120° spacings in stellar longitude. Given the complete lack of power at 12.3 d in RV periodogram, however, and the similar lack of power at 4.11 d in the stellar photometric periodogram, this seemed very unlikely, and we proposed that the planetary interpretation was indeed correct.

6.3 Addendum: 51 Peg b, ET-1, and the Addition Approximation

The observations of 51 Peg b and ET-1 were, in retrospect, somewhat fortuitous, in that they were not much affected by the addition approximation problem discussed in sections 2.6.1 and 5.10, which we had not at the time appreciated. The addition approximation appears as a non-linearity in the RV response that sets in when the scale of RV variation becomes larger than of order a few $\mathrm{km\,s^{-1}}$. Since exoplanetary RV signatures are on much smaller scales than this, it is primarily the motion of the Earth relative to the Solar System barycentre (which can be as large as 60 $\mathrm{km\,s^{-1}}$) that determines when the addition approximation becomes a problem. The positions on the sky and times of observation of 51 Peg and HD 102195 were such that the differential barycentric corrections, shown in figures figures 6-15 and 6-16, were relatively small over the lengths of the runs.

In general the largest variation between observations for each run was on the order of $2\,\mathrm{km\,s^{-1}}$ or less. Although this is large enough to cause some systematic effect, it would appear as a slow, approximately linear drift within a run, of at most of order ~ 20–$30\,\mathrm{m\,s^{-1}}$. This effect would easily have been interpreted as error in the phase-velocity scale: to correct a linear trend of $30\,\mathrm{m\,s^{-1}}$ over a signal of $2000\,\mathrm{m\,s^{-1}}$ due to barycentric motion requires an adjustment in Γ of 1.5%, similar to apparent variations that were indeed seen from run to run. The trend could therefore have been at least partially absorbed in calibrating Γ. Since the velocity semi-amplitude of both 51 Peg and HD 102195 was around $60\,\mathrm{m\,s^{-1}}$, a 1.5% error in Γ would only amount to differential errors of around $1\,\mathrm{m\,s^{-1}}$, which is insignificant compared to the photon errors in those measurements.

From figure 6-16 we can see, however, that in attempting to join runs together, the iodine approximation would contribute very significant random run-to-run offset errors, on the scale of many tens of meters (see section 5.10). This error source therefore likely explains why attempts to link the sets of observations between the various ET runs were largely unsuccessful, forcing us to allow floating offsets between each run.

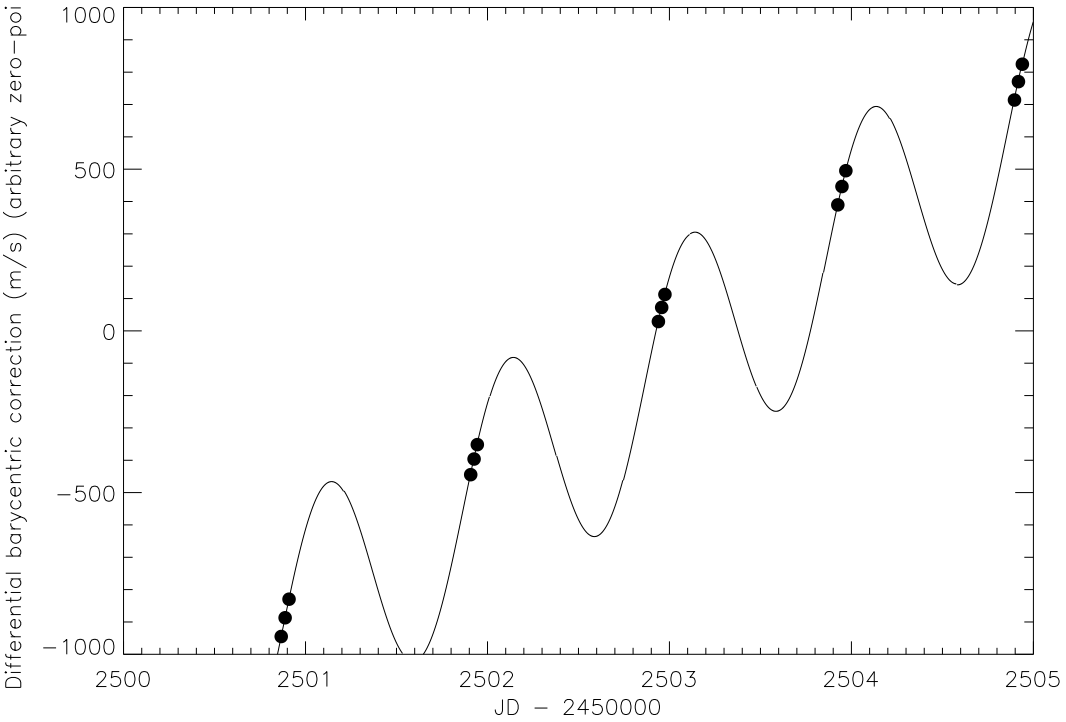

Figure 6-15. Differential barycentric corrections for early 51 Peg measurements with the KPNO ET. The continuous line shows the barycentric correction as a function of time. The points mark the actual times of observation. The arbitrary RV zero-point is chosen for clarity to indicate the scale of variation.

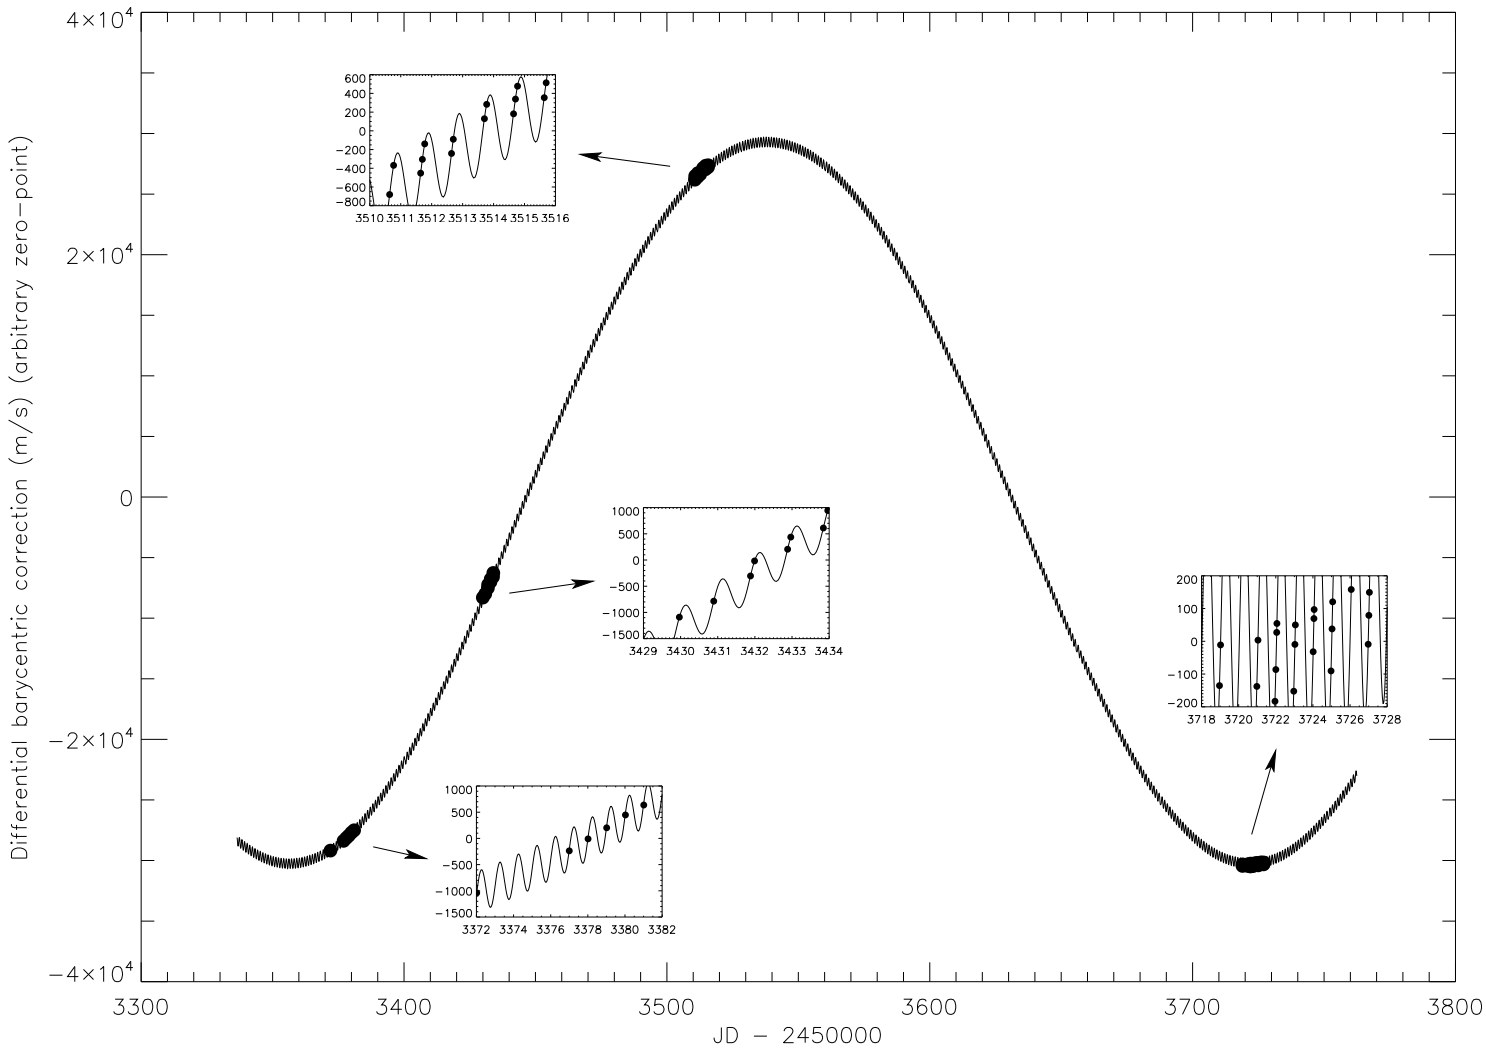

Figure 6-16. Differential barycentric corrections for HD 102195 measurements with the KPNO ET. As in figure 6-15, the continuous line shows the barycentric correction, and the points mark actual times of observation. Insets show zoomed in plots for each observing run. The y-axes all have arbitrary zero-points chosen for clarity to indicate the scale of variation.

CHAPTER 7
RESULTS FROM THE MULTI-OBJECT KECK ET AT APO

In March and April 2005 we demonstrated the feasibility of using the Sloan Digital Sky Survey (SDSS) 2.5 m telescope at Apache Point Observatory for multi-object Doppler measurements, with a 20-object DFDI instrument. Precisions of around $16\,\mathrm{m\,s^{-1}}$ were achieved on bright reference stars. From August 2005 to February 2006, the full scale 60-object W. M. Keck ET was designed and developed. First light was obtained in March 2006, and after some modifications and adjustments, initial science observations began, with further observing runs in April and May 2006. Further observing runs used much of the bright time during 2006, and continuing engineering between May and November 2006 steadily improved the precision of the results. The potential for a multi-object instrument was discussed in Ge (2002); Ge et al. (2002, 2003a); Mahadevan et al. (2003); preliminary results were reported in Ge et al. (2006b); Mahadevan et al. (2005); van Eyken et al. (2007).

Major engineering upgrades are also currently underway, including replicating the Keck ET to double the simultaneous target capacity to 120 objects, in anticipation of a 'Multi-object APO Radial Velocity Exoplanet Large-area Survey'[1] (MARVELS). MARVELS is to be conducted with the Keck ET at the SDSS telescope as part of the planned Sloan Digital Sky Survey III (SDSS III) in 2008–2014 (Ge et al. 2007). Here we report the early results from the prototype, and from observations with the full 60-object Keck ET up until November 2006 which formed the trial survey for the proposed pilot program that would run from 2006–2008.

A brief description of the instrument is given in section 3.3. The most up-to-date instrument description can be found in detail in Ge et al. (2006b); Wan et al. (2006);

[1] Previously referred to in earlier publications by the name 'All Sky Extrasolar Planet Survey' (ASEPS).

Zhao & Ge (2006). Data were reduced using the procedures outlined in chapter 4, the full multi-object pipeline being built upon the original KPNO ET pipeline during the course of the observations.

Target selection was led by Suvrath Mahadevan and Roger Cohen, and is discussed in detail in Mahadevan (2006). Early fields were mainly chosen around known planet-bearing or RV-stable reference stars. Aside from these reference stars, survey stars in the fields were chosen in the magnitude range $7.6 < V < 12.0$, and similarly to the single-object KPNO ET survey, limited where possible to slow-rotating ($v \sin i < 12\,\mathrm{km\,s^{-1}}$) main sequence or sub-giant stars of type F7–M0.

7.1 Prototype

The prototype multi-object instrument built in March and April 2005 represented a feasibility study for the concept. 20 simultaneous spectra were successfully obtained on-sky (Ge et al. 2005), of which 10 had good enough fringes to produce useful RV results. (The remaining 10 were unusable due to various effects – primarily low fringe visibility – caused by instrument misalignment.) The phase-velocity scale, Γ, was determined by choosing it to give the smallest residuals for the reference stars observed. Figure 7-1 shows results from two reference stars observed, 36 UMa (RV stable) and 55 Cnc (planet bearing), and an example search star, TYC 4985-485-1, which shows large RV variation, all observed in April. The predicted RV signals are overplotted for 36 UMa and 55 Cnc. For the latter, the innermost known planet, 55 Cnc e, is again neglected, as for the early observations reported during the detection of ET-1 (section 6.2.1.2). The outer planets again would have a negligible effect on this timescale. 36 UMa (magnitude $V = 4.8$) shows an RMS residual of $16.1\,\mathrm{m\,s^{-1}}$ and 55 Cnc ($V = 5.95$) shows an RMS of $15.5\,\mathrm{m\,s^{-1}}$; both (coincidentally) show photon errors of $11.7\,\mathrm{m\,s^{-1}}$, indicating that there are significant sources of error in addition to photon noise. The search star shows a strong RV variation of $\sim 2500\,\mathrm{m\,s^{-1}}$ over 5 days, and is an example of a probable binary system.

These results were reported in Ge et al. (2006b). Along with the other data obtained, they indicated that although there were problems to be ironed out, performance was sufficient to proceed with a full 60-object instrument.

7.2 May 2006 Results

During trial observations with the newly installed Keck Exoplanet Tracker in May 2006 we were able to successfully obtain 54 usable simultaneous fringing stellar spectra on eight different fields, of a quality sufficient to begin a survey for short-period hot-Jupiter type planets. Of the 60 spectra originally designed for, three were lost off the edges of the detector, three were lost due to the seam between the two stacked beamsplitters in the interferometer, and one suffered from extremely low throughput, probably because of a damaged fibre. A total of 432 stars were surveyed over a period of 12 nights.

Figure 7-2 shows an example raw-data frame, showing a full set of simultaneous stellar spectra. Data were taken to measure the apparent solar velocity reflected off the daytime sky (figure 7-3), providing a zero reference for which typical RMS scatters were around 18–27 $\mathrm{m\,s^{-1}}$ over the duration of the run, and 12–16 $\mathrm{m\,s^{-1}}$ (matching the photon limit) over periods of a few hours.[2] The phase-velocity scale was determined by allowing it to float and minimising the χ^2 of the residuals in the day sky data. Figure 7-4 shows example results obtained for three known planet-bearing stars used as references, including ET-1, showing performance close to the photon limit. From the known hot-Jupiter occurrence rate of 1% among solar-type stars, we expected ∼4 hot-Jupiter type planets to be present in the total sample. Among the search targets from the run (examples of

[2] Day sky data is a useful calibrator in that barycentric corrections for the Sun are always minimal, since the Earth's orbital motion is always essentially perpendicular to the line-of-sight to the Sun, so that the only significant motion is diurnal rotation, which is much smaller ($\sim$ 500 $\mathrm{m\,s^{-1}}$ at the equator, vs. $\sim$ 30 $\mathrm{km\,s^{-1}}$ orbital velocity). This helps avoid the effects of the addition approximation non-linearity.

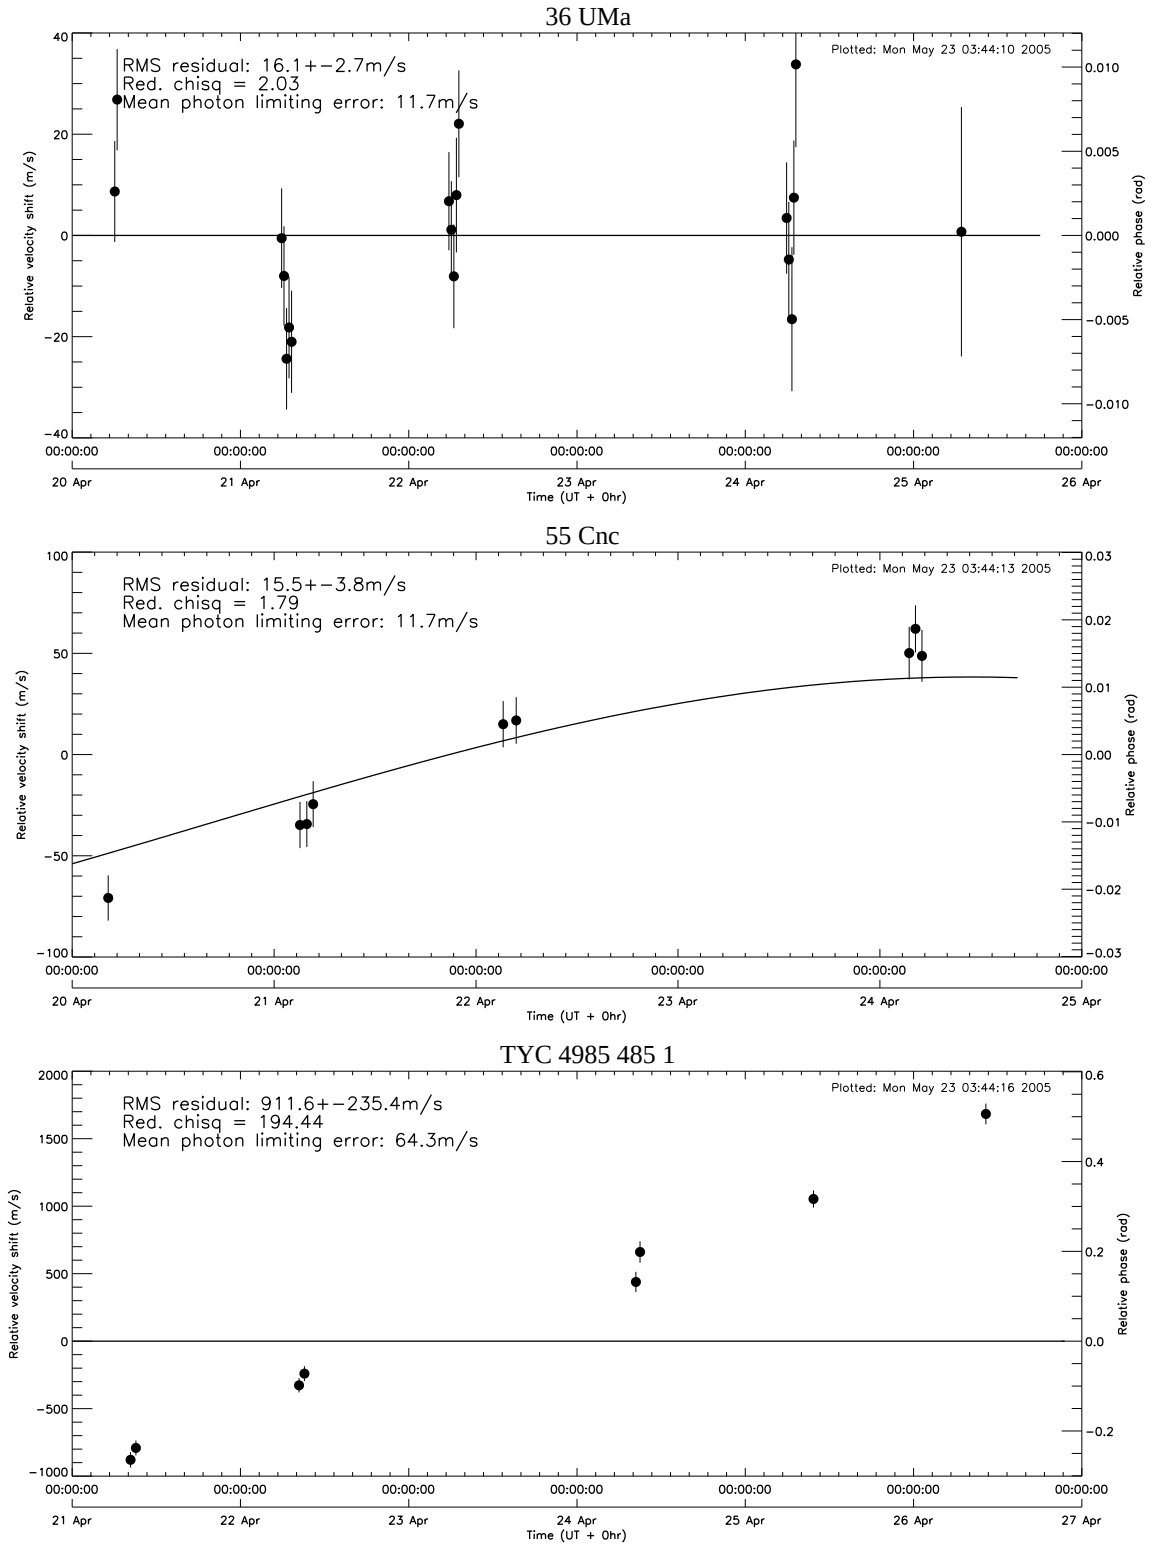

Figure 7-1. Results from the 20-object prototype instrument in April 2005, for RV-stable star 36 UMa, known planet-bearing star 55 Cnc, and an example survey star, TYC 4985 485 1. Expected RV curves are overplotted for 36 UMa and 55 Cnc. Error bars represent photon errors.

which are shown in figure 7-5), around 15 were chosen as interesting candidates for further followup. These preliminary results were reported in van Eyken et al. (2007).

Although we had obtained reasonable results, there was still clearly room for improvement at this point. Instrument throughput was around a third of the design specifications, increasing the photon errors correspondingly by a factor of $\sqrt{3}$. Fringe visibility was also significantly lower than expected, at around half that of the KPNO single-object ET. Since photon precision varies inversely with fringe visibility (equation 2–31), this also could cause a further factor of two loss in precision. The 12-day noise floor of $\sim 20\,\mathrm{m\,s^{-1}}$ seen in the solar data also suggested further sources of systematic error. These issues were to be targeted in the following engineering runs.

7.3 November 2006 Results

By November 2006, various significant engineering improvements had been made, including, among other things, improving the throughput, replacing the pair of stacked beamsplitters in the interferometer with a single-piece beamsplitter, realigning the optics and refocusing the instrument. These modifications, along with fine-tuning of the data reduction pipeline, led to an improvement of around a factor of two in the photon limit, and typically 3–5 times in RMS scatter over the May 2006 results. The Keck ET also now yielded 59 usable simultaneous fringing spectra.

Observations were taken over 5 nights on an observing run in November 2006, visiting 6 fields for a total of 354 searched stars and over 1700 data points. Γ was again determined by minimising χ^2 for the day-sky reference data taken. Figures 7-6 and 7-7 show measurements of two known planet-bearing reference stars, HD 209458 (magnitude $V = 7.70$) and HIP 14810 ($V = 8.59$), along with their predicted RV curves. Both show a good match to the predictions, with significantly smaller errors than in May, HD 209458 showing RMS residuals of $14.1\,\mathrm{m\,s^{-1}}$ and a mean photon limiting error of $8.5\,\mathrm{m\,s^{-1}}$. Although there were only two data points for HIP 14810, the difference in RV between the

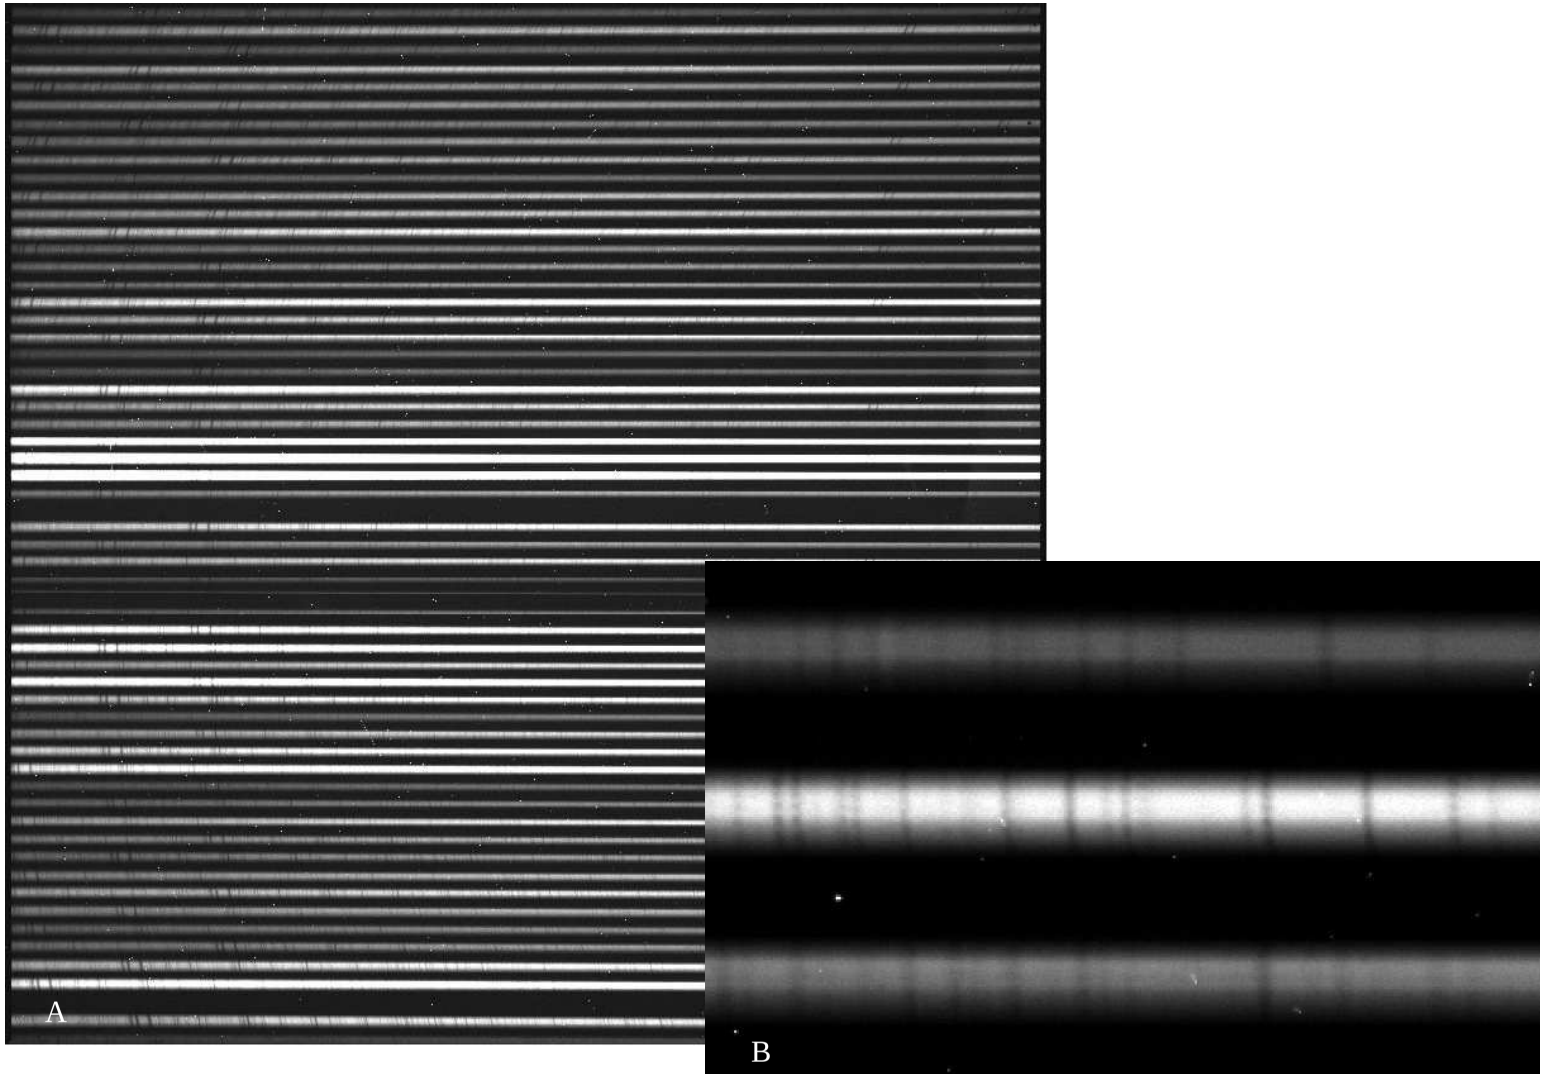

Figure 7-2. Example raw data frame from early Keck ET, May 2006, showing 54 usable simultaneous stellar spectra, without iodine. A) Full data frame. B) Magnified portion. Interferometer fringes can faintly be seen.

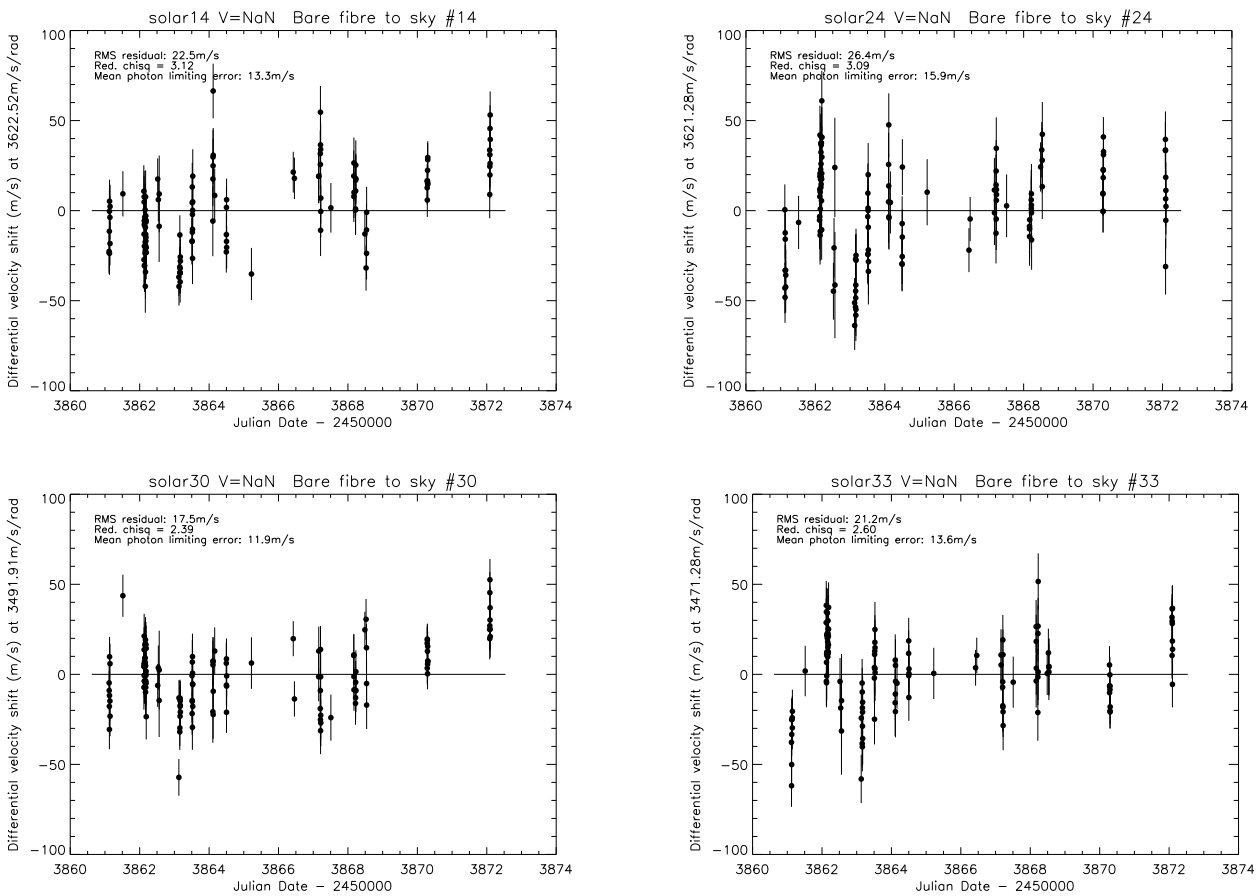

Figure 7-3. Example solar data from early Keck ET, May 2006, obtained by pointing fibre ends directly to daytime sky, showing the expected zero velocity drift. Short term (few hour) RMS matches the photon noise. Results are from simultaneous measurements with different fibres ('#' indicates spectrum number on the detector).

two points is close enough to the prediction to be a useful confirmation of performance, with a photon limiting error of $16.4\,\mathrm{m\,s^{-1}}$.

The general improvement in instrument precision can clearly be seen in the day-sky data taken during the run (figure 7-8). Most of the 59 spectra on the detector showed RMS scatters in the range $\sim$ 7–10 $\mathrm{m\,s^{-1}}$ over the five days of the run, down from 18–27 $\mathrm{m\,s^{-1}}$ in the May data. While a few had larger RMS scatters, some approached their photon limit, typically $6\,\mathrm{m\,s^{-1}}$ (down from 12–16 $\mathrm{m\,s^{-1}}$). The measurement of sky data simultaneously with all the fibres provides a useful test: in principle all the fibres should be measuring exactly the same results, within the noise. Any systematic differences

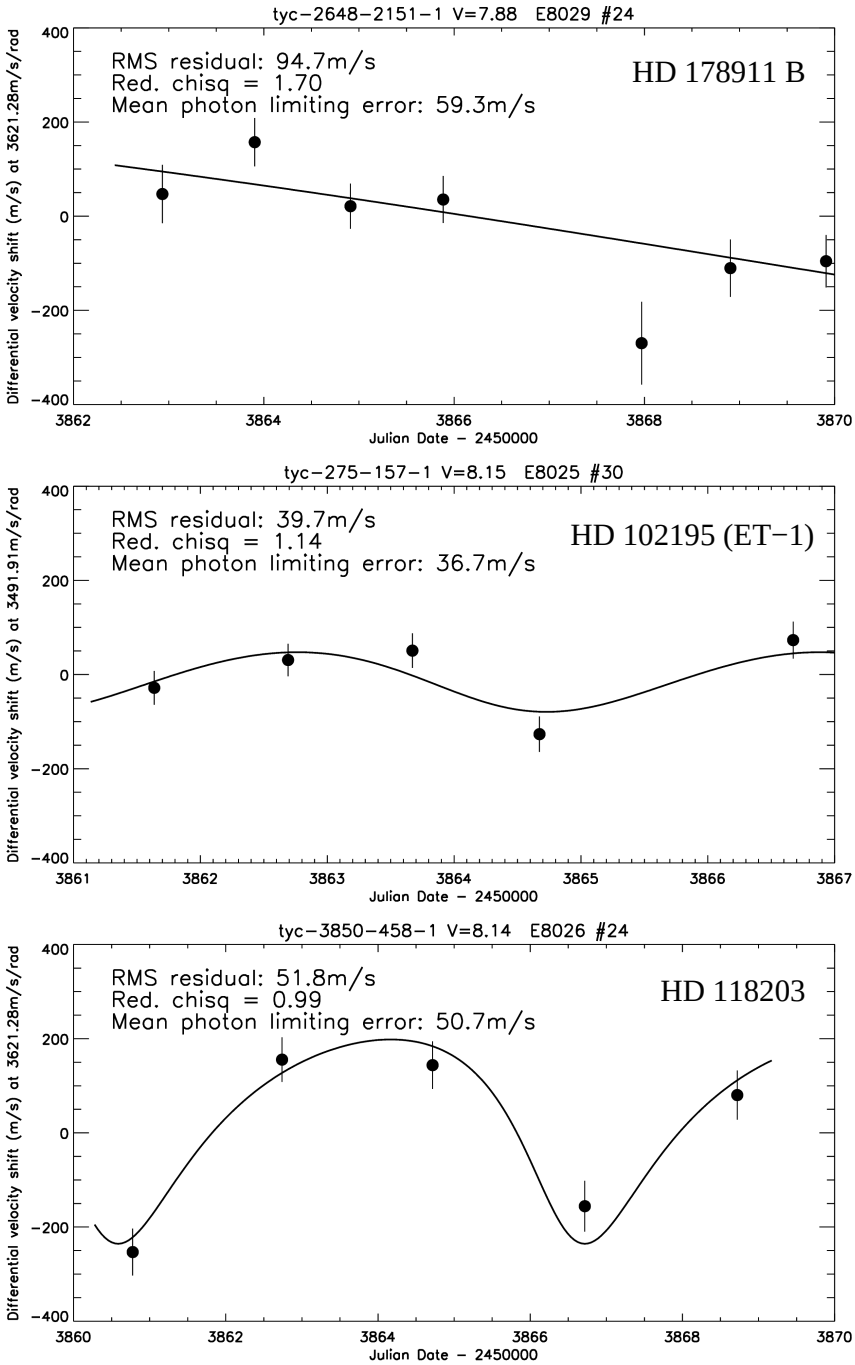

Figure 7-4. Known planet-bearing stars from early Keck ET data, May 2006, used as references: HD 178911 B, ($m \sin i = 6.3\, M_{\rm J}$, period=71.5 d), HD 102195 (ET-1, $m \sin i = 0.48\, M_{\rm J}$, period=4.1 d), and HD 118203 ($m \sin i = 2.13\, M_{\rm J}$, period=6.1 d). Predicted RV curves from previously published measurements are overplotted (parameters from Butler et al. (2006) and the online Extrasolar Planets Encyclopaedia, http://exoplanet.eu/).

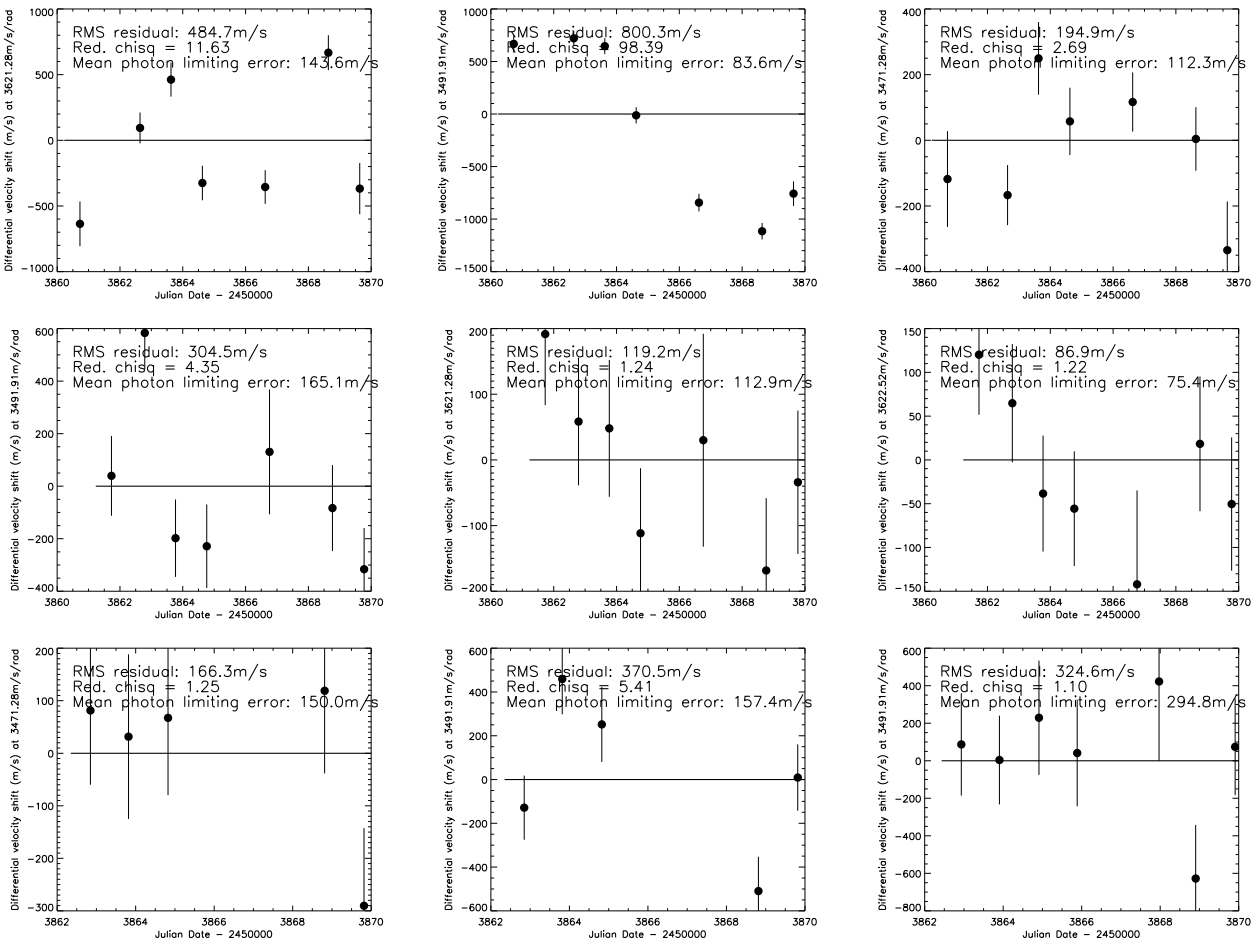

Figure 7-5. Small selection of different example search star results from early Keck ET data, May 2006. One shows an apparent sinusoidal RV signal. If not due to systematic errors, this could be a planet, or more likely, a spectroscopic binary. A Keplerian fit to the data suggests a companion of $m \sin i \sim 11\,M_{\rm J}$, period$\sim 12\,$d.

must therefore be due to instrument or pipeline issues. The causes of these differences seen are still under investigation.

7.4 Discussion

As of the November 2006 results, there was still room for improvement in the instrument. Among other things, instrument throughput could still be improved by $\sim$ 30%, and spectral resolution by a factor of $\sim$ 1.5. The interferometer-delay active locking, using only one-beam tracking, was also not precise enough over the large field required for all 59 spectra (see section 3.1), leading to poor correction for many of the

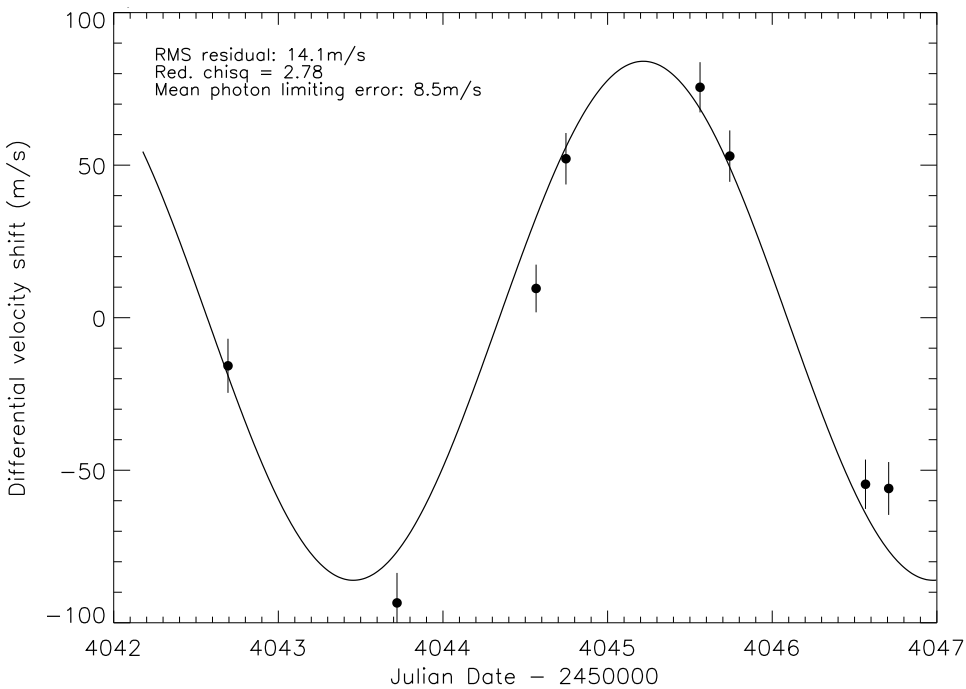

Figure 7-6. Keck ET measurements of HD 209458 from November 2006. Error bars represent photon limiting errors. Predicted RV curve is overplotted (orbital parameters taken from Naef et al. 2004).

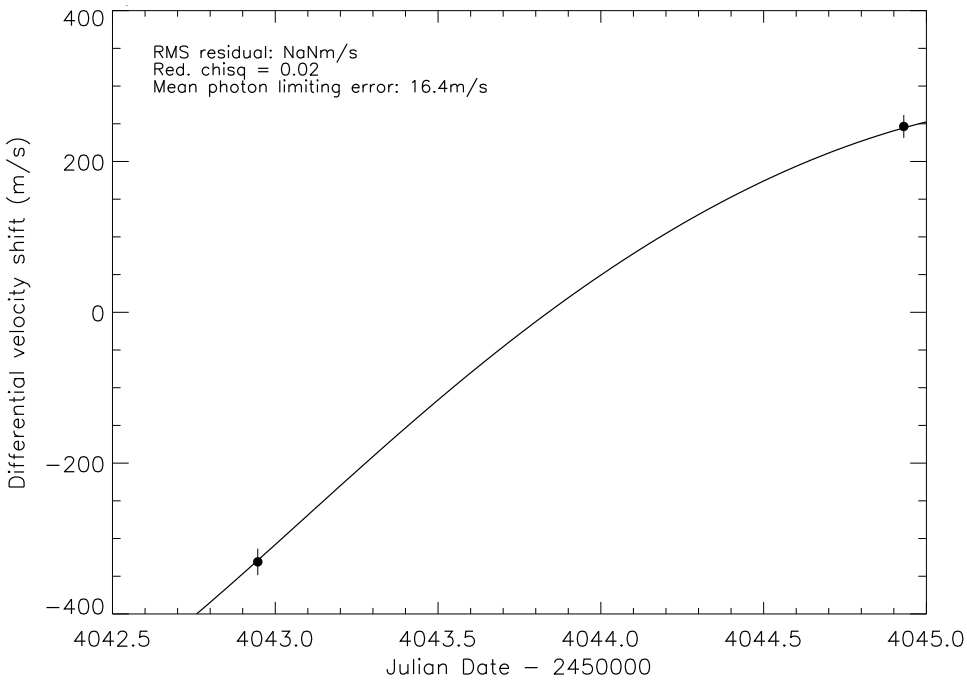

Figure 7-7. Keck ET measurements of HD 14810 (TYC 1231 1727 1) from November 2006. Error bars represent photon limiting errors. Predicted RV curve is overplotted (orbital parameters taken from Butler et al. 2006).

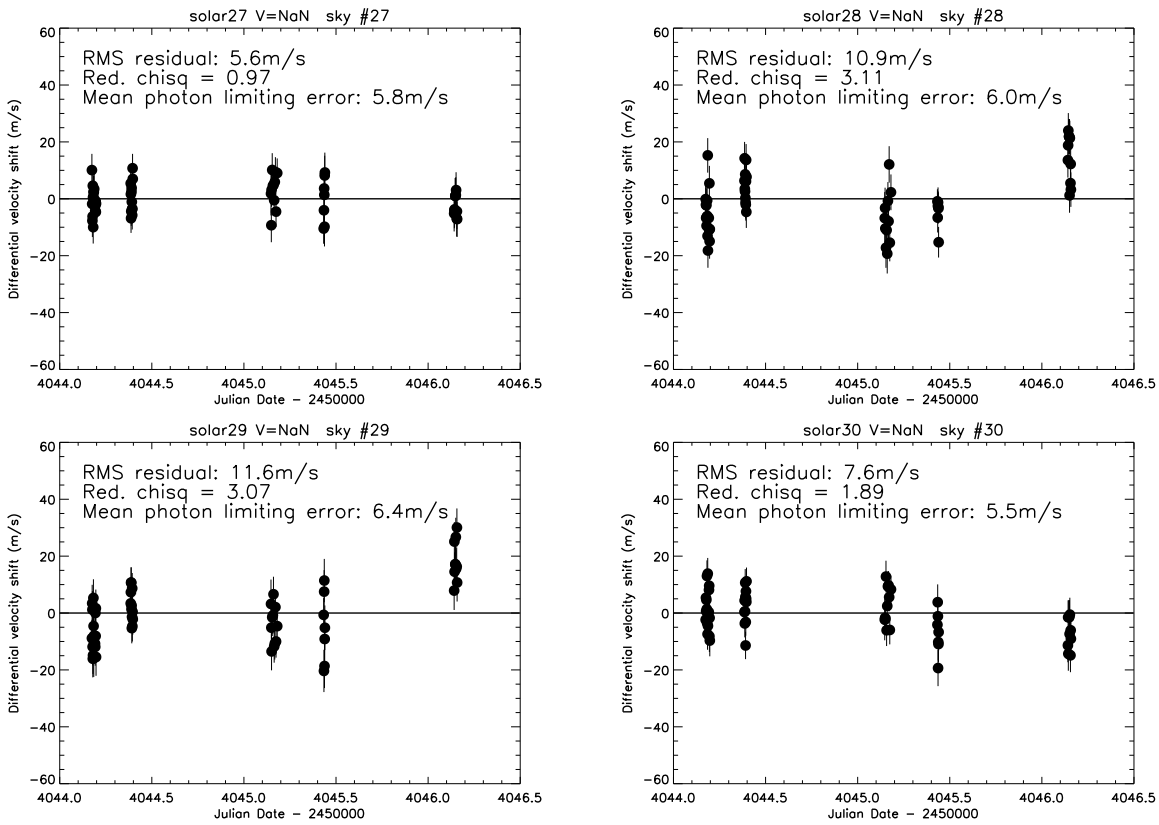

Figure 7-8. Representative selection of day-sky RV data from Keck ET, November 2006, from different fibres, showing expected zero velocity-shift. Results are for four simultaneous spectra near the centre of the detector, numbers 27–30 (indicated by '#' at top of plots).

spectra. Uncorrected instrument drift for these observations without the iodine reference was still at tens of $\mathrm{km\,s^{-1}}$, and a significant reduction in fringe visibility for many of the spectra was seen due the large phase drifts occurring during the length of exposures. The system has now been upgraded to two-beam phase locking, and fringes for all of the 59 objects are stable to $\sim 1/300$ waves, or $\sim 60\,\mathrm{m\,s^{-1}}$, over 5 hours.

Systematic errors were still seen in many instances. Sources of error, still under investigation, may have included moonlight contamination, scattered light cross-talk between spectra, and effects of distortion in the optical design. Most significantly, the addition approximation error (see sections 2.6.1 and 5.10) had not yet been addressed. It is likely that this last effect explained why we were able to get good results in some cases, but saw systematic errors in many others.

The data here represent a demonstration that the multi-object concept works, and the beginnings of an initial planet survey. We are currently searching and compiling the results of this trial survey to find the strongest planet candidates for further followup; we expect that a handful of genuine planets should be detectable in the data. As a byproduct, we are also compiling the clear stellar binary systems that are found in the data: since the signals from stellar binaries are so large (on scales of $\mathrm{km\,s^{-1}}$), they should be well above our error levels, and will not be buried by the addition approximation problem, which should only cause errors on the level of at most $100\,\mathrm{m\,s^{-1}}$. All the issues mentioned above are being addressed in ongoing engineering and software upgrades, in preparation for the full MARVELS survey to begin in mid-2008.

CHAPTER 8
MEASURING STELLAR ROTATION WITH A DFDI

8.1 Background

We saw in section 2.4 that the visibility of a fringe is given by the normalised absolute value of the complex Fourier transform of the input spectrum, evaluated at the interferometer delay $d = d_0$. For a single spectral line, the reciprocal scaling principle of Fourier transforms implies that as the line gets broader, its Fourier transform as a function of d will become narrower. At the fixed delay value d_0, we can therefore expect the fringe visibility due to the line to decrease as the line width increases. Stellar rotation has the effect of broadening spectral lines: it is therefore interesting to ask, can measurements of fringe visibility be used to infer the projected stellar rotation velocity using a DFDI instrument?

The broadening of a stellar spectral line about its central wavelength is caused by a combination of several effects (Carroll & Ostlie 1996). 'Natural broadening', due to the Heisenberg uncertainty principle, represents a fundamental physical limit to how narrow a line can be: the absorption is caused by photons exciting electrons to higher orbitals, and the limited time that an electron spends in its excited state means that there is a corresponding uncertainty in the energy of that orbital, and hence in the exact wavelength at which absorption occurs. 'Doppler', or 'thermal' broadening, is caused by a spread in the Doppler shifts of the absorbing atoms in the gas of the photosphere due to their thermal motion. 'Pressure' and 'collisional' broadening are caused by perturbations in the orbitals of the absorbing atoms due to close passage and collisions with other neutral atoms and ions, and are thus tied with the pressure of the gas in the photosphere. Other broadening effects include micro- and macro-turbulence, Doppler broadening caused by large-scale convective motions of the gas in the photosphere. Of these effects, the core of the line profile is dominated by thermal broadening and micro- and macro-turbulence in solar type stars. The combination of all these effects gives what we term the 'intrinsic'

line width for a star. Typical full-width half-maximum (FWHM) intrinsic line widths for a non-rotating solar-type spectrum are on the order of 0.1 Å.

Rotational broadening is caused by the fact that the starlight we observe is integrated over the full stellar disc on the sky (Gray 1992). The observed spectrum is a summation of the spectra contributed by each point on the stellar disc, each with its own projected line-of-sight velocity due to the rotation of the star, and weighted by the limb darkening profile across the disc of the star. The line-of-sight velocity varies from zero at the centre of the disk to a maximum $v \sin i$ at the limb on the equator, where v is the equatorial rotation velocity, and i is the angle of inclination between the line of sight and the rotation axis of the star. There is a degeneracy between v and i which means that only the combination $v \sin i$ is measurable using spectroscopic techniques. The total observed line profile is found to be a convolution of the intrinsic line profile and a rotation profile whose width is determined by the value of $v \sin i$, and whose precise shape is determined by the stellar limb-darkening profile (Gray 1992).

We address the question of whether we can measure $v \sin i$ by measuring simply the mean visibility of the fringes across a fringing spectrum. The exact form of a spectrum depends on effective temperature, T_e, and metallicity, $[M/H]$,[1] which both govern the atomic lines that are present and their respective depths; and the surface gravity, $\log g$, which determines the pressure in the photosphere, and therefore pressure and collisional broadening. For simplicity we neglect the effects of micro- and macroturbulence, and assume these to be included in the intrinsic line profile. Since all the spectral lines will be broadened in the same way by stellar rotation, the question is whether for given values of T_e, $\log g$ and $[M/H]$, and a given instrument configuration (wavelength coverage, delay d_0, and resolution R), we can usefully calibrate a scale which will map the measured

[1] Measured on a base-10 logarithmic scale relative to solar metallicity, representing an abundance across all the metals.

mean fringe visibility to a value of $v \sin i$. In addition we are interested in how accurately it can be done, and how precisely the stellar parameters need to be known in order to achieve this. In principle it may be possible to recover the other stellar parameters directly from low-resolution spectra obtained with ET, which could be created by appropriately processing the data to remove the fringes: the feasibility of this is under investigation by Roger Cohen (Cohen et al. 2006). If the stellar parameters are already known, a $v \sin i$ measurement would be very simple to achieve, and could potentially be a useful byproduct of a DFDI survey. We discuss here the results from a preliminary investigation into the feasibility of such measurements.

8.2 Theoretical Predictions

8.2.1 The Broadened Line Profile

To understand the behaviour of the fringe visibility as $v \sin i$ changes, it is helpful to try and perform a very rough theoretical analysis. We begin by considering a single spectral line, modelling it for simplicity as the convolution of two Gaussian functions, one representing the intrinsic line profile, and the other approximating the rotational broadening kernel. Working in wavenumber space, κ, we express the resulting spectrum, P, as:

$$P(\kappa) = [I(1 - g_{\rm i}(\kappa))] \otimes g_{\rm r}(\kappa), \tag{8–1}$$

where I is a constant representing the continuum flux level, and $g_{\rm i}$ and $g_{\rm r}$ are Gaussian functions with the subscripts i and r referring to the intrinsic line profile and the rotational broadening profile respectively. Writing out the Gaussians explicitly:

$$g_{\rm i}(\kappa) \equiv \delta_{\rm i} {\rm e}^{-\kappa^2/(2\omega_{\rm i}^2)} \quad ; \quad g_{\rm r}(\kappa) \equiv \frac{1}{\omega_{\rm r}\sqrt{2\pi}}\, {\rm e}^{-\kappa^2/(2\omega_{\rm r}^2)}. \tag{8–2}$$

$\delta_{\rm i}$ represents the fractional depth of the intrinsic line ($0 \leq \delta_{\rm i} \leq 1$). $\omega_{\rm i}$ and $\omega_{\rm r}$ are measures of the widths of the two Gaussians (actually, the dispersions of the Gaussians). The factor in front of the broadening profile is a normalisation factor so that $g_{\rm r}$ has unit area. To simplify the notation, we take the intrinsic line to be centred at $\kappa = 0$. (Shifting the centre

of the line to some other value $\kappa = \kappa_0$ will give exactly the same results under convolution, but simply shifted to the new line-centre.)

Since convolution is distributive, we can expand out equation 8–1:

$$\begin{aligned} P(\kappa) &= I[(1 \otimes g_{\rm r}) - (g_{\rm i} \otimes g_{\rm r})] \\ &= I[1 - (g_{\rm i} \otimes g_{\rm r})]. \end{aligned} \tag{8–3}$$

The profile of the broadened line is therefore just $g = g_{\rm i} \otimes g_{\rm r}$, which we can calculate by taking the Fourier transform over κ:

$$\begin{aligned} \widehat{g}(d) &= \widehat{g_{\rm i} \otimes g_{\rm r}} = \widehat{g}_{\rm i} \cdot \widehat{g}_{\rm r} \\ &= \delta_{\rm i}\omega_{\rm i}\sqrt{2\pi}\,{\rm e}^{-2\omega_{\rm i}^2(\pi d)^2} \;\cdot\; {\rm e}^{-2\omega_{\rm r}^2(\pi d)^2} \\ &= \delta_{\rm i}\omega_{\rm i}\sqrt{2\pi}\,{\rm e}^{-2\pi^2 d^2(\omega_{\rm i}^2+\omega_{\rm r}^2)}, \end{aligned} \tag{8–4}$$

where d is the conjugate variable to κ, which we later associate with the interferometer delay. Now reversing the Fourier transform to find g, we find:

$$g(\kappa) = \frac{\delta_{\rm i}\omega_{\rm i}}{\sqrt{\omega_{\rm i}^2 + \omega_{\rm r}^2}}\,{\rm e}^{-\kappa^2/[2(\omega_{\rm i}^2+\omega_{\rm r}^2)]} \tag{8–5}$$

In other words, the rotationally broadened line profile is itself approximated by a Gaussian, whose width is equal to the quadrature summation of the intrinsic and broadening profile widths. The depth of the line is proportional to the intrinsic line depth, $\delta_{\rm i}$, and scales inversely with the final line width by $\omega_{\rm i}/\sqrt{\omega_{\rm i}^2 + \omega_{\rm r}^2}$, so that the total flux under the line stays constant (consistent with the conservation of equivalent width expected from rotational broadening).

8.2.2 Fringe Visibility

Returning to the input spectrum due to the absorption line following equation 8–3, for a single rotationally broadened line, we can write:

$$P(\kappa) = I(1 - g(\kappa)) \tag{8–6}$$

To find the visibility due to this spectrum at wavelength channel j, we use the wavenumber version of the complex visibility formula from chapter 2 (equation 2–9):

$$\boldsymbol{\gamma} = \frac{\mathcal{F}[P(\kappa)w_j(\kappa)]_{d=d_0}}{\mathcal{F}[P(\kappa)w_j(\kappa)]_{d=0}} \tag{8–7}$$

We take the spectrograph response function, $w_j(\kappa)$,[2] to be centred on the wavelength of the absorption line, so that we are choosing to look at the particular channel j in the spectrum where the fringe due to the line lies. We note that by the Fourier shift theorem, shifting the absorption line to any centre $\kappa = \overline{\kappa}$ (along with its corresponding response function) must result only in a phase shift of $\boldsymbol{\gamma}$: since we are only concerned with the absolute visibility, $\gamma = |\boldsymbol{\gamma}|$, such a phase shift is of no consequence, and, without loss of generality, we can continue as if g and w are centred at $\kappa = 0$.

Substituting equation 8–6 into the complex visibility equation 8–7, we can write the Fourier transform in the numerator of the complex visibility equation 8–7 as

$$\mathcal{F}[P(\kappa)w(\kappa)]_{d=d_0} = \mathcal{F}[Iw(1-g)]_{d=d_0}, \tag{8–8}$$

where we have dropped the j subscript and assumed that the (κ) dependency is implicit in w and g. If we approximate the response function $w(\kappa)$ as a top-hat function of width Δw and unit height, and assume that this width is much larger than that of the line, then looking at figure 8-1 we see that we can assume $Iw(1-g) \approx I(w-g)$ so that:

$$\mathcal{F}[P(\kappa)w(\kappa)]_{d=d_0} \approx I(\widehat{w-g}) = I(\widehat{w} - \widehat{g}) = -I\widehat{g}. \tag{8–9}$$

[2] Note the distinction between w and ω. Also note that as shown in appendix A, w is essentially just the reverse of the LSF, so that the widths of w and the LSF are the same: for the purposes of the discussion in this chapter, the concepts of response function and LSF (or resolution element) can be used interchangeably.

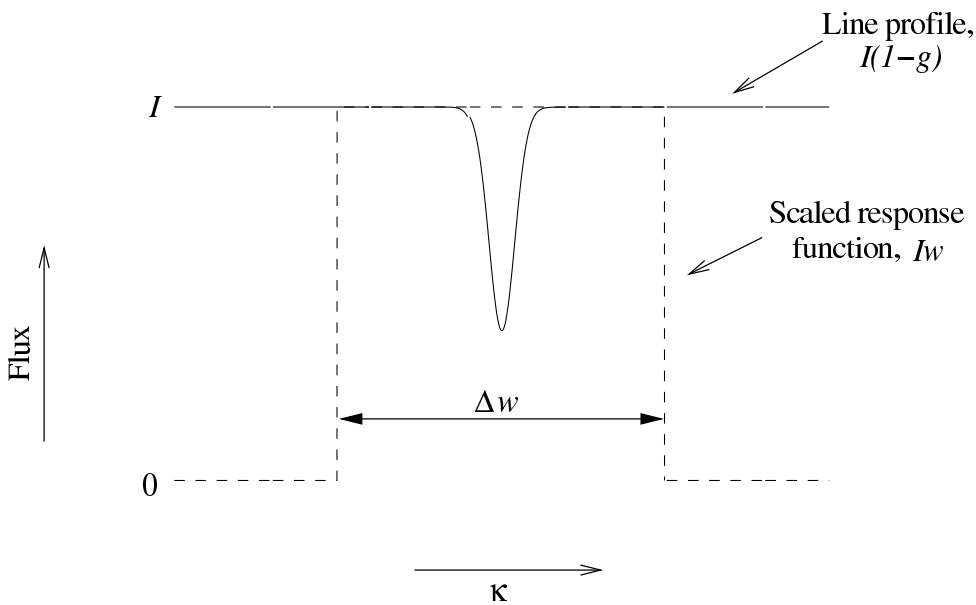

Figure 8-1. Top-hat response function centred on a single Gaussian absorption line.

The $\widehat{w}$ term vanishes because it is a pure comb term, and the interferometer delay d_0 (and/or the low pass filtering in the data reduction) is always chosen such that the comb is invisible and $\widehat{w} \to 0$ at $d = d_0$ (see section 2.6.1).

Using equation 8–4 we can substitute $\widehat{g}$ into equation 8–9 so that:

$$\mathcal{F}[P(\kappa)w(\kappa)]_{d=d_0} \approx -I\delta_{\rm i}\omega_{\rm i}\sqrt{2\pi}\,{\rm e}^{-2\pi^2 d^2(\omega_{\rm i}^2+\omega_{\rm r}^2)}. \tag{8–10}$$

Consider now the denominator of the complex visibility equation 8–7, which is the total flux under $P(\kappa)w(\kappa)$. Referring again to figure 8-1, and using the standard result for the area under a Gaussian, we can see that:

$$\begin{aligned}\mathcal{F}[P(\kappa)w(\kappa)]_{d=0} = \int Pw\ {\rm d}\kappa &= I\Delta w - I\int g(\kappa)\ {\rm d}\kappa \\ &= I(\Delta w - \delta_{\rm i}\omega_{\rm i}\sqrt{2\pi}).\end{aligned} \tag{8–11}$$

Finally substituting the numerator and denominator (equations 8–10 and 8–11) into the complex visibility equation 8–7 and taking the absolute value to determine the

absolute visibility gives:

$$\gamma = |\boldsymbol{\gamma}| = \frac{\delta_{\mathrm{i}}\omega_{\mathrm{i}}\sqrt{2\pi}\,\exp[-2\pi^2 d_0^2(\omega_{\mathrm{i}}^2 + \omega_{\mathrm{r}}^2)]}{\Delta w - \delta_{\mathrm{i}}\omega_{\mathrm{i}}\sqrt{2\pi}}. \tag{8–12}$$

Since wavelength is generally the preferred unit of measure in optical spectroscopy, we need to convert the width intervals from wavenumber to wavelength. For a wavenumber interval $\Delta\kappa$ corresponding to the wavelength interval $\Delta\lambda$ centred at wavelength $\overline{\lambda}$, we can approximate:

$$\Delta\kappa \approx \frac{-\Delta\lambda}{\overline{\lambda}^2}. \tag{8–13}$$

We can therefore write for the visibility:

$$\boxed{\gamma(\omega_{\lambda\mathrm{r}}) = \frac{\delta_{\mathrm{i}}\omega_{\lambda\mathrm{i}}\sqrt{2\pi}\,\exp[-2\pi^2 d_0^2(\omega_{\lambda\mathrm{i}}^2 + \omega_{\lambda\mathrm{r}}^2)]}{\overline{\lambda}^4(\Delta w_\lambda - \delta_{\mathrm{i}}\omega_{\lambda\mathrm{i}}\sqrt{2\pi})},} \tag{8–14}$$

where the subscript λ is used to indicate that the width measures have now been replaced with their wavelength equivalents. The FWHM of the intrinsic and rotational Gaussian profiles is related to the dispersion widths, ω, by:

$$\omega = \frac{\mathrm{FWHM}}{2\sqrt{2\ln 2}}. \tag{8–15}$$

A least χ^2 fit Gaussian profile to the more accurate rotational broadening profile given by Gray (1992, equation 17.2) (see figure 8-2) gives

$$\mathrm{FWHM_r} \approx 1.30\Delta\lambda_{\mathrm{L}} = 1.30\frac{\lambda}{c}\cdot v\sin i \tag{8–16}$$

where $\Delta\lambda_{\mathrm{L}}$ is the maximal Doppler shift at the limb of the stellar disc, related directly to $v\sin i$ via the Doppler shift equation, and c is the speed of light.

8.2.3 Predicted Relation

We are now in a position to plot an estimated curve relating visibility, γ, to $v\sin i$. Using equation 8–14 to 8–16, we take approximately solar parameters, with instrument parameters corresponding to those assumed in the simulations discussed later (section 8.3): we take $\mathrm{FWHM_i} = 0.1\,\text{Å}$ for the typical stellar intrinsic line width; $\Delta w = 0.97\,\text{Å}$

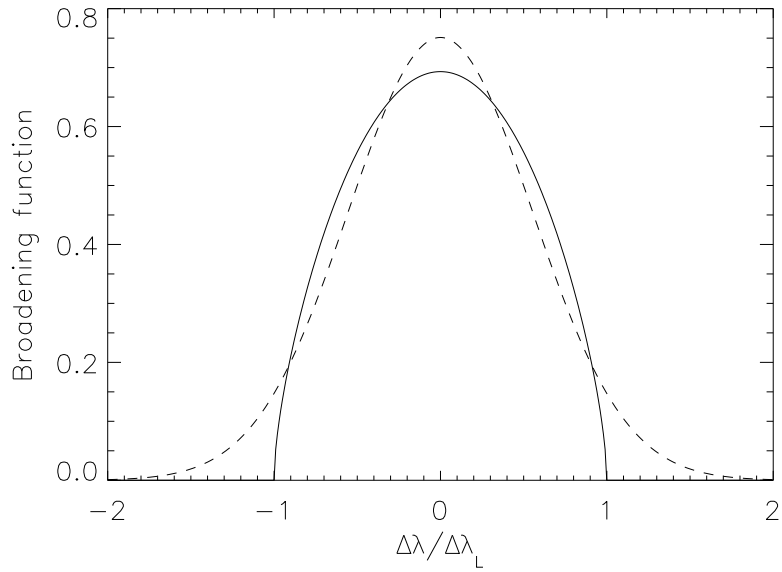

Figure 8-2. Gaussian approximation to a normalised rotational broadening profile. Solid line: broadening profile following Gray (1992, equation 17.2, with limb-darkening coefficient $\varepsilon = 0.6$). Dashed line: best-fit Gaussian. The y-axis is in units of the maximal stellar Doppler shift Δ_L, corresponding to the projected equatorial rotation velocity, $v \sin i$.

for the width of the LSF (since this is the width of the low-pass filter applied in the pre-processing for the simulated images) – the same as the width of the spectrograph response function; $d_0 = 7\,\mathrm{mm}$ for the interferometer delay; and $\overline{\lambda} = 5400\,\text{Å}$ as the mean wavelength across the spectrum: since the wavelength coverage of the simulated spectra (5000–5800 Å) is relatively small compared to $\overline{\lambda}$ this value is taken to be a reasonably representative.

To estimate a representative line depth, we take a synthetic stellar spectrum with near-solar parameters (as in section 8.3), and chop it into 0.97 Å segments, corresponding to the width of the LSF. We find a minimum flux level for each segment, and take a mean over all these values, giving a value for δ_i of 0.40. Where there are two lines within a resolution element, the linearity of Fourier transforms suggests that the resulting complex visibility $\boldsymbol{\gamma}$ will be a summation of the complex visibilities due to the two individual lines, scaled down by the additional flux loss below the continuum level. If one line is much deeper than the other, it will dominate the visibility; where the two are of similar

depths, the phases of the two fringe contributions will sometimes be in phase, so that both contribute to the visibility, and sometimes out of phase, so that the visibilities contributing to the total fringe cancels out. We might therefore expect the average resulting visibility to be on the same order as that of a single line, so that the estimate of $\delta_{\rm i}$ remains reasonable. Clearly multiple lines, line blending and segment boundaries which fall within line profiles complicate the issue, however, so this value is only representative.

Figure 8-3 shows the predicted visibility as a function of $v \sin i$ for these parameters, using the analytical visibility equation (equation 8–14). Looking at the form of the equation, we see that it is itself a Gaussian function of the rotational broadening, and hence of $v \sin i$, scaled approximately by the ratio of the area enclosed by the intrinsic line shape to the area under the spectrograph response function. Changing either the intrinsic line depth or the intrinsic line width has only the effect of scaling the Gaussian – neither changes the shape of the function. This fortunate circumstance means that all fringe visibilities across all channels will change by the same fractional amount under a change of $v \sin i$, so that it should be entirely meaningful to take a mean visibility over a complete spectrum to improve S/N. This mean value should itself also scale in exactly the same way.[3] The only parameter which can affect the shape of the visibility function in any way other than an overall scaling is the interferometer delay.

[3] The validity of taking a mean should also remain true for a non-Gaussian broadening kernel, because of the way the intrinsic and broadening profiles multiply in equation 8–4, which leads to the numerator of the visibility equation (equation 8–14): the broadening kernel must depend only on $v \sin i$, and the intrinsic profile cannot have any $v \sin i$ dependency by definition. The denominator of the equation, the total flux within the response function, must also remain approximately independent of the rotational broadening because the broadening is flux conserving.

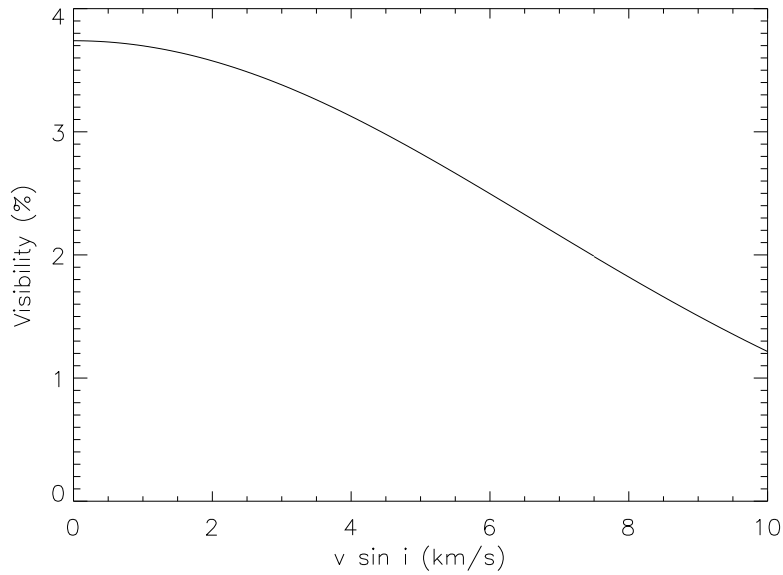

Figure 8-3. Analytically predicted visibility vs. $v \sin i$ curve for a Sun-like star observed with an ET-type instrument operating at 5000–5800 Å with a post-filtering resolution element of width 0.97 Å.

8.3 Simulations

8.3.1 Description

A grid of simulated DFDI spectra with different stellar parameters was created by Suvrath Mahadevan using high-resolution synthetic model spectra provided by Roger Cohen, for the purposes of investigating whether it is possible to recover the various stellar parameters from ET spectra. The high-resolution synthetic spectra were created using the software 'SPECTRUM' (Gray & Corbally 1994) which combines a Kurucz model atmosphere (Castelli & Kurucz 2003) with its own spectral line list to calculate a normalised spectrum. These were then processed through Mahadevan's instrument simulation code to produce fringing spectra similar to those that would be seen by the ET instruments.

The grid covered the ranges $T_e = 4000$–$6500\,\mathrm{K}$ in steps of $250\,\mathrm{K}$; $\log g = 2.0$–$5.0\,\mathrm{dex}$ (cgs units, $\log\,\mathrm{cm\,s^{-2}}$), in steps of $0.5\,\mathrm{dex}$; $[M/H] = -2.5 - +0.5\,\mathrm{dex}$ in steps of $0.5\,\mathrm{dex}$, and additionally $[M/H] = +0.2\,\mathrm{dex}$; and $v \sin i = 0, 2, 5$ and $10\,\mathrm{km\,s^{-1}}$. Random Gaussian

noise was added to all the simulations to simulate photon noise at the level S/N = 50 per pixel.

The simulated fringing spectra covered 5000–5800 Å with a linear wavelength scale and a resolution $R = 6000$ at 5000 Å given by a uniform Gaussian LSF of FWHM=0.83 Å (4.27 pixels). The spectra covered 4096 pixels in the dispersion direction and 52 in the slit direction. The simulations were then run through part of the standard data reduction pipeline, beginning with low-pass filtering (since the spectra were otherwise 'perfect' spectra), at a cutoff period of 5.0 pixels, followed by fringe fitting to obtain the whirls. For each spectrum, a mean fringe visibility was calculated across the full wavelength coverage.

8.3.2 Results and Dependence of $v \sin i$ on Stellar Parameters

Figure 8-4 shows plots of mean visibility against the four stellar parameters, representing cuts across the simulation grid all centred on the grid point that corresponds to the closest match to solar values. For each free parameter on the x-axis, the remaining parameters are fixed at $T_e = 5750\,\mathrm{K}$, $\log g = 4.5$, $[M/H] = 0.0$, and $v \sin i = 2.0\,\mathrm{km\,s^{-1}}$ (cf. true solar values, $T_e = 5770\,\mathrm{K}$, $\log g = 4.44$, $[M/H] = 0.00$, and $v \sin i = 1.7\,\mathrm{km\,s^{-1}}$, from Valenti & Fischer 2005).

Figure 8-5 shows similar plots of mean visibility against $v \sin i$, this time showing contours of T_e, $\log g$, and $[M/H]$. In each case, the unplotted parameters are again held at the nearest-to-solar values in the grid.

Since the simulations have photon noise added, it is necessary to estimate the resulting errors in the mean visibility measurements, $\Delta\overline{\gamma}$. We take:

$$\Delta\overline{\gamma} = \frac{\overline{\Delta\gamma_j}}{\sqrt{n}}, \tag{8–17}$$

where $\Delta\gamma_j$ is the visibility error in channel j as returned by the fringe fitting, over which we take a mean across all j, and n is the number of *independent* wavelength channels (i.e. the number of resolution elements within the range of spectral coverage, ~ 800 for a 4096

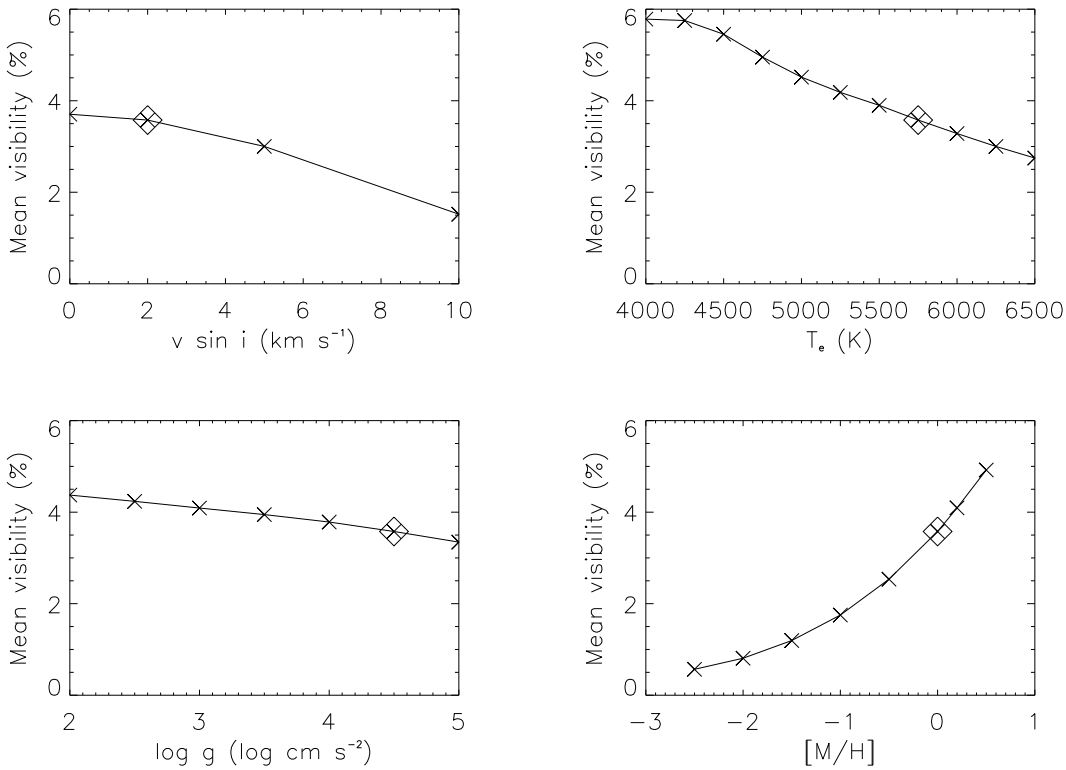

Figure 8-4. Fringe visibility dependency on different stellar parameters from simulations. Diamond marks nearest-to-solar grid point. For each free parameter, the remaining parameters are held fixed at the values nearest to solar in the simulation grid. Error bars are on the order of 0.016% in visibility.

pixel-wide spectrum and 5.0 pixel resolution element). This equation gives typical errors for all the simulated data of $\Delta\overline{\gamma} \sim 0.016\%$ for S/N=50.

It is clear from the plots that the mean visibility varies monotonically with $v \sin i$. If the other three stellar parameters are known, therefore, then it should be possible to estimate $v \sin i$ for the star. In order to gain a rough quantitative estimate of how precisely each parameter needs to be known to obtain a given precision in $v \sin i$, we can calculate an average gradient for visibility across $v \sin i$, $\partial\overline{\gamma}/\partial(v \sin i)$, and across each of the other three parameters, $\partial\overline{\gamma}/\partial X_i$, where X_i represents the parameters with the subscript i taking values $1, 2, 3$, corresponding to each parameter (not to be confused with the orbital inclination in $v \sin i$). We can then calculate:

$$\frac{\partial(v \sin i)}{\partial X_i} = \frac{\partial(v \sin i)}{\partial \overline{\gamma}} \cdot \frac{\partial \overline{\gamma}}{\partial X_i} \tag{8–18}$$

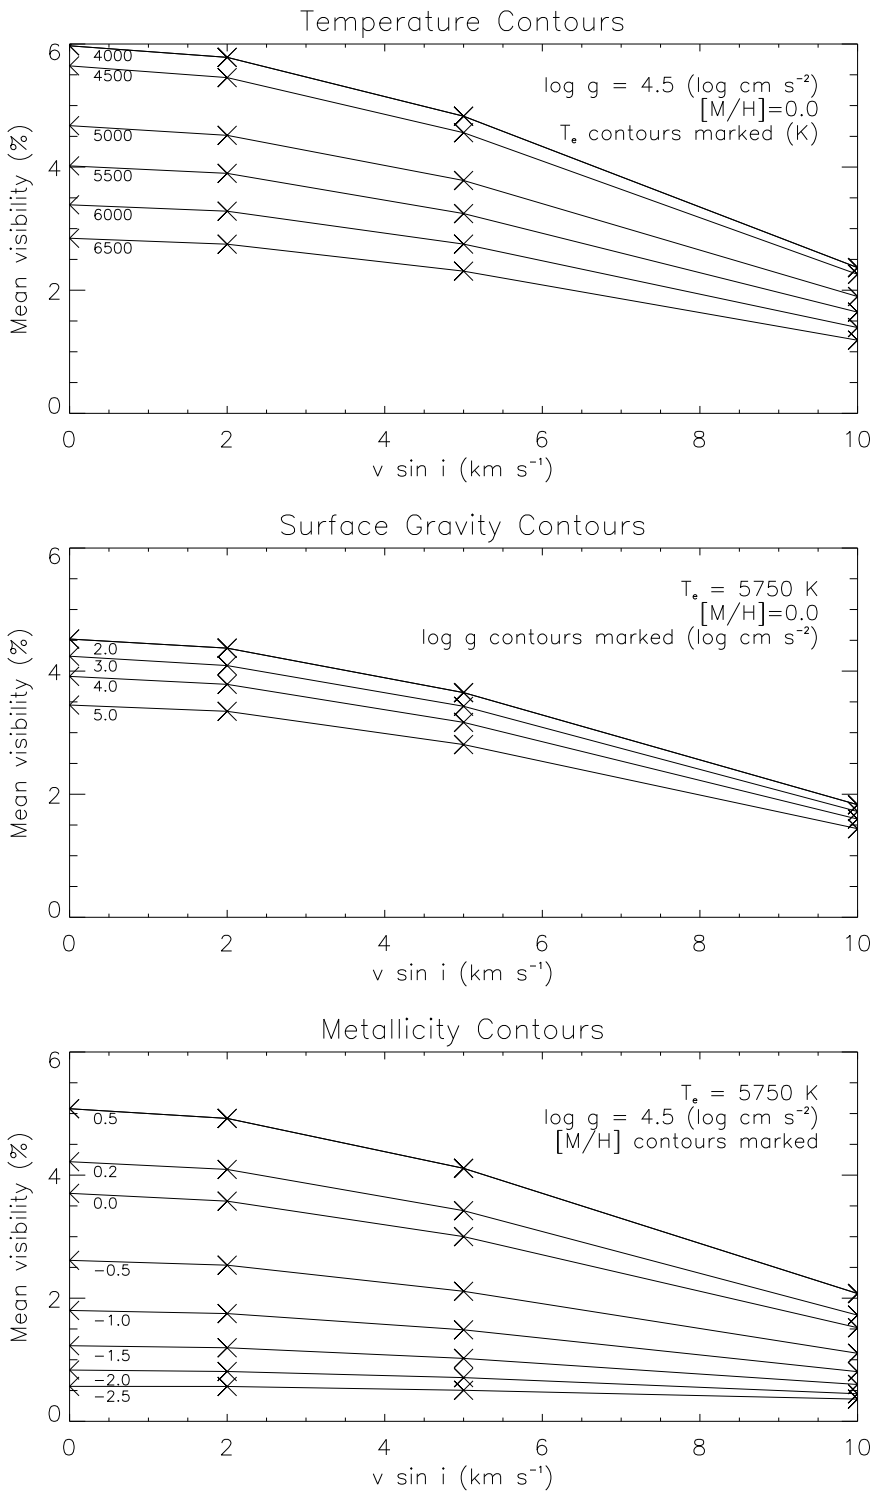

Figure 8-5. Contours of visibility vs. $v \sin i$ from simulations for different stellar parameters. Contours are shown for effective temperature, surface gravity and metallicity (in K, dex and dex, respectively). Each contour is labelled with its respective value down the left hand side. Error bars are on the order of 0.016% in visibility.

Table 8-1. Dependencies of mean visibility, $\overline{\gamma}$, and $v \sin i$ on stellar parameters X, where X represents $v \sin i$, T_e, $\log g$ and $[M/H]$.

Parameter, X	Units of X	$\partial\overline{\gamma}/\partial X$ (%/unit X)	$\partial(v \sin i)/\partial X$ ($\mathrm{km\,s^{-1}}$/unit X)
$v \sin i$	$\mathrm{km\,s^{-1}}$	-0.218	—
T_e	K	-0.00121	0.00556
$\log g$	dex (cgs)	-0.342	1.57
$[M/H]$	dex	1.45	-6.65

We calculate these gradients by taking cuts across the grid in all four parameter directions, centred at the nearest-to-solar grid point: the mean gradient is taken to be that of a straight line passing through the visibilities at the smallest and largest parameter values in each dimension (i.e., equivalently, the mean gradient of each of the plots in figure 8-4). The results of these calculations are summarised in table 8-1. We note that, apart from a difference in interferometer delay, most instrumental effects – which will most likely affect the spectrograph resolution and therefore the width of the response function w – will end up only scaling the visibility measurements (as seen in equation 8–14). Such a scaling in visibility will cancel out in equation 8–18, so that the values for $\partial(v \sin i)/\partial X$ should stay broadly the same. These dependency measures are of course only means across the grid, however, and therefore only provide rough estimates of the dependencies. A more exact treatment would consider the partial differentials at each point in the grid.

Finally, we note that we can estimate the effect of uncertainty in the visibility measurements due to photon noise by taking the reciprocal of the $v \sin i$ gradient from table 8-1, $\partial(v \sin i)/\partial\overline{\gamma} = -4.59\,\mathrm{km\,s^{-1}}$ per percent visibility. For the previously calculated error level of $\Delta\overline{\gamma} = 0.016\%$ at S/N = 50 (comparable to actual observed values as well as the simulations), this corresponds to $\Delta(v \sin i) \sim 0.07\,\mathrm{km\,s^{-1}}$, quite small compared to the typical errors expected due to uncertainties in the stellar parameters (particularly in the case of parameters determined using ET – see section 8.5).

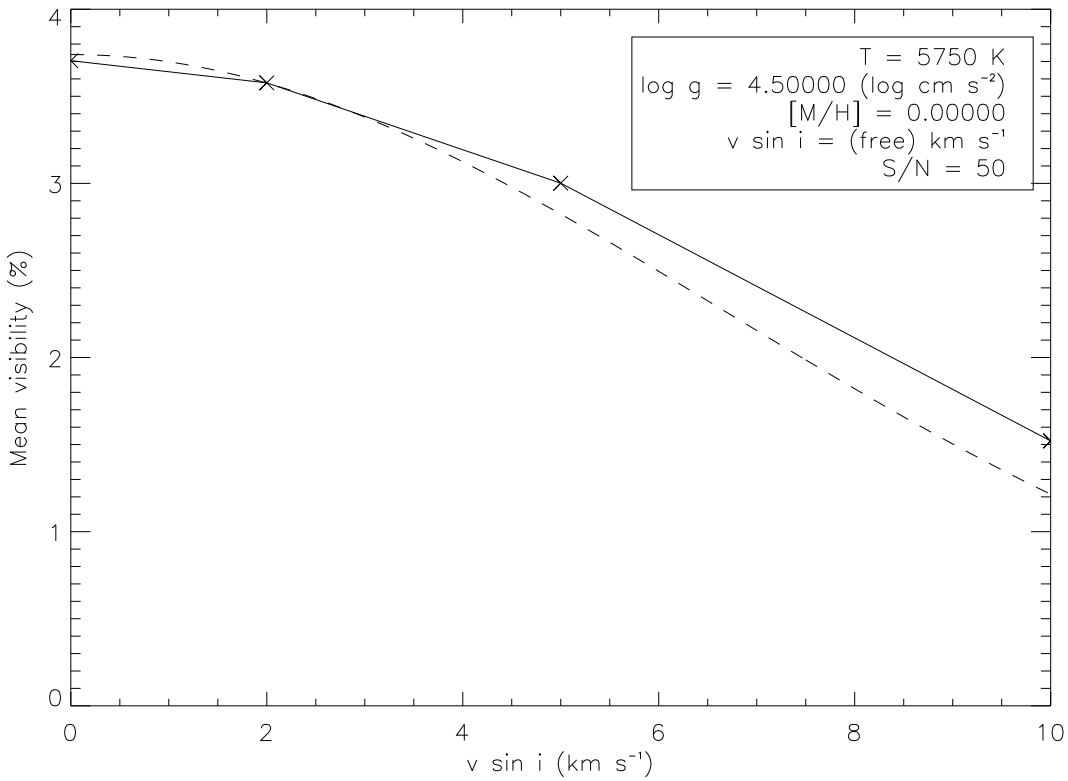

Figure 8-6. Simulated data overplotted on the theoretical curve from figure 8-3. Solid line with crosses: simulated data. Dashed line: analytical prediction. Stellar parameters for the simulations are for near-solar values, as listed in the legend.

8.3.3 Comparison with Theoretical Predictions

Figure 8-6 shows the theoretically predicted visibility–rotation-rate relation from section 8.2 overplotted on the corresponding data from the simulations for the nearest-to-solar parameter set. Following the previous arguments, we would probably expect some scaling difference between the two curves due to differences in shape of the spectrograph response profile, the approximate nature of the choice of representative line depth, and so forth. In fact, a remarkably good agreement is seen without any change in scaling. This may in fact somewhat coincidental given the approximate values used in the analytical calculations, but it nevertheless shows that a good estimate of the behaviour can be obtained analytically.

Since a true rotational broadening profile is rather different in shape from a Gaussian, we expect some deviation in the analytical model from reality. At small $v \sin i$, where

the intrinsic line profile dominates over the broadening profile (up to values of order the equivalent Doppler width of the intrinsic line, $\sim 4\,\mathrm{km\,s^{-1}}$), the shape should be primarily the Fourier transform of the intrinsic line profile, that is, approximately a Gaussian, so that the prediction should be reasonably good. At larger $v \sin i$, however, the shape should start to look more like the transform of the broadening profile. From figure 8-2 we can see that a more accurate broadening profile drops off more steeply and is narrower than the Gaussian approximation: we therefore can expect the *transform* to drop off more slowly, so that at large $v \sin i$, we should see a visibility for 'real' spectra larger than predicted by the analytical model. This behaviour is indeed seen in the overplotted curves in figure 8-6. A more accurate analytical curve should, in principle, be reasonably easy to obtain by performing a numerical Fourier transform of the real rotational broadening profile rather than using the Gaussian approximation.

8.4 Observations

A set of data taken in January 2006 with the single-object ET on the KPNO 2.1 m telescope was reduced by Roger Cohen using the standard single-object pipeline. A number of the targets observed during the run were chosen by Cohen specifically for the purpose of investigating the possibility of measuring stellar rotation rates. These included eight targets in particular which we investigate here: HD 105405, HD 134044, HD 3861, HD 130087, HD 3674, HD 12414, HD 17037 and daytime sky, a common calibration reference used as a proxy for the sun in ET observation runs. All of these targets had known stellar parameters, being listed in Valenti & Fischer (2005) (hereafter 'VF05'). Of the eight, a subset of four (HD 134044, HD 3861, HD 3674, and HD 17037) were particularly closely matched in T_e, $\log g$ and $[M/H]$, differing only in $v \sin i$. We use these targets here to perform a mean visibility analysis and compare with the simulations in section 8.3.

Simple mean fringe visibilities were obtained in the same way as for the simulations, as were error bars on the mean visibilities. To obtain a good match with the known stellar

parameters, the grid was interpolated using three-dimensional linear interpolation across the three parameters T_e, $\log g$ and $[M/H]$, to give cuts across $v \sin i$ at any chosen set of values for the first three parameters within the range of the grid.

Using the measurements listed in VF05, the well-matched target subset had mean values: $\overline{T}_e = 6197\,\mathrm{K}$ (range 6148–6223 K); $\overline{\log g} = 4.27\,\mathrm{dex}$ (range 4.23–4.32 dex); and $\overline{[M/H]} = 0.0525\,\mathrm{dex}$ (range 0.04–0.07 dex). Figure 8-7 shows a plot of mean visibility against $v \sin i$ for the subset of targets, where the values of $v \sin i$ and the corresponding errors are also taken from VF05. Overplotted is the curve expected from an interpolation of the simulations to match the mean values for this data set. Clearly the true visibility is over-predicted; a second overplotted curve shows the same simulations scaled in visibility by a factor of 0.59 to give a better match to the data: a good fit is now seen.

The reason for the mismatch may be due to several effects. There may be a difference in effective resolution between the simulations and the data – although the resolution elements (and hence the response function) should match well in width after filtering, the LSF is Gaussian in the simulations, but may be closer to a top-hat in the real data. It is also likely that the fringes in the real data are somewhat undersampled by the instrument point spread function (PSF) since the fringe density is rather high in this data set: undersampling can lead to a distinct reduction in visibility. Finally, there may be a certain amount of defocus or aberration in the instrument, which would again blur the fringes out. All of these effects can be expected to scale down the visibility, and it is therefore not unreasonable to have to allow a consequent scaling factor.

Figure 8-8 shows the full data set, including those from figure 8-7, overplotted with contours from the simulation grid interpolated to match the parameters for the various other targets. This time all the contours are scaled by the same factor of 0.59 as before. A good match is seen in all cases, consistent with the error bars; HD 12414 falls just outside its $v \sin i$ error bar but only by a statistically reasonable amount given the total number of data points.

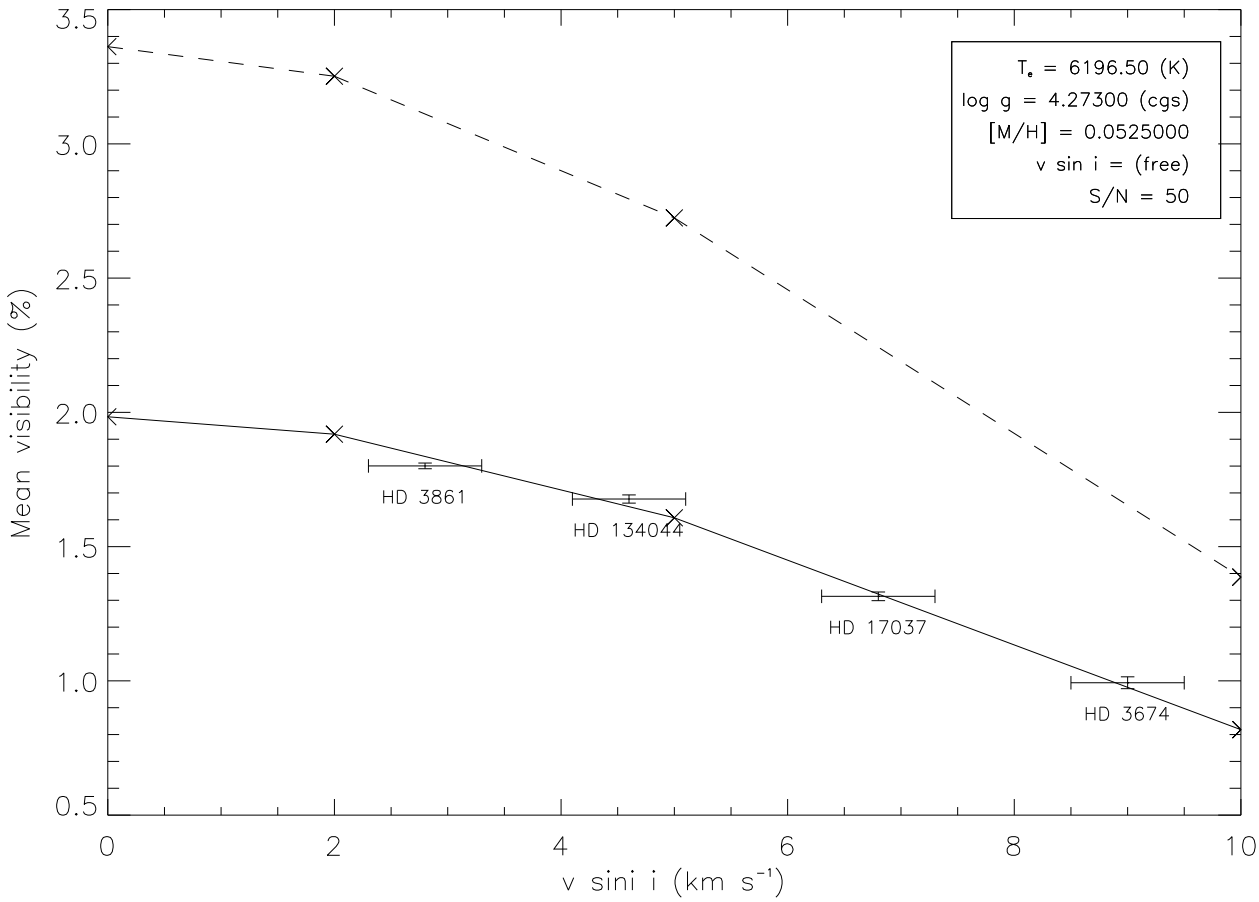

Figure 8-7. Real data and overplotted simulations for a well-matched subset of observed targets. The simulated data (dashed/solid line with crosses) are from grid interpolations at stellar parameters matching the mean parameters of the four targets (see legend). The dotted line represents the un-scaled simulations; the solid line represents the same simulations scaled in visibility by a factor of 0.59.

The only exception is the day-sky data point, which seems to lie a rather long way off the solar $v \sin i$ contour. The reason for this discrepancy is not yet clear, although it is not unreasonable to expect that the character of the solar spectrum reflected off the Earth's atmosphere may be rather different from a directly observed solar spectrum. If the spectrum is integrated over layers of the sky which are moving at different velocities due to varying wind speeds, for example, the spectral lines may become broadened. In fact the spectrum was also actually taken during twilight: it is known that when the Sun is low in the sky, differential sky transmission across the solar disc can affect the spectral lines quite

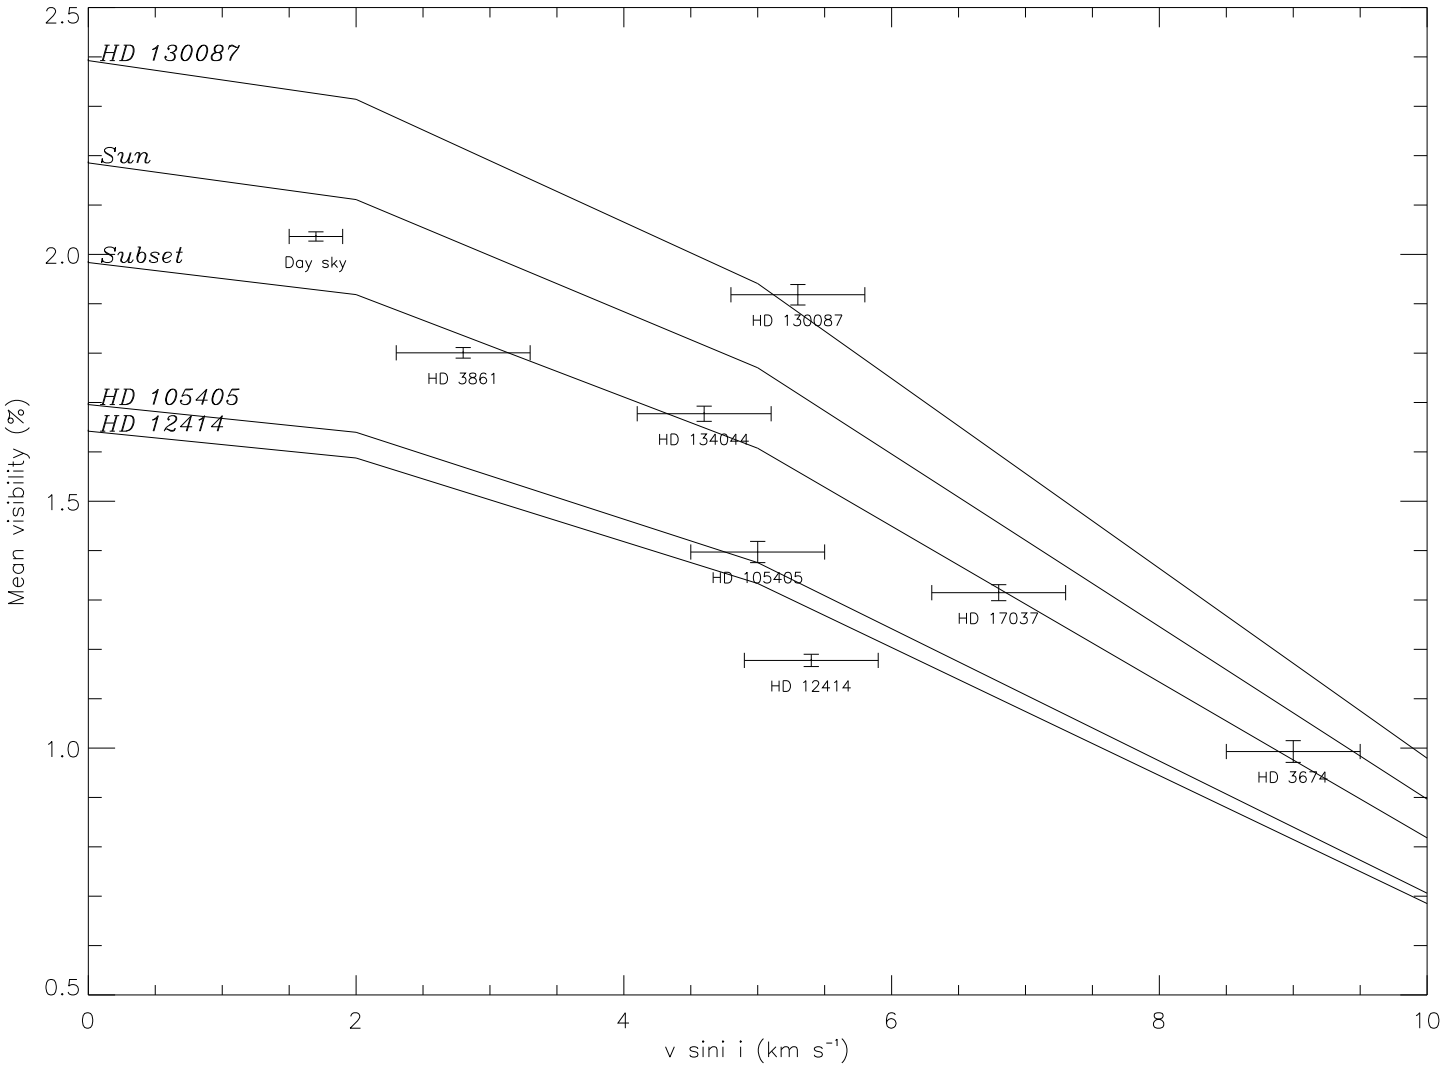

Figure 8-8. All eight targets with overplotted visibility contours from simulations. Contours are chosen to match the target stellar parameters as labelled in italics. *'Subset'* indicates contour for the mean parameter values used for the well-matched target subset as shown in figure 8-7.

strongly (Deming et al. 1987). Over the 5 minute length of the exposure, there is a good chance that the differential extinction was also changing rather rapidly, so it is perhaps not surprising that this data point does not match terribly well.

8.5 Conclusions

Rotational velocity should be measurable provided that the other stellar parameters, effective temperature, surface gravity, and metallicity, are known well enough. The important question is how well those parameters will be known in practice. Preliminary investigations by Roger Cohen into attempting to use simulated ET data to recover low resolution spectra and determine the three needed parameters suggests typical uncertainties of $\Delta T_e = 295\,\mathrm{K}$, $\Delta(\log g) = 0.48\,\mathrm{dex}$, and $\Delta[M/H] = 0.44\,\mathrm{dex}$ (for S/N = 200 per pixel, although he finds that in practice, the results with real data at S/N > 100 are actually a little better).

Given estimates of the uncertainties in the stellar parameters, ΔX_i, where we use the index i to refer to each of the parameters so that $X_1 = T_e$, $X_2 = \log g$, and $X_3 = [M/H]$, we can estimate the total uncertainty in the projected rotational velocity $\Delta(v \sin i)$. If we assume the ΔX_i are all independent, we can use the standard combination of errors:

$$[\Delta(v \sin i)]^2 = \sum_i [\Delta X_i]^2 \left[\frac{\partial(v \sin i)}{\partial \Delta X_i} \right]^2 . \tag{8–19}$$

Although the assumption of independence may not be ideal, this at least gives us a first-order estimate. Using Cohen's estimates gives a total uncertainty in the rotational velocity measurement, $\Delta(v \sin i)$, of $3.4\,\mathrm{km\,s^{-1}}$, dominated by the uncertainty in the metallicity measurement.

Given that these are preliminary results, there may be some room for improvement in the precision: experimenting with different ways of weighting the averaging of the fringe visibilities might help to reduce the metallicity dependence, for example. Since rotational broadening and broadening by the instrument LSF both conserve equivalent width, it should also be possible to determine equivalent widths of the spectral lines from the ET

spectra: with some assumptions on the intrinsic line widths, this could be used to estimate line depths and perhaps provide some extra information to remove degeneracy and further reduce the dependence of $v \sin i$ on the other stellar parameters.

Whether or not further improvements are possible, the results shown here nonetheless suggest that even with a basic estimate of stellar parameters from the ET instrument, it should be possible to make reasonable measurements of projected stellar rotational velocity.

CHAPTER 9
CONCLUSIONS AND CONTRIBUTIONS

9.1 Work Remaining

9.1.1 RV measurement

One of the most significant problems in the ET project remains the addition approximation non-linearity. With systematic errors on the scale of up to $100\,\mathrm{m\,s^{-1}}$ dominating, it is hard to try and tackle other smaller sources of error. The ideal answer would be to find a mathematically exact solution, or at least a more accurate approximation. Another very good solution would be 'combined beam' superposition of the reference spectrum, where a ThAr or tungsten-illuminated iodine spectrum is literally added to the stellar spectrum, for example by splicing two input fibres together into one. In this way the addition approximation is no longer an approximation – it is exact. Combined beam superposition would add photon noise to the stellar spectrum, although provided the two spectra are well balanced in flux, the additional noise may not be much worse than that due to flux loss when using an iodine absorption cell in the beam path.

Balancing fluxes between star and reference over 60 simultaneous targets, however, may present a challenge. Current efforts by Suvrath Mahadevan to model the error term using high-resolution iodine and synthetic stellar spectra show some promise. In this approach, a grid of corrections across velocity and stellar parameter space could in principle be applied. If instrument stability can be controlled well enough, simply running parallel reference spectra alongside the stellar spectra, or alternatively, bracketing stellar exposures in time with reference exposures may provide a third solution.

Another useful benefit of correcting the addition approximation error term, if it can be done, would be to allow one to meaningfully chop up the spectra into small, essentially wavelength-independent sections. At this point, the errors that would dominate each section because of the addition approximation are likely to be so large that such chopping would be largely fruitless. However, if chopping can be done, it would allow

for a comparison of the scatter in the results from section to section, giving a very useful measure of internal errors in the RV measurements (a trick commonly employed with traditional echelle-based RV surveys).

As surveying down to very faint magnitudes during bright-sky time continues, moonlight contamination is also likely to become an important issue. Approaches to remove moonlight contamination, either at the pre-processing stage, or by modelling it as an additional third spectrum along with the iodine/star superposition, have yet to be explored.

A more complete and thorough total error budget would still aid the ET project. The question of why the apparent noise floor remains in the day-sky data from the Keck ET, for example, is still open. As the software for producing simulated ET data by Suvrath Mahadevan becomes more sophisticated, it should become a lot easier to start assessing the level of various effects, simply by simulating one kind of instrument problem at a time and running the results through the pipeline.

Improvements still needed in the data reduction pipeline include moving away from the 'bulk-shift compensation' technique used in the algorithm for separating stellar and iodine reference components from the combined data, since the method has the potential to introduce systematic errors of its own. More reliable approaches are in principle relatively easy to implement, but have suffered from convergence problems in the χ^2 minimisation when finding best-fit solutions. With sufficiently accurate initial guesses and more refined searching of χ^2 space, convergence should not pose major problems. Correction of image distortion in the spectra, previously handled simply by chopping off bad sections, is also worth addressing: this would recover a not-insignificant number of wasted photons.

On a more pragmatic front, the pipeline needs to be moved forward from a development orientation to an application orientation. There is much to be streamlined, optimised, and made more user-friendly.

9.1.2 Stellar rotation

The $v \sin i$ project is still in its early stages. Measurements of more targets with known stellar parameters with both the Keck ET and the single-object ET will help calibrate the visibility function more accurately for the real instruments, and to expand the measurements out to a broader stellar parameter space. Ideally the technique will be proven with some blind measurements of $v \sin i$ for cases where the correct value is not initially known. Regarding the analytical predictions, it would be of interest to model a more accurate rotational broadening profile to see if a better match to the real data can be obtained. More investigation is also needed to fully understand why the visibilities from the single-object ET are so much lower than those from the simulations. A fuller treatment of the effects of photon noise would be of benefit; and as discussed previously, we plan to investigate methods of reducing the sensitivity of the $v \sin i$ measurements to changes in the other stellar parameters, whether by intelligent weighting of the visibility averaging to select those lines that are less affected, or by incorporating extra information that can be obtained from the relative flux measured from channel to channel in the spectrum.

9.2 Future Possibilities

It is interesting to consider other avenues for the DFDI approach. Could the interferometer comb itself be used as a fiducial reference, for example, instead of iodine or ThAr? Changes in the interferometer delay will shift the phase of the comb, and so it can in principle track instrument drift. The problem with this so far has lain in the symmetry of the comb: as a simple example, if the image on the detector were to drift in a direction exactly parallel to the comb, the stellar fringes would appear to shift in phase and in the dispersion direction, and yet the comb would appear not to have changed at all, leading to the possible but incorrect conclusion that the shift is wholly intrinsic to the star. If the image on the detector can reliably be stabilised to sufficient accuracy in either the slit or the dispersion direction (or both), then this knowledge will provide the

necessary information to break the degeneracy between stellar and instrument shift, and the intrinsic RV can be measured. This would be a major step forward in the field of Doppler RV measurement, allowing common-beam calibration with neither flux loss (as in iodine absorption) nor photon noise addition (as in ThAr superposition).

Another idea that has not yet been pursued is whether the DFDI technique could be used to search stars of earlier A–F spectral types. These stars have been almost completely avoided in RV searches because of their lack of strong stellar absorption lines, meaning that there is much less Doppler shift information encoded in their spectra (an exception is the work in Galland et al. (2005) and subsequent papers). Furthermore, these stars generally have higher rotational velocities, on the scale of 10–100 $\mathrm{km\,s^{-1}}$ (Gray 1992), again reducing the Doppler information content. The exceptionally strong Balmer lines in these spectra are too deep and intrinsically broadened too be of much use for precision measurements with a DFDI instrument. Nonetheless, there still remain many very shallow metal lines, and it is in principle possible to tune a DFDI instrument to be sensitive to these rotationally broadened lines. Early experiments with simulated data using a spectrum of Vega (an extreme A0V case) suggested that for a rotational velocity, $v \sin i$, of 25 $\mathrm{km\,s^{-1}}$, a precision of 50 $\mathrm{m\,s^{-1}}$ might be achievable. At a more extreme $v \sin i$ of 150 $\mathrm{km\,s^{-1}}$, the achievable precision might be around 340 $\mathrm{m\,s^{-1}}$. Although these are very much lower precisions than for later-type stars, it is still a velocity space that has been largely unexplored: several–Jupiter-mass planets in few-day orbits should still be detectable, as should brown dwarfs. Since there are hints that more massive stars may harbour more massive planets (Sato et al. 2007; Ida & Lin 2005), this would be an interesting area of investigation.

9.3 Summary and Current Status

The ET instruments have convincingly demonstrated the capacity of dispersed fixed-delay interferometry for exoplanet detection. Both the single-object ET at KPNO and the multi-object Keck ET at APO are able to routinely uncover the RV signals

of known exoplanets. The discovery of the new planet, ET-1, with the single-object ET proved a major milestone. At the time of writing, we have been able to push best short-term precisions down to the 2–3 $\mathrm{m\,s^{-1}}$ mark with bright reference stars using this instrument. As of mid-2007, the best short-term (2-day) noise floor for the Keck ET has reached down to $\leq 5.6\,\mathrm{m\,s^{-1}}$ on day-sky data, with 2-month RMS scatter typically at 11–13 $\mathrm{m\,s^{-1}}$. Typical photon limiting precisions for the Keck ET were then $\sim 50\,\mathrm{m\,s^{-1}}$ at visual magnitude $V = 11$–12, with the best around $\sim 30\,\mathrm{m\,s^{-1}}$; typical precisions at $V \sim 8$ were around 7–9 $\mathrm{m\,s^{-1}}$. Such precisions are adequate for finding planets with $M \sin i$ of order $1\,M_{\mathrm{J}}$ or more in few-day orbits (i.e. hot-Jupiters) down to $V = 12$. A handful of promising candidates from the pilot survey are currently being followed up as they become observable, and we hope to be able to report a few new planets soon. The precision is also more than adequate for uncovering stellar binary and brown dwarf companions, and we are now looking through our current results to catalogue such findings.

We expect to push the survey precision further, with a goal of $3.4\,\mathrm{m\,s^{-1}}$ at $V = 8$ and $21\,\mathrm{m\,s^{-1}}$ at $V = 12$ (Ge et al. 2007). The Keck ET instrument is still being upgraded: the interferometer and interferometer mounts have recently been replaced to improve stability; the optics have been realigned to improve throughput, aberration and slit-illumination homogeneity; and the slit width has been slightly reduced to try and mitigate the problems with aliasing of the interferometer comb.

The measurement of projected stellar rotation ($v \sin i$) by measuring the average fringe visibility for a stellar spectrum also appears promising. Even at the suggested precision levels of $\sim 3.4\,\mathrm{km\,s^{-1}}$ from the preliminary investigations, this can still be a useful byproduct of the ET surveys. It is reasonable to expect that there is room for improvement in this precision level.

As the ET instruments' overall precision and reliability improves in preparation for the planned MARVELS survey (Ge et al. 2007), we hope that the ET instruments will be

able to make a very significant contribution to the number of extrasolar planets found over the coming years.

9.4 Contributions

Work on a project like ET cannot be other than a group effort, and wherever there are substantial contributions beyond my own work, I have tried to explicitly credit people within the text. My own work within the ET project has consisted primarily in the development of the data reduction procedure, including the new reduction algorithms and coding of the majority of the pipeline (except where otherwise stated). I was also responsible for all reduction of the ET data from early on in the project, until Scott Fleming, Stephen Kane, and Brian Lee kindly took over the running of data through the pipeline at the start of 2006 for the KPNO data and in late 2006 for the Keck ET data.

The specific mathematical analysis of the principles of the DFDI approach outlined in this work is also my own, though some of the initial background material is drawn from several sources – again I have tried to state all such sources explicitly in the text. Apart from the photon limiting formula (equation 2–31) taken from Ge (2002), the discussion of analytical error analysis, including the precise form of the addition approximation cross-talk terms, is my own, and to my knowledge, new.

The investigation into the measurement of stellar rotation velocities represents a start on a new project that I have taken on. The semi-empirical investigation into the fringe visibility dependence on stellar parameters is, to my knowledge, new work. The relevant mathematical analysis presented here is also my own independent work, though I understand that Dr. Jian Ge has previously done some somewhat similar calculations, and had also previously proposed the idea of using fringe measurements for measuring stellar line widths and abundance studies. It should also be noted that Roger Cohen also performed some early closely related work on measuring stellar rotations (along with other stellar parameters) with the ET instruments, using a somewhat different approach (Cohen et al. 2006). I am indebted to Suvrath Mahadevan and Roger for their grid of

simulated ET spectra, which they had produced expressly for the purpose of investigating the measurement of stellar parameters with the ET instruments; and also to Roger for planning and taking the actual ET data which I was later able to use for the purposes of the stellar rotation project. Roger also ran those particular data through the pipeline to produce ready-made whirl data from which I was able to produce all the mean fringe visibility information.

The design and construction of the ET instruments were lead by other members of the ET team, but I was able to assist with all the major instrument installation and engineering runs at Kitt Peak and Apache Point, as well as the early prototyping run at the Hobby Eberly Telescope. I have also been involved in discussions and been able to provide input during all the stages of the conceptual development of the instruments (the initial concept of the multi-point for the interferometer phase-locker, for example, was a particular contribution I was able to make). Finally, I have spent a total of somewhere between 100–200 nights observing with the ET instruments in their various incarnations at both the KPNO 2.1 m telescope and the Sloan telescope.

None of this work, however, would have been possible without the hard work of all the ET team members, past and present, working on the instrumentation, observing, and providing much input on the data reduction software development: Dr. Jian Ge, the P.I., who started the project and has guided it tirelessly throughout; Suvrath Mahadevan, my fellow classmate and partner in crime, whose simulated data and continual sharing of ideas have been of fundamental importance; Brian Lee and Scott Fleming, who helped alleviate the burden of running the pipeline so I could work on other things; Craig Warner, who wrote the ET electronic log software, the graphical interface for the pipeline, and replaced the early IRAF parts of the reduction procedure with improved IDL code; Stephen Kane who has worked extensively on the analysis of the RV data from the ET instruments, and also helped with running some of the ET data through the pipeline; Roger Cohen and Mari-Cruz Galvez who have worked on the measurement of stellar parameters with ET;

Curtis DeWitt and Dan McDavitt, who worked on much of the early instrumentation and instrument control; engineers Xiaoke Wan, Bo Zhao, Abishek Hariharan, Jerry Friedman, Deqing Ren, Mike Zugger and John Groot; Andrew Vanden Heuvel, Pengcheng Guo and Justin Crepp, who have put in much observing time and provided much practical input for the ET project. Almost all have at one time or another run the observing gauntlet with the various incarnations of the ET instrument – some far beyond the call of good health!

I would also like to thank Bill Cochran and Dave Erskine for some very helpful input. This list of those who have contributed to the project is not intended to be exhaustive or exclusive, and I apologise for any unintended omissions here and elsewhere in this work.

9.5 Credits

Parts of this work have been reproduced from the Astrophysical Journal and Astrophysical Journal Letters: Figures 1, 2, and 3, tables 1 and 2, and some adapted parts of main text, from "First Planet Confirmation with a Dispersed Fixed-Delay Interferometer," (van Eyken et al. 2004a); and Figures 5, 8, and 9 (top), and tables 1, 3 and 5, from "The First Extrasolar Planet Discovered with a New-Generation High-Throughput Doppler Instrument" (Ge et al. 2006a) (©2004 and 2006: The American Astronomical Society. All rights reserved. Printed in U. S. A.).

The author has been a Visiting Astronomer at the Kitt Peak National Observatory, National Optical Astronomy Observatory, which is operated by the Association of Universities for Research in Astronomy, Inc. (AURA) under cooperative agreement with the National Science Foundation. This research has made use of the SIMBAD database, operated at CDS, Strasbourg, France. The Hobby-Eberly Telescope (HET) is a joint project of the University of Texas at Austin, Pennsylvania State University, Stanford University, Ludwig-Maximillians-Universität München, and Georg-August-Universität Göttingen. The HET is named in honour of its principal benefactors, William P. Hobby and Robert E. Eberly. IRAF is distributed by the National Optical Astronomy Observatories, which are operated by the Association of Universities for Research in Astronomy, Inc.,

under cooperative agreement with the National Science Foundation. This work has made use of the IDL Astronomy User's Library (Landsman 1993),[1] and Craig Markwardt's IDL library.[2]

The author gratefully acknowledges support from the JPL Michelson Fellowship program funded by NASA; from the fellowship provided by Penn State University; from the SPIE scholarship program; and from Kitt Peak National Observatory for travel support for many nights of observing at the 2.1 m telescope.

The multiple-object ET instrument was supported by the W. M. Keck Foundation and the University of Florida. The development of the single-object ET was supported by NSF, NOAO and the Pennsylvania State University.

[1] See also http://idlastro.gsfc.nasa.gov/ (last accessed Oct 2007).

[2] http://cow.physics.wisc.edu/~craigm/idl/idl.html, (last accessed Oct 2007).

APPENDIX A
THE SPECTROGRAPH RESPONSE FUNCTION AND THE LSF

The response function w due to the spectrograph optics is very closely related to the instrument line spread function (LSF). If we assume the line spread function is approximately identical in form at closely separated channels x on the detector, where x represents pixel column number, then we can define the LSF as $L(x, x_0(\lambda_0)) = L_0(x - x_0(\lambda_0))$ where L represents the normalised envelope of intensity spread across pixels x on the detector due to monochromatic light of wavelength λ. x_0 is the central position of the LSF for wavelength λ_0, and in general $x(\lambda)$ represents the wavelength calibration mapping wavelength to detector position.

The response function at an infinitesimally wide position on the detector is given by writing the contribution from each overlapping LSF at that position. The contribution from wavelength λ_0 at position x_1 is given by $W(\lambda_0, x_1) = L(x_1, x_0(\lambda_0)) = L_0(x_1 - x_0(\lambda_0))$. Therefore as a continuous function of general wavelength λ, we can write the response function at position x_1 as

$$W(\lambda, x_1) = L_0(x_1 - x(\lambda)). \tag{A–1}$$

We see that this is really just the LSF reversed (since the x term is now negative).

Now we extend this to the total contribution at pixel number j where j is an integer. Let $t_j(x)$ represent the pixel response function, describing the normalised throughput of the pixel across its width as a function of x. Then we can write $W(\lambda, x')t_j(x')\,\mathrm{d}x' = L_0(x' - x(\lambda))t_j(x')\,\mathrm{d}x'$. Summing over all x', we have the full response function w_j at pixel column j, given by:

$$w_j(\lambda) = \int L_0(x' - x(\lambda))t_j(x')\,\mathrm{d}x', \tag{A–2}$$

i.e., essentially a convolution of the response function with the pixel response function. To the extent that the width of the pixel is narrow compared to the LSF (i.e. that the image is well over-sampled), then to a reasonable approximation, t_j is close to a delta function, and the instrument response function at position x is approximately just the reversed

LSF. Analogous arguments can be followed in wavenumber (κ) space instead of wavelength space, simply replacing λ with κ to obtain the same exactly the same results as a function of κ.

APPENDIX B
FRINGE FORMATION: AN ALTERNATIVE VIEWPOINT

We can view the formation of the fringes as given by equation 2–9 in another way. If we write out the Fourier transforms explicitly as integrals, we obtain:

$$\boldsymbol{\gamma}_j \;=\; \frac{\mathcal{F}[P_\kappa w_{\kappa j}]_{d=d_0}}{\mathcal{F}[P_\kappa w_{\kappa j}]_{d=0}} \;=\; \frac{\int P(\kappa) w_j(\kappa) \mathrm{e}^{-2\pi\kappa d_0}\,\mathrm{d}\kappa}{\int P(\kappa) w_j(\kappa)\,\mathrm{d}\kappa}, \tag{B–1}$$

where we have dropped the κ subscripts. If we assume that the spectrograph response function at channel j, $w_j(\kappa)$, is uniform across the whole spectrum (i.e. for all j), then we can express it as a wavenumber-shifted version of a 'universal' spectrograph response function, w_0, centered at $\kappa = 0$, shifted so that its centre is at the wavenumber of the channel in question, κ_j. From appendix A, we know that the spectrograph response function is just the reverse of the LSF, so we can write:

$$w_j = w_0(\kappa - \kappa_j) = L_0(\kappa_j - \kappa), \tag{B–2}$$

where L_0 represents the 'universal' LSF, also centred at $\kappa = 0$. We can therefore rewrite equation B–1:

$$\boldsymbol{\gamma}_j = \frac{\int P(\kappa) \mathrm{e}^{-2\pi\kappa d_0} L_0(\kappa_j - \kappa)\,\mathrm{d}\kappa}{\int P(\kappa) L_0(\kappa_j - \kappa)\,\mathrm{d}\kappa}. \tag{B–3}$$

The integrals are just convolutions over the variable κ, themselves functions of κ_j:

$$\boldsymbol{\gamma}_j = \frac{[(P(\kappa)\mathrm{e}^{-2\pi\kappa d_0}) \otimes L_0(\kappa)](\kappa_j)}{[P(\kappa) \otimes L_0(\kappa)](\kappa_j)}. \tag{B–4}$$

If $\boldsymbol{\gamma}_j$ is the complex visibility, then the intensity as a function of delay and wavenumber κ_j must be given by:

$$\begin{aligned} I(\kappa_j, d_0) &= I_\mathrm{t}(\kappa_j) + I_\mathrm{t}(\kappa_j)\Re\{\boldsymbol{\gamma}_j\} \\ &= I_\mathrm{t}(\kappa_j) + I_\mathrm{t}(\kappa_j)\,\Re\left\{\frac{(P(\kappa)\mathrm{e}^{-2\pi\kappa d_0}) \otimes L_0(\kappa)}{P(\kappa) \otimes L_0(\kappa)}\right\}, \end{aligned} \tag{B–5}$$

where I_t represents the total flux contributing to the channel, and we have dropped the explicit notation of the (κ_j) functional dependence for the convolutions. Since the

denominator of the expression for $\boldsymbol{\gamma}_j$ still also represents the total flux in the channel, we can substitute it for $I_{\rm t}$. The total flux must be a real quantity, as are the line spread function, L_0, and the input spectrum, $P(\kappa)$, so we can take them outside the $\Re\{\}$ operator, giving us:

$$\begin{aligned} I(\kappa_j, d_0) &= P(\kappa) \otimes L_0(\kappa) \quad + \quad (P(\kappa)\,\Re\{\mathrm{e}^{-2\pi\kappa d_0}\}) \otimes L_0(\kappa) \\ &= P(\kappa) \otimes L_0(\kappa) \quad + \quad (P(\kappa)\cos(-2\pi\kappa d_0)) \otimes L_0(\kappa). \end{aligned} \tag{B–6}$$

Finally, because convolution is distributive, we can write:

$$\begin{aligned} I(\kappa_j, d_0) &= (P(\kappa)(1 + \cos(-2\pi\kappa d_0)) \otimes L_0(\kappa) \\ &= [P(\kappa) \cdot T(\kappa)] \otimes L_0(\kappa), \end{aligned} \tag{B–7}$$

where we define $T(\kappa) \equiv 1 + \cos(-2\pi\kappa_j d_0)$, which can be thought of as the interferometer transmission function, equivalent to the interferogram that would be obtained for pure white light and an infinite resolution spectrograph (exactly as in equation 2–8). In other words, we have simply the input spectrum multiplied with the interferometer transmission function, and then convolved with the LSF due to the spectrograph. Thinking in two dimensions, to match the wide-slit format of the actual ET spectra, we can replace the LSF to its (in some senses) two dimensional equivalent, the instrument point spread function (PSF). Exactly the same results can be derived in frequency space, simply by substituting ν for κ.

This way of looking at fringe formation is the approach used by Erskine (2003) and followed by Mahadevan (2006). The same formalism is also used by Mahadevan to produce simulated DFDI spectra.

REFERENCES

Armitage, P. J. 2007, ApJ, 665, 1381

Baranne, A., Queloz, D., Mayor, M., Adrianzyk, G., Knispel, G., Kohler, D., Lacroix, D., Meunier, J.-P., Rimbaud, G., & Vin, A. 1996, A&AS, 119, 373

Barker, L. M., & Hollenbach, R. E. 1972, J. Appl. Phys., 43, 4669

Barker, L. M., & Schuler, K. W. 1974, J. Appl. Phys., 45, 3692

Beckers, J. M., & Brown, T. M. 1978, in Future Solar Optical Observations Needs and Constraints, ed. G. Godoli, 189

Beuzit, J.-L., Mouillet, D., Oppenheimer, B. R., & Monnier, J. D. 2007, in Protostars and Planets V, ed. B. Reipurth, D. Jewitt, & K. Keil, 717

Brown, T. M., Noyes, R. W., Nisenson, P., Korzennik, S. G., & Horner, S. 1994, PASP, 106, 1285

Butler, R. P., & Marcy, G. W. 1996, ApJ, 464, L153

Butler, R. P., Marcy, G. W., Williams, E., Hauser, H., & Shirts, P. 1997, ApJ, 474, L115

Butler, R. P., Marcy, G. W., Williams, E., McCarthy, C., Dosanjh, P., & Vogt, S. S. 1996, PASP, 108, 500

Butler, R. P., Wright, J. T., Marcy, G. W., Fischer, D. A., Vogt, S. S., Tinney, C. G., Jones, H. R. A., Carter, B. D., Johnson, J. A., McCarthy, C., & Penny, A. J. 2006, ApJ, 646, 505

Campbell, B., & Walker, G. A. H. 1979, PASP, 91, 540

Carroll, B. W., & Ostlie, D. A. 1996, An Introduction to Modern Astrophysics (Reading, MA: Addison-Wesley)

Castelli, F., & Kurucz, R. L. 2003, in IAU Symp., 210, Modelling of Stellar Atmospheres, ed. N. Piskunov, W. W. Weiss, & D. F. Gray, 20P

Charbonneau, D., Brown, T. M., Burrows, A., & Laughlin, G. 2007, in Protostars and Planets V, ed. B. Reipurth, D. Jewitt, & K. Keil, 701

Cochran, W. D., & Hatzes, A. P. 1993, in ASP Conf. Ser., 36, Planets Around Pulsars, ed. J. A. Phillips, S. E. Thorsett, & S. R. Kulkarni, 267

Cohen, R., Mahadevan, S., & Ge, J. 2006, in BAAS, 38, 1105

Connes, P. 1985, Ap&SS, 110, 211

Correia, A. C. M., Udry, S., Mayor, M., Laskar, J., Naef, D., Pepe, F., Queloz, D., & Santos, N. C. 2005, A&A, 440, 751

da Silva, R., Udry, S., Bouchy, F., Mayor, M., Moutou, C., Pont, F., Queloz, D., Santos, N. C., Ségransan, D., & Zucker, S. 2006, A&A, 446, 717

David, M., & Verschueren, W. 1995, A&AS, 111, 183

Deming, D., Espenak, F., Jennings, D. E., Brault, J. W., & Wagner, J. 1987, ApJ, 316, 771

Douglas, N. G. 1997, PASP, 109, 151

Edser, E., & Butler, C. P. 1898, Philos. Mag., 46, 207

Erskine, D. J. 2003, PASP, 115, 255

Erskine, D. J., Edelstein, J., Lloyd, J., & Muirhead, P. 2006, in Proc. SPIE, 6269, Ground-Based and Airborne Instrumentation for Astronomy, ed. I. S. McLean & M. Iye, 62692P

Erskine, D. J., & Ge, J. 2000, in ASP Conf. Ser., 195, Imaging the Universe in Three Dimensions, ed. W. van Breugel & J. Bland-Hawthorn, 501

Fischer, D. A., & Valenti, J. 2005, ApJ, 622, 1102

Forrest, A. K., & Ring, J. 1978, in High Resolution Spectrometry, ed. M. Hack, 462

Frandsen, S., Douglas, N. G., & Butcher, H. R. 1993, A&A, 279, 310

Galland, F., Lagrange, A.-M., Udry, S., Chelli, A., Pepe, F., Queloz, D., Beuzit, J.-L., & Mayor, M. 2005, A&A, 443, 337

Gatewood, G., & Eichhorn, H. 1973, AJ, 78, 769

Ge, J. 2002, ApJ, 571, L165

Ge, J., Erskine, D. J., & Rushford, M. 2002, PASP, 114, 1016

Ge, J., Mahadevan, S., Lee, B., Wan, X., Zhao, B., van Eyken, J., Kane, S., Guo, P., Ford, E., Fleming, S., Crepp, J., Cohen, R., Groot, J., Galvez, M. C., Liu, J., Agol, E., Gaudi, S., Ford, H., Schneider, D., Seager, S., Weinberg, D., & Eisenstein, D. 2007, in Extreme Solar Systems, ed. D. Fischer, F. Rasio, S. Thorsett, & A. Wolszczan, ASP Conf. Ser., in press

Ge, J., Mahadevan, S., van Eyken, J., Dewitt, C., & Shaklan, S. 2003a, in ASP Conf. Ser., 294, Scientific Frontiers in Research on Extrasolar Planets, ed. D. Deming & S. Seager, 573

Ge, J., van Eyken, J., Mahadevan, S., Dewitt, C., Cohen, R., Fleming, S., Guo, P., vanden Heuval, A., Kane, S., Dermott, S., Kron, R., Strauss, M., Gunn, J., Seager, S., & Schneider, D. 2005, in Protostars and Planets V, 8410

Ge, J., van Eyken, J., Mahadevan, S., DeWitt, C., Kane, S. R., Cohen, R., Vanden Heuvel, A., Fleming, S. W., Guo, P., Henry, G. W., Schneider, D. P., Ramsey, L. W., Wittenmyer, R. A., Endl, M., Cochran, W. D., Ford, E. B., Martín, E. L., Israelian, G., Valenti, J., & Montes, D. 2006a, ApJ, 648, 683

Ge, J., van Eyken, J. C., Mahadevan, S., DeWitt, C., Ramsey, L. W., Shaklan, S. B., & Pan, X. 2003b, in Proc. SPIE, 4838, Interferometry for Optical Astronomy II, ed. W. A. Traub, 503

Ge, J., Wan, X., Zhao, B., Hariharan, A., Mahadevan, S., van Eyken, J., Guo, P., McDavitt, D., DeWitt, C., Cohen, R., Fleming, S. W., Kane, S., Crepp, J., & Shaklan, S. 2006b, in Proc. SPIE, 6269, Ground-Based and Airborne Instrumentation for Astronomy, ed. I. S. McLean & I. Masanori, 62692E

Geake, J. E., Ring, J., & Woolf, N. J. 1959, MNRAS, 119, 616

Goodman, J. W. 1985, Statistical Optics (New York: Wiley)

Gorskii, S. M., & Lebedev, V. P. 1977, Izv. Krym. Astrofizika Obs., 57, 228

Gosling, J., Joy, B., & Stelle, G. 1996, The Java language specification (The Java Series, Reading, MA: Addison-Wesley, —c1996)

Gould, A., Udalski, A., An, D., Bennett, D. P., Zhou, A.-Y., Dong, S., Rattenbury, N. J., Gaudi, B. S., Yock, P. C. M., Bond, I. A., Christie, G. W., Horne, K., Anderson, J., Stanek, K. Z., DePoy, D. L., Han, C., McCormick, J., Park, B.-G., Pogge, R. W., Poindexter, S. D., Soszyński, I., Szymański, M. K., Kubiak, M., Pietrzyński, G., Szewczyk, O., Wyrzykowski, Ł., Ulaczyk, K., Paczyński, B., Bramich, D. M., Snodgrass, C., Steele, I. A., Burgdorf, M. J., Bode, M. F., Botzler, C. S., Mao, S., & Swaving, S. C. 2006, ApJ, 644, L37

Gray, D. F. 1992, The Observation and Analysis of Stellar Photospheres (2nd ed.) (Cambridge; New York: Cambridge University Press)

Gray, R. O., & Corbally, C. J. 1994, AJ, 107, 742

Gray, R. O., Corbally, C. J., Garrison, R. F., McFadden, M. T., & Robinson, P. E. 2003, AJ, 126, 2048

Griffin, R., & Griffin, R. 1973, MNRAS, 162, 243

Hajian, A. R., Behr, B. B., Cenko, A. T., Olling, R. P., Mozurkewich, D., Armstrong, J. T., Pohl, B., Petrossian, S., Knuth, K. H., Hindsley, R. B., Murison, M., Efroimsky, M., Dantowitz, R., Kozubal, M., Currie, D. G., Nordgren, T. E., Tycner, C., & McMillan, R. S. 2007, ApJ, 661, 616

Hanisch, R. J., Farris, A., Greisen, E. W., Pence, W. D., Schlesinger, B. M., Teuben, P. J., Thompson, R. W., & Warnock, A., III. 2001, A&A, 376, 359

Harlander, J., Reynolds, R. J., & Roesler, F. L. 1992, ApJ, 396, 730

Harvey, J., & The GONG Instrument Team. 1995, in ASP Conf. Ser., 76, GONG 1994. Helio- and Astro-Seismology from the Earth and Space, ed. R. K. Ulrich, E. J. Rhodes, Jr., & W. Dappen, 432

Heacox, W. D. 1986, AJ, 92, 219

Hilliard, R. L., & Shepherd, G. G. 1966, J. Opt. Soc. Am. (1917-1983), 56, 362

Ida, S., & Lin, D. N. C. 2004a, ApJ, 604, 388

Ida, S., & Lin, D. N. C. 2004b, ApJ, 616, 567

Ida, S., & Lin, D. N. C. 2005, ApJ, 626, 1045

Jacquinot, P. 1960, Rep. Prog. Phys., 23, 267

Kane, S. R., Collier Cameron, A., Horne, K., James, D., Lister, T. A., Pollacco, D. L., Street, R. A., & Tsapras, Y. 2004, MNRAS, 353, 689

Kley, W., Peitz, J., & Bryden, G. 2004, A&A, 414, 735

Landsman, W. B. 1993, in ASP Conf. Ser., 52, Astronomical Data Analysis Software and Systems II, ed. R. J. Hanisch, R. J. V. Brissenden, & J. Barnes, 246

Lawson, P. R. 2000, in Principles of Long Baseline Stellar Interferometry, ed. P. R. Lawson (Pasadena: NASA JPL), 113

Lomb, N. R. 1976, Ap&SS, 39, 447

Mahadevan, S. 2006, Ph.D. thesis, University of Florida, Gainesville, Florida

Mahadevan, S., Ge, J., van Eyken, J., Cohen, R., Dewitt, C., & Wan, X. 2005, in Protostars and Planets V, 8224

Mahadevan, S., Ge, J., van Eyken, J. C., DeWitt, C., & Shaklan, S. B. 2003, in Proc. SPIE, 5170, Techniques and Instrumentation for Detection of Exoplanets, ed. D. R. Coulter, 184

Marcy, G., Butler, R. P., Fischer, D., Vogt, S., Wright, J. T., Tinney, C. G., & Jones, H. R. A. 2005, Prog. Theor. Phys. Supp., 158, 24

Marcy, G. W., & Butler, R. P. 1992, PASP, 104, 270

Marcy, G. W., & Butler, R. P. 1996, ApJ, 464, L147

Marcy, G. W., Butler, R. P., Fischer, D. A., Laughlin, G., Vogt, S. S., Henry, G. W., & Pourbaix, D. 2002, ApJ, 581, 1375

Marcy, G. W., Butler, R. P., Williams, E., Bildsten, L., Graham, J. R., Ghez, A. M., & Jernigan, J. G. 1997, ApJ, 481, 926

Martínez Fiorenzano, A. F., Gratton, R. G., Desidera, S., Cosentino, R., & Endl, M. 2005, A&A, 442, 775

Mayor, M., & Queloz, D. 1995, Nature, 378, 355

McArthur, B. E., Endl, M., Cochran, W. D., Benedict, G. F., Fischer, D. A., Marcy, G. W., Butler, R. P., Naef, D., Mayor, M., Queloz, D., Udry, S., & Harrison, T. E. 2004, ApJ, 614, L81

McCarthy, C. 1995, Master's thesis, San Fransisco State University, San Fransisco, California

McMillan, R. S., Moore, T. L., Perry, M. L., & Smith, P. H. 1993, ApJ, 403, 801

McMillan, R. S., Moore, T. L., Perry, M. L., & Smith, P. H. 1994, Ap&SS, 212, 271

Michelson, A. A. 1903, Light Waves and their Uses (Chicago: University of Chicago Press)

Morrison, D. 2001, Astrobiology, 1, 3

Mosser, B., Maillard, J.-P., & Bouchy, F. 2003, PASP, 115, 990

Mullally, F., & Winget, D. 2006, in BAAS, 38, 1129

Mullally, F., Winget, D. E., & Kepler, S. O. 2006, in ASP Conf. Ser., 352, New Horizons in Astronomy: Frank N. Bash Symposium, ed. S. J. Kannappan, S. Redfield, J. E. Kessler-Silacci, M. Landriau, & N. Drory, 265

Naef, D., Mayor, M., Beuzit, J. L., Perrier, C., Queloz, D., Sivan, J. P., & Udry, S. 2004, A&A, 414, 351

Pepe, F., Mayor, M., Delabre, B., Kohler, D., Lacroix, D., Queloz, D., Udry, S., Benz, W., Bertaux, J.-L., & Sivan, J.-P. 2000, in Proc. SPIE, 4008, Optical and IR Telescope Instrumentation and Detectors, ed. M. Iye & A. F. Moorwood, 582

Pepe, F., Mayor, M., Queloz, D., Benze, W., Bertaux, J.-L., Bouchy, F., Lovis, C., Mordasin, C., Santos, N., Sivan, J.-P., & Udry, S. 2005, Messenger, 120, 22

Perryman, M. A. C. 2000, Rep. Prog. Phys., 63, 1209

Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, Numerical recipes in C. The Art of Scientific Computing (2nd ed.) (Cambridge: Cambridge University Press)

Raymond, S. N., Mandell, A. M., & Sigurdsson, S. 2006, Science, 313, 1413

Sato, B., Izumiura, H., Toyota, E., Kambe, E., Takeda, Y., Masuda, S., Omiya, M., Murata, D., Itoh, Y., Ando, H., Yoshida, M., Ikoma, M., Kokubo, E., & Ida, S. 2007, ApJ, 661, 527

Scargle, J. D. 1982, ApJ, 263, 835

Seager, S., Turner, E. L., Schafer, J., & Ford, E. B. 2005, Astrobiology, 5, 372

Segura, A., Kasting, J. F., Meadows, V., Cohen, M., Scalo, J., Crisp, D., Butler, R. A. H., & Tinetti, G. 2005, Astrobiology, 5, 706

Shepherd, G. G., Gault, W. A., Miller, D. W., Pasturczyk, Z., Johnston, S. F., Kosteniuk, P. R., Haslett, J. W., Kendall, D. J. W., & Wimperis, J. R. 1985, Appl. Opt., 24, 1571

Silvotti, R., Schuh, S., Janulis, R., Solheim, J.-E., Bernabei, S., Østensen, R., Oswalt, T. D., Bruni, I., Gualandi, R., Bonanno, A., Vauclair, G., Reed, M., Chen, C.-W., Leibowitz, E., Paparo, M., Baran, A., Charpinet, S., Dolez, N., Kawaler, S., Kurtz, D., Moskalik, P., Riddle, R., & Zola, S. 2007, Nature, 449, 189

Sozzetti, A. 2005, PASP, 117, 1021

Struve, O. 1952, Observatory, 72, 199

Stumpff, P. 1980, A&AS, 41, 1

Tody, D. 1986, in Proc. SPIE, 627, Instrumentation in astronomy VI, ed. D. L. Crawford (Bellingham, WA: SPIE), 733

Tody, D. 1993, in ASP Conf. Ser., 52, Astronomical Data Analysis Software and Systems II, ed. R. J. Hanisch, R. J. V. Brissenden, & J. Barnes, 173

Tremaine, S., & Zakamska, N. L. 2004, in AIP. Conf. Proc., 713, The Search for Other Worlds, ed. S. S. Holt & D. Deming (New York: AIP), 243

Udry, S., Fischer, D., & Queloz, D. 2007, in Protostars and Planets V, ed. B. Reipurth, D. Jewitt, & K. Keil, 685

Udry, S., Mayor, M., Naef, D., Pepe, F., Queloz, D., Santos, N. C., Burnet, M., Confino, B., & Melo, C. 2000, A&A, 356, 590

Valenti, J. A., & Fischer, D. A. 2005, ApJS, 159, 141

van de Kamp, P. 1963, AJ, 68, 515

van Eyken, J. C., Ge, J., Mahadevan, S., & DeWitt, C. 2004a, ApJ, 600, L79

van Eyken, J. C., Ge, J., Mahadevan, S., Dewitt, C., Cohen, R., vanden Heuval, A., Fleming, S., Henry, G., & Martin, E. 2005, in Protostars and Planets V, 8579

van Eyken, J. C., Ge, J., Mahadevan, S., DeWitt, C., Friedman, J., Finnerty, P., Ren, D., & Zugger, M. 2004b, in Proc. SPIE, 5492, Ground-Based Instrumentation for Astronomy, ed. A. F. M. Moorwood & M. Iye, 445

van Eyken, J. C., Ge, J., Mahadevan, S., DeWitt, C., & Ren, D. 2003, in Proc. SPIE, 5170, Techniques and Instrumentation for Detection of Exoplanets, ed. D. R. Coulter, 250

van Eyken, J. C., Ge, J., Wan, X., Zhao, B., Hariharan, A., Mahadevan, S., Dewitt, C., Guo, P., Cohen, R., Fleming, S. W., McDavitt, D., Crepp, J., Warner, C., Kane, S., Leger, F., & Pan, K. 2007, in Rev. Mex. AA Ser. Conf., 29, First Light Science with the GTC, 151

Verschueren, W., & David, M. 1999, A&AS, 136, 591

Vogt, S. S., Marcy, G. W., Butler, R. P., & Apps, K. 2000, ApJ, 536, 902

Walker, G. A. H., Walker, A. R., Irwin, A. W., Larson, A. M., Yang, S. L. S., & Richardson, D. C. 1995, Icarus, 116, 359

Wan, X., Ge, J., Guo, P., Zhao, B., Hariharan, L., & McDavitt, D. 2006, in Proc. SPIE, 6269, Ground-Based and Airborne Instrumentation for Astronomy, ed. I. S. McLean & I. Masanori, 62692T

Wolszczan, A., & Frail, D. A. 1992, Nature, 355, 145

Zhao, B., & Ge, J. 2006, in Proc. SPIE, 6269, Ground-Based and Airborne Instrumentation for Astronomy, ed. I. S. McLean & I. Masanori, 62692U

BIOGRAPHICAL SKETCH

After narrowly avoiding being sold in a two-for-one bundle along with his mother to a Masai warrior in Africa (for the respectable sum of seven goats), Julian van Eyken was eventually born in 1976 in rural Somerset, England, where he grew up under the care of two of the best parents in the world. As a teenager he attended Wells Cathedral School, where he specialised in physics, mathematics and music. He graduated from Cambridge University with a degree in natural sciences in 1998, and went on to pursue his PhD in astronomy on a fellowship at Penn State University, where he was later awarded two SPIE scholarships, and a Michelson Fellowship by NASA JPL in 2004, for his work on the ET instruments. Later that year he moved to the University of Florida to complete his PhD working on the same project. He can do the Vulcan salute with *both* hands, knows his African bush creature noises as well as his farm animal sounds, and is always a little suspicious about what turn life might take next.