

Dissertation
submitted to the
Combined Faculties of the Natural Sciences and Mathematics
of the Ruperto-Carola-University of Heidelberg, Germany
for the degree of
Doctor of Natural Sciences

Put forward by
Paul HEEREN
born in: Kulmbach, Germany
Oral examination: 20.07.2021

Testing Planet Candidates around Giant Stars: Computation and Analysis of High Precision Radial Velocities

Referees:

Prof. Dr. Andreas QUIRRENBACH
Prof. Dr. Henrik BEUTHER

“The fact that we live at the bottom of a deep gravity well, on the surface of a gas covered planet going around a nuclear fireball 90 million miles away and think this to be normal is obviously some indication of how skewed our perspective tends to be.”

Douglas Adams

RUPRECHT-KARLS-UNIVERSITÄT HEIDELBERG

Abstract

Testing Planet Candidates around Giant Stars: Computation and Analysis of High Precision Radial Velocities

by Paul HEEREN

The radial velocity (RV), or Doppler, technique is one of the most successful methods in the search for exoplanets; with more than two decades of RV measurements acquired for some stars, and thanks to a precision around 1 m s^{-1} and better reached by modern spectrographs, it allows to explore an ever greater variety of planetary systems. In this PhD dissertation, I present my contributions to the RV survey of G- and K-giant stars, which is conducted by the Exoplanet Group at the Landessternwarte (LSW) Heidelberg. The aim is to track planet candidates in the sample, and thus strengthen our understanding of planet occurrence rates around these types of stars and discern between planetary signatures and false positives caused by intrinsic stellar variations.

My work can be split into two parts: First, I was involved in the Waltz telescope project, which will act as a successor to the CAT telescope at Lick observatory and allow to continue the RV survey of giant stars with an LSW-owned telescope. I describe my work on opto-mechanical components which enabled to reach first light on-sky, and I present results from early observations. My main task in the project was the build-up of the Waltz DRS (data reduction software), which will be used to reduce acquired spectra and extract RVs. I detail the structure of the software and implementation of mathematical methods, and discuss first test results on early Waltz and archived Lick spectra.

The second part of my work concerns the analysis of the highly eccentric stellar binary ϵ Cygni, which is part of the K-giant sample. RV data obtained at Lick, with the SONG telescope on Tenerife, and from the literature show short-period variations in addition to the long signal caused by the stellar companion, which might hint at a planetary companion on a S-type orbit around the primary. I present Keplerian and dynamical models and constrain the orbit of the stellar companion; however, in combination with a stability analysis of the system, the models deem the planet hypothesis to be highly unlikely. I examine possible alternative explanations for the short-period RV variations and find tidally induced stellar oscillations as a plausible cause.

Zusammenfassung

Prüfung von Planetenkandidaten um Riesensterne: Berechnung und Analyse von hochpräzisen Radialgeschwindigkeiten

Die Radialgeschwindigkeits-Technik, oder auch Doppler-Technik, ist eine der erfolgreichsten Methoden zum Auffinden von Exoplaneten. Mit mehr als zwei Jahrzehnten gesammelter Messungen für manche Sterne, und mit Hilfe moderner Spektrographen, die eine Präzision von 1 m s^{-1} und besser erreichen, erlaubt diese Methode eine immer größere Vielfalt von Planetensystemen zu erforschen. In dieser Dissertation stelle ich meine Beiträge zur Radialgeschwindigkeitsstudie von G- und K-Riesensternen vor, die von der Exoplanetengruppe an der Landessternwarte (LSW) Heidelberg durchgeführt wird. Sie hat zum Ziel, Planetenkandidaten in der Stichprobe aufzuspüren, und damit unser Verständnis der Planetenvorkommen um diese Sterntypen zu verbessern und zwischen von Planeten verursachten Signalen und falsch-positiven Signaturen aufgrund von stellaren Variationen zu unterscheiden.

Meine Arbeit lässt sich in zwei Bereiche aufteilen: Einerseits war ich am Waltz-Teleskop-Projekt beteiligt, das als Nachfolger des CAT-Teleskops am Lick-Observatorium dienen und damit eine Weiterführung der Radialgeschwindigkeitsmessungen von Riesensternen an einem Teleskop der LSW erlauben wird. Ich beschreibe meine Arbeit an optisch-mechanischen Komponenten, mit denen erste Beobachtungen am Himmel erreicht werden konnten, und ich präsentiere Resultate dieser Beobachtungen. Meine Hauptaufgabe im Projekt war es, die Datenreduktionssoftware Waltz DRS aufzubauen, die zur Reduktion der aufgenommenen Spektren und Extraktion der Radialgeschwindigkeiten eingesetzt wird. Ich beschreibe den Aufbau der Software und die Einbindung der mathematischen Methoden, und diskutiere erste Testergebnisse an frühen Waltz- und archivierten Lick-Spektren.

Der zweite Teil meiner Arbeit betrifft die Analyse des stark exzentrischen Doppelsystems ϵ Cygni, das Teil der K-Riesen-Stichprobe ist. Radialgeschwindigkeitsmessungen vom Lick-Observatorium, vom SONG-Teleskop auf Teneriffa, und aus der Literatur zeigen neben dem langen Signal, das durch den stellaren Begleiter verursacht wird, auch kurzperiodische Variationen, die auf einen planetaren Begleiter auf einem S-Orbit um den Primärstern hindeuten könnten. Ich präsentiere Kepler- und dynamische Modelle und grenze den Orbit des stellaren Begleiters ein; in Kombination mit einer Stabilitätsanalyse des Systems lassen die Modelle die Planetenhypothese jedoch als sehr unwahrscheinlich erscheinen. Ich untersuche mögliche alternative Erklärungen für die kurzperiodischen Radialgeschwindigkeitsvariationen und halte zeiteninduzierte stellare Oszillationen für eine plausible Ursache.

Acknowledgements

I would not have been able to perform the work presented in this dissertation, let alone write it, without the help of a large network of people around me who supported me in many different ways, some more direct, others probably without even knowing about their helping me. The latter will hopefully know after having read my acknowledgements.

Science is teamwork, and I have discussed my research with more people than I can possibly recall. Of greatest importance and influence on my work however was the feedback of my supervisors, Andreas Quirrenbach and Sabine Reffert: Their detailed suggestions helped to open up many alleged dead-ends, and I am grateful for the chances they have given me, and for the responsibilities they have trusted me with. Furthermore, I want to thank Martin Kürster, third advisor on my IMPRS thesis committee, who always has an open ear for ones worries and knows how to find an encouraging response. Also I am thankful to Henrik Beuther, who agreed to act as second referee of this thesis, and to Joachim Wambsganss and Fred Hamprecht, for joining my defence examination panel. And to the proofreaders of this work, Sabine Reffert and Steffi Yen: Thank you so much for your invested time, and the helpful and very constructive comments!

During my PhD, I dove head-on into the field of planetary dynamics, and my gratitude goes out to Trifon Trifonov and Man Hoi Lee who would always readily provide me with learning material and answers to my questions. For my work on the Waltz telescope project, I got a lot of fruitful advice (e.g. about how not to break things) from Walter Seifert, Otmar Stahl, Julian Stürmer and Rob Harris, who I would like to thank here. Still, the Waltz instruments would not be standing and functional without the workshops of the LSW, and my thanks go out to Jochen Tietz, Edwin Lutz, and particularly Lutz Geuer for his competent work (talking about "Lutz-precision"). My work on the Waltz DRS profited greatly from the help and advice of several experts in the field: René Tronsgaard Rasmussen, who shared his pyodine code with me; Rafael Brahm and Andres Jordan, who guided me when I started working with the CERES package; and notably Frank Grundahl, who not only shared his many experiences (and some frustrations) with RV analysis codes with me, but also supplied me with SONG data to test the Waltz DRS on.

A productive workspace however not only requires work-related help, but also a socially comfortable environment — and for creating this, I want to thank all my colleagues, fellow PhD students (viva la generacion 13), bachelor and master students, and the institute staff members. Particularly, I am grateful for having met all the wonderful characters of the LSW gang; heated discussions, pranks gotten out-of-hand, collective lunching (and dining, some drinking), puzzles, movies, birthday gifts and songs...

Finally, with a global pandemic raging during the final year of my PhD, I am thankful to my flatmates who created a home worth of spending much more than just one lockdown in. I also got by with a little help of my friends, with whom I can take everything with the earnest sarcasm it deserves. I want to thank Anujeema, for feeling so close despite being so far. And I am grateful for my family, who always reassured me of my ability to write this dissertation (even whilst not being fully seated). Most of all to my parents: You taught me to marvel at life, and to follow my heart rather than any outer expectations. Thank you!

Contents

Abstract	ix
Acknowledgements	xi
Contents	xiii
1 Introduction	1
1.1 The Radial Velocity (RV) method	4
1.1.1 Orbital parameters from radial velocities	5
1.1.2 Design and performance of high-resolution spectrographs	7
1.1.3 Reduction of spectra	10
1.1.4 The iodine-cell method	14
1.2 Exoplanets around evolved stars	17
1.2.1 The Lick RV survey of evolved stars	17
1.2.2 Occurrence rate of planets as function of stellar mass and metallicity	18
1.2.3 Overall occurrence rate and period distribution of planets	19
1.3 RV variations beyond the Keplerian model	21
1.3.1 Solar-like oscillations	22
1.3.2 Gravity-mode and mixed-mode oscillations	22
1.3.3 Stellar spots	23
1.3.4 Dynamical analysis of RV measurements	23
2 The Waltz Telescope Project: An overview and recent updates	27
2.1 Introduction	27
2.2 Instrument description	28
2.2.1 Telescope and dome	28
2.2.2 The frontend	29
2.2.3 The calibration unit	30
2.2.4 The spectrograph	31
2.2.5 Control software & electronics	33
2.3 Performance of the system	34
2.3.1 Test of the tracking stability	34
2.3.2 Spectrograph detector characteristics	35
2.3.3 Analysis from calibration spectra	36
2.3.4 Stellar spectra	38
2.4 Modifications of the system	39
2.4.1 Spectrograph detector mount	39
2.4.2 Calibration fiber	40
2.4.3 Design modification and assembly of the FFU	41

2.4.4	FFU alignment and efficiency measurements	43
2.5	Summary	45
3	A data reduction software (DRS) for the Waltz Telescope	47
3.1	Introduction	47
3.2	General structure of the DRS	48
3.3	Reduction of the spectra	49
3.3.1	Overview of the modified CERES code	50
3.3.2	Workflow of the reduction process	52
Pre-processing using calibration frames	52	
Tracing the Echelle orders & scattered light subtraction	53	
Optimally extracting the orders	54	
Wavelength calibration of extracted spectra	55	
Creation of science output spectra	59	
3.3.3	Barycentric correction	62
3.3.4	Test results of the spectrum reduction	64
Simple versus optimal extraction	64	
Global wavelength solution	65	
Velocity drifts of ThAr spectra	67	
3.4	Radial velocity analysis	68
3.4.1	Overview of the upgraded pyodine code	69
3.4.2	Description and implementation of the model	70
LSF models	71	
Wavelength & continuum models	73	
I2 FTS and stellar template spectrum	74	
The combined model and parameter sets	74	
Chunking algorithms	75	
3.4.3	The fitting procedure of the chunk models	77
3.4.4	Constraining and altering starting parameters	78
3.4.5	Pixel weights	79
Telluric mask	80	
Bad pixel mask	82	
3.4.6	Template creation	82
3.4.7	Combination of chunk velocities	84
SONG velocity weighting algorithm	86	
Velocity combination algorithm based on the dop code	88	
3.4.8	Workflow of the observation modelling	88
3.4.9	Workflow of the template creation	90
3.5	Testing of the RV analysis code	92
3.5.1	Results from simulated spectra	93
The Single-Gaussian LSF analysis	95	
The Multi-Gaussian LSF analysis	96	
3.5.2	Early results from Lick observations	98
Modelling procedures and parameters	99	
Stellar template creation for HIP 96459	100	
Modelling of I2 observations of HIP 96459	104	
RV results for Pollux	107	
RV results for HIP 36616	110	
3.5.3	An early test on SONG spectra	113
3.6	Computational performance of the Waltz DRS	116
Reduction of spectra	117	

Extraction of RVs	118
3.7 Conclusion and outlook	118
4 Is there a planet in the highly eccentric stellar binary ϵ Cygni?	123
4.1 Introduction	124
4.2 Stellar properties	124
4.3 Observations	126
4.4 Analysis of the RV data	128
4.4.1 Determining the orbit of the close stellar companion	128
4.4.2 GLS periodogram of the RV residuals	130
4.4.3 Fitting a double-Keplerian model to the data	132
4.4.4 Investigating the temporal evolution of the Keplerian signal	134
4.4.5 Fitting a dynamical model to the RV data	136
4.4.6 The co-orbital scenario	137
4.5 Dynamical stability analysis	138
4.5.1 Theoretical considerations about the orbital configuration	138
4.5.2 Stability analysis of coplanar orbits	139
4.5.3 Stability analysis of mutually inclined orbits	140
4.6 Possible alternative explanations for the short-period RV variations	142
4.6.1 Hierarchical triple	142
4.6.2 Stellar spots	144
4.6.3 Oscillatory convective modes	146
4.6.4 Potential stellar oscillations through the heartbeat phenomenon	147
4.7 Summary & conclusions	150
5 Conclusions and future prospects	153
A Waltz DRS parameters	155
B Measured Radial Velocities of ϵ Cygni	159
Own publications	163
Bibliography	165

List of Figures

1.1	Cumulative number of exoplanet detections over time.	2
1.2	Distribution of exoplanets discovered to date.	3
1.3	Radial velocity measurements and best Keplerian model of 51 Pegasi.	5
1.4	Schematic cross-section of a blazed diffraction grating.	8
1.5	Schematic optical design of an Echelle spectrograph.	9
1.6	Echelle spectrum of a Quartz-Tungsten halogen lamp.	10
1.7	Schematic depiction of the optimal extraction algorithm.	12
1.8	Basic steps in the I2-cell method.	16
1.9	Planet occurrence rate as a function of metallicity and stellar mass in the Lick survey.	18
1.10	Distribution of MS and evolved stars that host giant planets detected by the RV method.	20
2.1	Model of the Waltz frontend.	29
2.2	Model of the Waltz spectrograph.	32
2.3	Measured tracking error of the Waltz telescope.	34
2.4	Spectrograph detector characteristics.	36
2.5	Results from calibration spectra.	37
2.6	Results from observations of the star Arcturus.	38
2.7	Modified detector mount of the Waltz spectrograph.	39
2.8	Pinhole and calibration light focus.	40
2.9	Design of the modified FFU.	42
2.10	Laboratory results from the FFU.	43
2.11	Estimated efficiencies of the FFU for different seeing conditions.	45
3.1	Structure of the Waltz DRS.	49
3.2	Organisation of the modified CERES package.	51
3.3	Flowchart of the reduction of science spectra.	52
3.4	Order tracing results.	53
3.5	Analysis plots of the initial wavelength calibration.	57
3.6	Reduction result of an observation of HIP 37826 (1).	61
3.7	Test results of the BVs computed by barycorrpy on observations from the Lick survey.	64
3.8	Reduction result of an observation of HIP 37826 (2).	65
3.9	Result from the global wavelength calibration computed on a ThAr frame.	66
3.10	Measured velocity drifts of wavelength solutions of ThAr spectra.	67
3.11	Structure of the modified pyodine package within the Waltz DRS.	69
3.12	Exemplary LSFs.	73

3.13	Schematic representation of the three different chunking algorithms.	76
3.14	Telluric features in the I2 wavelength region.	81
3.15	Flowchart of the extraction of radial velocities.	89
3.16	Flowchart of the creation of stellar templates.	91
3.17	Simulated spectra.	94
3.18	One chunk of the simulated modelling.	95
3.19	Results from simulated observations, with the <code>SingleGaussian</code> LSF model.	96
3.20	Results from simulated observations, with the <code>MultiGaussian_Lick</code> and <code>MultiGaussian</code> LSF models.	97
3.21	Modelling results of O-star spectra for the template creation of HIP 96459.	101
3.22	Median evaluated LSF of all chunks in the template creation of HIP 96459.	102
3.23	Four exemplary deconvolved stellar template chunks of HIP 96459.	103
3.24	Differences of wavelength parameters between the Lick and Waltz DRS templates.	104
3.25	RV timeseries results for HIP 96459.	105
3.26	Analysis results of RVs of HIP 96459.	106
3.27	RVs and Keplerian model results for the star Pollux.	107
3.28	GLS periodograms of the RV timeseries of Pollux.	109
3.29	RVs and Keplerian model results for the star HIP 36616.	111
3.30	RVs (from SONG spectra) and Keplerian model results for the star ϵ Cyg.	114
3.31	GLS periodograms of the RV timeseries of ϵ Cyg.	116
4.1	Uncorrected SONG RV measurements from the last week of our asteroseismic campaign on ϵ Cyg.	127
4.2	RV measurements of ϵ Cyg, and residuals to the Single-Keplerian models.	129
4.3	GLS periodograms of the residuals of the RV measurements after removal of the binary signal.	132
4.4	RV measurements of ϵ Cyg and best-fit model, phase-folded by the best-fit period of the putative planet.	133
4.5	Evolution of the orbital period and semi-amplitude of the presumed planet.	135
4.6	GLS periodograms of the residuals of the best dynamical models.	137
4.7	Lomb-Scargle periodogram of the RV residuals and the corresponding demodulated RVs.	138
4.8	Survival times for test particles placed into the ϵ Cyg stellar binary.	141
4.9	Exemplary RVs for a simulated hierarchical triple.	143
4.10	HIPPARCOS photometry for ϵ Cyg.	145
4.11	BRITTE photometry for ϵ Cyg.	146
4.12	Primary stellar radii over binary orbital periods for a number of heartbeat systems.	148
4.13	Primary stellar radii and tidal forcing ratio over periastron distances for heartbeat systems.	149

List of Tables

1.1	Examples of radial velocities.	7
2.1	Properties of the Waltz Spectrograph	31
2.2	Cumulative efficiencies of the FFU	44
3.1	FITS output formats of reduced science spectra in the Waltz main reduction routine	60
3.2	Overview of free parameters of the sub-models, and overall number of free model parameters n	75
3.3	Properties of simulated observations	93
3.4	Template creation for HIP 96459, median LSF parameter results of all chunks	102
3.5	Keplerian parameters for the planetary companion to Pollux, from the Waltz DRS and Lick RVs	108
3.6	Keplerian parameters for the HIP 36616 system, from the Waltz DRS and Lick RVs	112
3.7	Keplerian parameters for the HIP 102488 system, from the Waltz DRS and iSONG RVs extracted from SONG spectra	115
4.1	Stellar properties of ϵ Cyg A	125
4.2	Keplerian parameters of the ϵ Cyg binary system, from the single- and double-Keplerian models	131
A.1	Parameters used for reduction of Waltz spectra, as contained in the parameter input file	155
A.1	continued.	156
A.1	continued.	157
A.2	Parameters used for the I2 analysis of Lick spectra with the Waltz DRS	157
A.3	Parameters used for the I2 analysis of SONG spectra with the Waltz DRS	157
A.3	continued.	158
B.1	RV measurements, taken with the Hamilton Spectrograph at Lick Observatory	159
B.1	continued.	160
B.2	RV measurements of the SONG Tenerife node (for the high-cadence asteroseismic RV measurements only the nightly median is included)	160
B.2	continued.	161
B.2	continued.	162

List of Abbreviations

ADU	Analog-to-Digital Units
BVC	Barycentric Velocity Correction
CARMENES	Calar Alto high-Resolution search for M dwarfs with Exoearths with Near-infrared and optical Echelle Spectrographs
CAT	Coudé Auxiliary Telescope
CCD	Charge-Coupled Device
CCF	Cross-Correlation Function
CERES	Collection of Elemental Routines for Echelle Spectra
DEC	Declination
DRS	Data Reduction Software
ELT	Extremely Large Telescope
ESPRESSO	Echelle Spectrograph for Rocky Exoplanets and Stable Spectroscopic Observations
EXPRESS	EXoplanets around Evolved Stars
FFU	Fiber-Feeding Unit
FWHM	Full Width at Half-Maximum
GLS	Generalized Lomb-Scargle (periodogram)
HARPS	High Accuracy Radial Velocity Planet Searcher
HB	Horizontal Branch
HIPPARCOS	High Precision Parallax Collecting Satellite
HIRES	High Resolution Echelle Spectrometer
JWST	James Webb Space Telescope
LED	Light-Emitting Diode
LM	Levenberg-Marquardt (minimization)
LSF	Line Spread Function
LSP	Long Secondary Period (star)
LSW	Landessternwarte Heidelberg
MCMC	Markov-Chain Monte Carlo
MS	Main Sequence
PDF	Probability Density Function
PMT	Photo-Multiplier Tube
PPPS	Pan-Pacific Planet Search
PSF	Point Spread Function
RA	Right Ascension
RGB	Red Giant Branch
RV	Radial Velocity
SERVAL	Spectrum Radial Velocity Analyser
SONG	Stellar Observations Network Group
ThAr	Thorium-Argon

*To Regina & Bernhard,
for their caring.*

INTRODUCTION

At some point in my childhood, I discovered the part of my parents' bookshelf that was home to a great number of science fiction novels — and I became fascinated with them. Classics such as "The Mote in God's Eye" by Larry Niven and Jerry Pournelle, "The Rolling Stones" by Robert A. Heinlein, and others bore witness of the great expectations that the public had about the universe, human space explorations and of course possible alien life in the middle of the last century. They were sparked by the rapid evolution in science and technology, and it only seemed to be a matter of time until many of the ideas and worlds described in these books would become reality.

Looking back today, at the wild imaginations about the number and nature of planets outside the Solar System, it seems hard to believe that even only 40 years ago we had no actual knowledge about the occurrence of these so-called exoplanets. Of course this is not due to ignorance on the side of scientists; speculations about the possibility of other worlds in the universe go back at least to the old Greeks, and with astronomical instruments becoming ever larger and more precise people in the modern era started thinking about how to finally test the exoplanet hypothesis. The general problem is: Stars are generally much brighter than their planetary companions, and detections of exoplanets through direct imaging therefore are only possible in very few cases even with the most modern instruments. In this context, the very short but accurate reasoning put forward by Struve (1952), about the possibilities of spectrographic and photometric detections of massive planets orbiting close-in to their host stars, is especially notable: It focused precisely on the two most successful exoplanet detection methods today, the so-called Radial Velocity (RV) or Doppler method and the Transit method, and their respective strengths.

Still, claimed exoplanet detections of that time, such as around Barnard's star on the basis of astrometric measurements (van de Kamp, 1969), were generally proven wrong as time passed on. The first confirmed exoplanet detection was announced as late as 1992, when the radio astronomers Alexander Wolszczan and Dale Frail recorded periodic variations in the pulse arrival times of the millisecond pulsar PSR 1257+12 (Wolszczan and Frail, 1992). These regular shifts with an amplitude of just ± 15 ps are caused by the gravitational interaction of two planets orbiting the dead star, thus pulling it ever so slightly around the system's barycenter. It is no surprise that pulsar timing delivered the first verified exoplanet detection – this technique reached the required measurement precision very early on as compared to other methods. In contrast, promising signatures in radial velocity measurements of that time, such as for the stars γ Cephei (Campbell, Walker, and Yang, 1988) and Pollux (Hatzes and Cochran, 1993), were too spurious to be clearly indicative of a sub-stellar companion: Given the measurement uncertainties, other possible explanations for the observed radial velocity signatures (such as of stellar origin) could




FIGURE 1.1: Cumulative number of exoplanet detections over time, indicated by detection method. *Credit:* NASA Exoplanet Archive.

not be ruled out. The planet hypothesis in these cases was only confirmed much later (Hatzes et al., 2003; Hatzes et al., 2006; Reffert et al., 2006).

The first confirmed discovery of an exoplanet around a solar-type star is attributed to Mayor and Queloz (1995), who noticed variations in their radial velocity measurements of the main-sequence star 51 Peg, caused by a Jupiter-mass companion on a very tight orbit with a period of just 4.2 d (see Figure 1.3).¹ Such an orbital architecture was quite unexpected at the time, as theories of planet formation suggested that massive planets should form at much larger orbital separations from their host stars. Today we know that these so-called "Hot Jupiters" migrated from their birthplaces in the outer parts of protoplanetary disks into closer orbits due to interaction with the disk material.

Following the discovery by Mayor and Queloz (1995), which was also the first confirmed one using the Doppler method, the still very young research field of exoplanets rapidly developed: During the next two decades, the number of new exoplanet detections increased nearly every year (compare Figure 1.1). Thanks to high-resolution spectrographs such as HARPS in La Silla/Chile (Mayor et al., 2003), HIRES at Keck Observatory/Hawaii (Vogt et al., 1994), or the Hamilton spectrograph at Lick Observatory/California (Vogt, 1987), the Doppler method remained the most successful detection technique for exoplanets until 2013 (see Figure 1.1).

This changed however with the operation of Kepler, the first space telescope mainly

¹For that discovery, and "for contributions to our understanding of the evolution of the universe and Earth's place in the cosmos" (Nobel Prize Outreach AB, 2019), the authors were awarded with one half of the Nobel Prize in Physics in 2019.




FIGURE 1.2: Distribution of exoplanets discovered to date, indicated by detection method. *Credit:* NASA Exoplanet Archive.

dedicated to search for exoplanets, starting in 2009: Using the Transit method, Kepler was able to glance at many stars at once and scan them simultaneously for transit events of possible planetary companions. The photometry measurements alone led to more than 3000 confirmed planet detections to date; in combination with discoveries from the Doppler method and other techniques, we today stand at a total of 4383 known exoplanets in 3249 systems². Thanks to these discoveries, we have learned about the wide variance of planetary systems, and generally about the occurrence of planets within our home galaxy: Notably, concerning the overall occurrence rate of Earth-analogues around Sun-like stars, Petigura, Howard, and Marcy (2013) derived an estimate of roughly 11%, and more recently Hsu et al. (2019) found an upper limit of 27%. Furthermore, many so-called Super-Earths and Mini-Neptunes have been discovered by Kepler, which have radii between $\sim 1 - 1.5 R_{\oplus}$ and $2 - 3 R_{\oplus}$, respectively, and are not prevalent in our own Solar System. There seems to exist a distinct difference in the formation scenario for these two types of planets, as many studies have detected a significant drop in the occurrence rate for the intermediate radius regime (e.g. Fulton et al., 2017). Finally, detections from microlensing surveys suggest that on average each star in the Milky Way hosts at least one planet (Cassan et al., 2012).

After the gold-rush-era of the exoplanet field, with many discoveries made on the basis of few observations, it has now entered a phase in which detection limits raised by instrumental and astrophysical effects, such as stellar oscillations or magnetic activity, are pushed. The next generation of instruments such as the James Webb Space Telescope (JWST)³ or the HIRES spectrograph at the forthcoming ELT (Marconi et

²<https://exoplanetarchive.ipac.caltech.edu/>, as of May 6th, 2021

³<https://www.jwst.nasa.gov/>

al., 2021) promise another large step in instrumental precision, however these facilities only offer a limited amount of observing time. As some stars require either very frequent observations or a survey over a long time-span to better understand ambiguous signatures and confirm or reject planet discoveries, the community will also have a need of less prominent instruments that offer better accessibility.

Throughout my work I have participated in the Waltz project, building and testing a low-cost, institute-owned telescope and spectrograph dedicated for radial velocity measurements of evolved stars, which will thus allow to continue a survey started in 1999 at Lick observatory. Chapter 2 will present the current status of the project and some of my contributions to it. A large amount of my work was invested in the development of the data reduction software (DRS) for the project, and in Chapter 3 I will describe its structure, algorithms and show first test results. Chapter 4 finally treats the radial velocity analysis of the K-giant and stellar binary system ϵ Cygni, which serves as an example of the importance of long-baseline observations and advanced analysis tools for finding ambiguities. This chapter serves as an introduction to the reader, and Chapter 5 summarises conclusions from my work and gives an outlook on future prospects.

1.1 The Radial Velocity (RV) method

Stars do not stand still in space, but move with respect to each other⁴. The component of the velocity vector of a star that is directed towards an observer (e.g. on Earth or a space telescope) is called the line-of-sight or radial velocity (RV). Due to the Doppler effect this radial velocity introduces a shift of the observed wavelengths of the star's emitted light as seen by the observer, $\Delta\lambda$, described by

$$z = \frac{\lambda_{\text{obs}} - \lambda}{\lambda} \equiv \frac{\Delta\lambda}{\lambda} = \frac{v}{c}, \quad (1.1)$$

where λ_{obs} and λ are the observed and laboratory wavelengths of the light, respectively, c is the speed of light, and v the relative velocity between the star and the observer along the line of sight⁵. The relative wavelength shift is then called the Doppler shift z . When using the RV method to search for exoplanets, we are only interested in the change of Doppler shift over time and do not care about absolute values; then, instead of laboratory wavelengths λ , we can insert any other reference wavelength frame λ_{ref} in the above expression, for instance a template spectrum of the same star, and thus calculate the relative Doppler shift and radial velocity with respect to that template.

It is important to consider though that the observing instrument itself is moving (e.g. with a one-year period around the sun, and a 24-hour period around the Earth's axis), which causes daily and yearly variations of measured radial velocities of several 10 km s^{-1} that definitely need to be accounted for. Additional smaller-scale phenomena, such as gravitational time dilation due to objects of the Solar System, might also be corrected for depending on the desired radial velocity precision (see Wright and Eastman, 2014, for a well-founded analysis). We will sum up all of

⁴A nice visualization of the relative movement of stars within the Milky Way was created based on measurements taken by the satellite mission *Gaia*: https://www.esa.int/ESA_Multimedia/Videos/2020/12/Gaia_s_stellar_motion_for_the_next_1.6_million_years.

⁵Note that we are using the equation of the non-relativistic Doppler effect here, because in most circumstances the condition $v \ll c$ is fulfilled.




FIGURE 1.3: Radial velocity measurements (black dots) and the best Keplerian model (solid line) of the solar-type star 51 Pegasi. The data has been phase-folded to the derived orbital period of the model. *Credit: Mayor and Queloz (1995).*

these effects hereafter under the so-called barycentric velocity v_B . If an accuracy of $\sim 3 \text{ m s}^{-1}$ suffices, the true radial velocity of the target star v_* can then be determined from the measured velocity shift of the spectrum v_{meas} through the simple relation $v_* = v_{\text{meas}} + v_B$ (more details are given in Section 3.3.3).

1.1.1 Orbital parameters from radial velocities

Radial velocities of stars as measured by a distant observer are influenced by various astrophysical phenomena and can be split up into different components: The star's orbital movement around the center of the Milky Way leads to a systemic radial velocity, which usually only changes very slowly; furthermore, intrinsic stellar sources, such as stellar spots and oscillations, can induce periodic radial velocity variations over a wide range of periods (see Section 1.3). For now, let us neglect these effects and assume a star that is orbited by one or more companions, and is therefore forced to move about the barycenter of the system. If the companions' orbits are not purely constrained to the plane of the sky as seen by the observer, so their orbital inclinations i are non-zero, the star's radial velocity will vary periodically over time: When the star moves towards the observer its light appears blue-shifted, and it will be red-shifted when the star moves away. By measuring this periodic variation and analyzing the signal, one can deduce certain parameters of the system. In the following, I will use some simplifications to arrive at meaningful relations; a more thorough derivation can be found in Lovis and Fischer (2011).

Let's assume the simple case of a star with only one companion. Kepler's Third Law states

$$a^3 = \frac{P^2}{4\pi^2} G(m_* + m_p) \quad , \quad (1.2)$$

where G is the gravitational constant, m_* and m_p are the masses of the star and companion, respectively, and P is the orbital period. The semi-major axis of the relative orbit of the two components is the sum of their individual semi-major axes about the barycenter: $a \equiv a_* + a_p$. Due to equilibrium of forces, their barycentric orbits obey

$$a_* m_* = a_p m_p \quad \Rightarrow \quad a = a_* + a_p = \frac{(m_* + m_p)}{m_p} a_* \quad . \quad (1.3)$$

If we first constrain our thoughts to strictly circular orbits, it is obvious that the orbital velocity of the star is simply the length of its orbit divided by the orbital period: $v_* = 2\pi a_* P^{-1}$. The maximum (absolute) radial velocity of the star as seen by the observer is called the semi-amplitude and depends on the inclination i of the orbit: $K_* = v_* \sin i$. In the case of orbits with non-zero eccentricity e , the relation needs to be modified and becomes

$$K_* = v_* \sin i = \frac{2\pi a_* \sin i}{P\sqrt{1-e^2}} \quad . \quad (1.4)$$

Using Equations 1.2, 1.3 and 1.4, we can eliminate the semi-major axis and obtain the relation of the so-called mass function:

$$f(m_p) = \frac{(m_p \sin i)^3}{(m_* + m_p)^2} = \frac{P}{2\pi G} K_*^3 \sqrt{(1-e^2)^3} \quad . \quad (1.5)$$

The orbital period P , semi-amplitude K_* and orbital eccentricity e can be deduced from the radial velocity measurements. However, even if we have knowledge of the stellar mass m_* (e.g. through spectroscopic or asteroseismic analysis), the mass function only allows to determine the *minimum mass* of the planet $m_{p,\min} \equiv m_p \sin i = m_p \sin 90^\circ$, as the orbital inclination i is not an observable quantity in the RV method. Thus, for a given semi-amplitude K_* , the true planetary mass can be considerably larger than the minimum mass if the inclination of the orbit is small. Still, for a random distribution of star-planet systems, it is statistically more likely to observe systems roughly edge-on ($i \approx 90^\circ$) rather than face-on ($i \approx 0^\circ$); for example, there is a probability of 87% that the inclination will lie between 30 and 90°, corresponding to an actual planet mass of 1 to 2 times the minimum mass derived from RVs (compare Quirrenbach, 2006; Fischer et al., 2014). We can conclude therefore that distributions of measured $m_p \sin i$ represent the true planetary mass distributions quite well.

By re-arranging above relation to solve for the semi-amplitude of the radial velocity signature K_* , and expressing the parameters in suitable units we can write (compare also Lovis and Fischer, 2011):

$$K_* \approx \frac{28.4376 \text{ m s}^{-1}}{\sqrt{1-e^2}} \frac{m_p \sin i}{M_J} \left(\frac{m_* + m_p}{M_\odot} \right)^{-\frac{2}{3}} \left(\frac{P}{1 \text{ yr}} \right)^{-\frac{1}{3}} \quad , \quad (1.6)$$

or, by substituting Kepler's Third Law (Equ. 1.2) for the period:

$$K_* \approx \frac{28.4376 \text{ m s}^{-1}}{\sqrt{1-e^2}} \frac{m_p \sin i}{M_J} \left(\frac{m_* + m_p}{M_\odot} \right)^{-\frac{1}{2}} \left(\frac{a}{1 \text{ AU}} \right)^{-\frac{1}{2}} \quad . \quad (1.7)$$

TABLE 1.1: Examples of radial velocities (assuming $e = 0$, $i = 90^\circ$, and just the one planet in the system)

System	m_* [M_\odot]	a [AU]	m_p [M_\oplus]	K_* [m s^{-1}]
Jupiter & Sun	1.0	5.20	317.8	12.47
Hot Jupiter & Sun	1.0	0.10	317.8	89.93
Earth & Sun	1.0	1.00	1.0	0.09
Super-Earth & Sun	1.0	1.00	5.0	0.45
Neptune & Sun	1.0	30.11	17.2	0.28
Earth in HZ around M-dwarf ^a	0.1	0.04	1.0	1.41
Gas giant around K-giant star ^b	1.9	1.7	2.9	45.89

^a Habitable zone (HZ) according to Kopparapu et al. (2013)

^b These parameters are comparable to the Pollux system (Reffert et al., 2006)

This expression helps us to receive a proper understanding of the magnitudes of the radial velocity variations for given systems. Table 1.1 displays some prominent examples: An alien observer measuring the line-of-sight velocity of the Sun would notice a semi-amplitude of 12.47 m s^{-1} due to Jupiter's gravitational pull; the presence of Earth in the Solar System however only leads to a semi-amplitude of 9 cm s^{-1} . Hot Jupiters, which orbit very close-in to their host star, induce much larger radial velocity amplitudes due to their small semi-major axes; and less massive stars such as M-dwarfs generally allow the detection of less massive planets at the same radial velocity precision.

It is obvious that, in order to be sensitive to planetary companions around solar-type stars, one needs to measure radial velocities with a precision of at least a few 10 m s^{-1} . Even then, only gas giants orbiting close-in to their host stars can be detected — expanding the parameter space to massive planets on larger orbits, or rocky planets close to the habitable zone around M-dwarfs, instead requires a precision on the order of 1 m s^{-1} . According to Equation 1.1, this corresponds to a Doppler-shift on the order of 10^{-9} , or a wavelength shift of just a few 10^{-5} \AA (at wavelengths around 6000 \AA , where solar-type stars emit most of their light). Measuring such tiny wavelength shifts reliably is truly challenging from an instrumental point of view, and requires excellent data reduction methods.

1.1.2 Design and performance of high-resolution spectrographs

Spectrographs used in radial velocity surveys need to combine two important features – namely a high resolving power, and a broad wavelength range. In the following I will give a concise overview of these aspects and instruments, concentrating on the so-called Echelle spectrographs that are widely used for radial velocity measurements. The Waltz spectrograph that I have worked with in the Waltz project has also been constructed following these design principles (see Section 2.2.4). A complete theoretical background on optical properties of spectroscopic instruments can be found in Eversberg and Vollmann (2015).

The spectral resolving power R , indicating the ability of an instrument to resolve features of different wavelengths, is an important quantity of a spectrograph and is given by

$$R = \frac{\lambda}{\Delta\lambda} \quad , \quad (1.8)$$

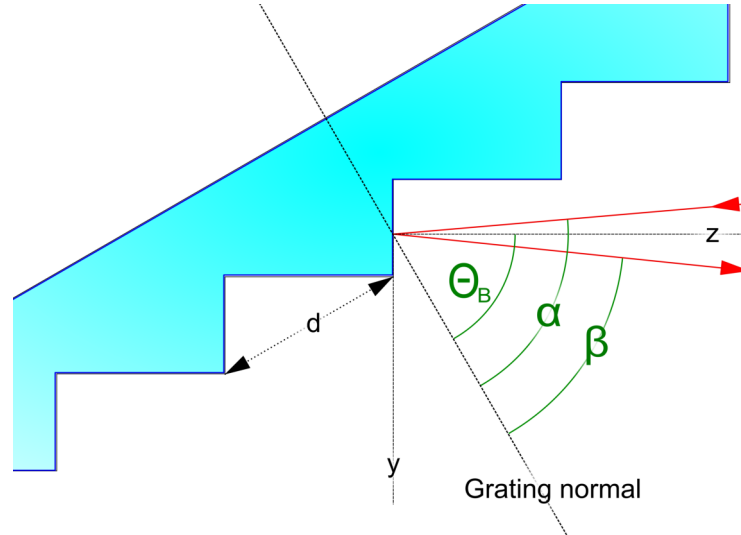


FIGURE 1.4: Schematic cross-section of a blazed diffraction grating, with blaze angle Θ_B , groove width d , and incidence and diffraction angles α and β . *Credit: Based on Figure 5.4 in Eversberg and Vollmann (2015).*

with λ being the observed wavelength and $\Delta\lambda$ the spectral resolution (the wavelength difference between two resolved features). For a spectrograph using a grating as main dispersion element, the grating equation describes the dispersion of light of wavelength λ , hitting the grating under an angle of incidence α , into order numbers n :

$$n \cdot \lambda = d \cdot (\sin \alpha + \sin \beta) \quad , \quad (1.9)$$

where d is the groove width of the grating, and β is the diffraction angle. In practise, many spectrographs are constructed in so-called quasi-Littrow configuration: They use a grating with grooves that are blazed under a blaze angle Θ_B , and the incidence and diffraction angles are perpendicular to the surfaces of the grating grooves (so $\Theta_B = \alpha \approx \beta$, compare Figure 1.4).⁶ We can then modify the grating equation accordingly, and derive the angular dispersion of the instrument from it:

$$\frac{d\beta}{d\lambda} = \frac{2 \cdot \tan \Theta_B}{\lambda} \quad . \quad (1.10)$$

Finally, using Equations 1.9 and 1.10, we can write for the resolving power of the instrument:

$$R = 2 \tan \Theta_B \cdot d\beta^{-1} = 2 \tan \Theta_B \cdot \frac{f_{\text{Cam}}}{s_{\text{CCD}}} \quad . \quad (1.11)$$

Here, f_{Cam} is the focal length of the camera imaging the spectrum onto the CCD, and s_{CCD} is the width of the imaged entry slit (or fiber cross-section) in the CCD focal plane. It is apparent that for a given f_{Cam} (and generally for given spectrograph dimensions), high resolving powers can be achieved by reducing the slit width as much as possible. For telescopes without an adaptive-optics (AO) system however,

⁶This way the efficiency of the grating is maximized.

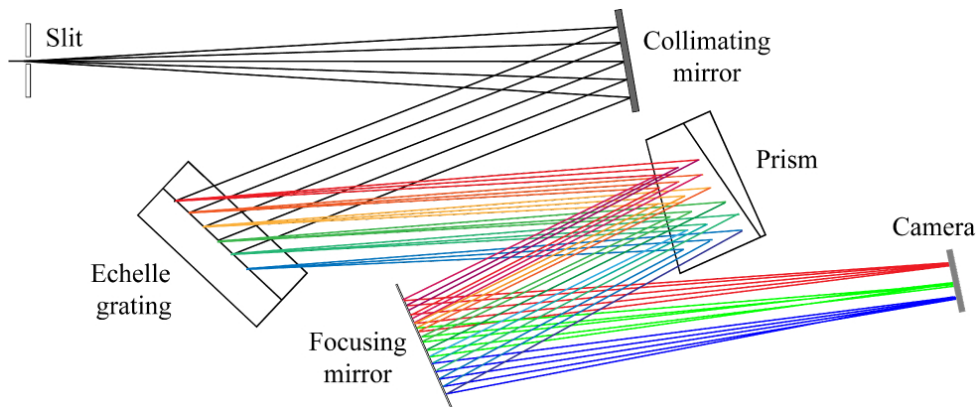


FIGURE 1.5: Schematic optical design of an Echelle spectrograph. Credit: <https://www.osapublishing.org/oe/fulltext.cfm?uri=oe-26-26-34131&id=403158>

the seeing induced by the atmosphere places a lower limit on the slit width — decreasing it further would mean losing light. Many spectrographs therefore make use of an image slicer that splits and re-arranges the telescope point-spread-function (PSF) to allow for a narrower slit (typically by a factor of 2 to 4, depending on the number of slices produced; compare Bowen, 1938; Pierce, 1965). In fiber-coupled spectrographs the image slicer usually sits behind the fiber output.

The second path to high resolving powers, and thus to resolving ever finer wavelength intervals $\Delta\lambda$, is using gratings with large blaze angles Θ_B . Typical spectrographs in RV surveys are equipped with so-called R2 or R4 gratings, corresponding to blaze angles of 63.5° and 75° , respectively, and achieve resolving powers between 50,000 and 150,000. This corresponds to resolutions of roughly $5 \cdot 10^{-2} \text{ \AA}$, which, given the low signal-to-noise ratios (S/N) of typical observations, still does not suffice to measure radial velocity shifts with the required precision and accuracy on a single absorption line.

Instead, RV spectrographs are constructed such that they combine the high resolutions stated above with a broad spectral wavelength range, in order to allow measurements on many individual spectral features and thus reduce the combined uncertainty. In most cases this is achieved by using Echelle gratings as main dispersion elements, which are characterized by large blaze angles and large groove widths (typically $3 \sim 50 \mu\text{m}$). This way, spectrographs working at visual wavelengths achieve their peak efficiencies at high order numbers of $n = 60 \sim 130$, where the orders strongly "overlap" — meaning that a given wavelength in two consecutive orders is separated only by a very small angle $\Delta\beta$ (near-infrared spectrographs usually use slightly lower order numbers). A second dispersion element, the so-called cross-disperser, is then placed such that the overlapping orders are separated perpendicularly to the grating diffraction axis, as depicted in Figure 1.5; prisms, gratings, grisms or more recently Volume Phased Holographic (VPH) gratings are often used as cross-disperser. When imaged on a CCD chip, one then receives a typical Echelle spectrum, where wavelengths vary over orders and within each order (see Figure 1.6).

In summary, according to Olling (2004), the strong dependence of the radial velocity precision of a spectrograph σ_{RV} on the spectral resolving power R and on the wavelength range $\Delta\lambda$ (in $[\text{\AA}]$) accessible to extract spectral information can be approximated as




FIGURE 1.6: An Echelle spectrum of a Quartz-Tungsten halogen lamp, obtained with the Waltz spectrograph and a DSLR camera.

$$\sigma_{\text{RV}} [\text{m s}^{-1}] \approx f_{\lambda} \cdot \left(\frac{400}{S/N} \right) \cdot \left(\frac{50,000}{R} \right) \cdot \sqrt{\frac{90}{\Delta\lambda}} \cdot \sqrt{0.862 + 1.415D_A} \quad , \quad (1.12)$$

where D_A is the spectroscopic dispersion in units [$\text{\AA}/\text{pix}$]. The function f_{λ} describes the density of absorption lines for a star of given spectral type, and varies between 1.84 in the near-infrared and roughly 1 at visual wavelengths. Evaluating this expression for the Waltz spectrograph described in Section 2.2.4 results in a radial velocity precision of $\sigma_{\text{RV}} \approx 3 \text{ m s}^{-1}$ (assuming $S/N = 100$, and $\Delta\lambda = 90 \text{ \AA}$).

The above relation denotes the best possible precision for a given setup; in order to actually achieve that, all noise sources need to be well understood and, where necessary and possible, either eliminated or calibrated for. These noise sources include (but are not limited to): pixel-to-pixel inhomogeneities (both in quantum efficiency and spatial extend); spatial and temporal variability of the spectrograph PSF, for example due to temperature variations; or contamination of the spectrum by telluric lines stemming from Earth's atmosphere. Extracting high-precision radial velocity information from stellar spectra therefore requires comprehensive reduction codes that correctly incorporate and weight all these factors. In the following sections I will introduce the most important aspects and techniques used in the reduction of spectra. Chapter 3, where I present the data reduction code of the Waltz project, then builds upon these methods in more detail.

1.1.3 Reduction of spectra

Reduction algorithms for Echelle spectra all follow some basic steps: First, the acquired spectrum frames need to be pre-processed, then the Echelle orders need to be

extracted from the 2D frames, a wavelength calibration needs to be performed, and finally, if required, radial velocities can be determined from the extracted spectra.

Whereas basic image processing such as bias- and dark-frame corrections can be performed just like for "normal" astronomical observations, flat-fielding is done differently due to the special layout of the spectrum: All the light from the source is concentrated within the Echelle orders, while the pixels between the orders ideally only contain noise⁷. As each order usually spans several pixels in cross-correlation direction, corresponding to the height of the imaged object (the entrance slit or fiber), light of a given wavelength is spread out over a vertical pixel column (called the spatial direction hereafter); therefore, for each order and at each pixel position along the main dispersion direction, the flux values from the corresponding pixel column need to be combined to receive a one-dimensional spectrum.

To achieve this, first the orders need to be traced, that is their positions on the detector need to be found. This is usually achieved by fitting polynomials to the pixels with maximum counts of each order. Then, there are different possible approaches to estimate the combined flux values for each wavelength bin in an order, \hat{F}_λ . The most simple one is a sum over the (already background-subtracted) flux values of each pixel along the spatial direction i within the order, $F_{i,\lambda}$:

$$\hat{F}_\lambda = \sum_i F_{i,\lambda} \quad . \quad (1.13)$$

This simple extraction algorithm produces good one-dimensional flux estimates for spectra that are dominated by the flux of the object; however, when the background noise is significant, simple extraction does not deliver the result of lowest variance. Several works in the literature aimed at developing routines that allow for better extraction in these cases of low S/N, most of which rely on some sort of weighted sum of pixel fluxes:

$$\hat{F}_\lambda = \sum_i W_{i,\lambda} F_{i,\lambda} \quad . \quad (1.14)$$

For the simple extraction algorithm, the weights $W_{i,\lambda}$ are unity at pixel positions i within the order, and zero elsewhere. Horne (1986) instead based the calculation of weights on an estimate of the fraction of the true object flux falling into each pixel, normalized by the total flux. In this algorithm, as the true flux is not known a priori, it is assumed that the flux fractions vary smoothly over λ (i.e. in dispersion direction); the flux fractions of the pixels within a spectrum order are then modelled through a number of low-order polynomials running along the detector rows, and normalized in spatial direction (i.e. cross-dispersion direction).

The algorithm by Horne (1986) has some great advantages over other methods: It does not make any prior assumption about the exact shape of the spatial profile of the slit, which makes it widely applicable (in contrast to, e.g., Urry and Reichert, 1988, who used a Gaussian profile). Additionally, by implementing an iterative modelling of the polynomials and employing a rejection cycle, excess flux due to cosmic ray hits can be excluded from the extraction.

⁷Of course, in reality there will be some amount of scattered light that also falls between the orders; this is a noise source as we can not extract any wavelength information from it, and it can be taken care of by subtracting a smoothed image of the inter-order flux values.




FIGURE 1.7: *Left*: A section of one order of a continuum spectrum obtained with the Waltz spectrograph. The green solid line is the trace of the order, green dotted lines denote polynomials used to compute weights for the optimal extraction algorithm. The red box shows a pixel column for a wavelength λ . *Right*: Flux values for pixels within one pixel column (red crosses) along with the polynomial values $P_{i,\lambda}$ calculated in the optimal extraction algorithm (schematic).

Marsh (1989) later generalized the algorithm to also allow the extraction of highly distorted spectra: By modelling the polynomials not along the detector rows, but actually parallel to the order traces, they do not need to account for the curvature of the orders themselves. Figure 1.7, left, shows one subset of an order of a continuum-source, and depicts the polynomials (green dotted lines), positioned parallel to the order trace (green solid line); the flux values of the pixel column in the red rectangle are plotted on the right (red crosses), and the flux estimate as modelled by the polynomials is shown by the green solid line. The optimal extraction algorithm by Marsh (1989) is used in the Waltz DRS, and a more in-depth description can be found in Section 3.3.2.

Other routines use slightly different approaches, for example to construct the mean spatial profile with splines (Hinkle et al., 2000; Piskunov and Valenti, 2002). Still, in nearly all advanced spectrum reduction routines today the fundamental principles of smoothing and weighting the pixel values is a common feature. This can be performed directly on the science object spectra (especially in the case of slit spectrographs, where the spatial profile might vary strongly over time), or on a masterflat, created from several individual flatfield spectra of a continuum source; in the latter case, the weights calculated from the masterflat can then be used to extract the science object spectra.

After extraction of the spectral orders, a wavelength calibration needs to be performed, assigning reference wavelengths to all pixels along the dispersion direction. Commonly this is achieved by using spectra of calibration lamps as reference, filled with a gas that produces a well-defined set of narrow emission lines. For high-resolution Echelle spectrographs working at visual wavelengths, a lamp with a mix of thorium (Th) and argon (Ar) is most often used. With the help of an atlas that contains the order and approximate pixel positions of all significant ThAr lines in a specific spectrograph, as well as their laboratory wavelengths, the lines and their positions in any ThAr spectrum of the same instrument can be identified and the respective wavelengths assigned. A global wavelength solution for all orders and pixels of the spectrum can then be found by fitting a two-dimensional polynomial to

the identified lines, preferably in an iterative process to exclude outliers. This polynomial can be an expansion of the grating equation for small deviation angles, of the form

$$\lambda(x, m) = \frac{1}{m}P(x, m) \quad , \quad (1.15)$$

where x and m are the pixel and order numbers of the Echelle spectrum, respectively, and $P(x, m)$ is the two-dimensional polynomial (Baranne et al., 1996). The polynomial degree in x - and m -directions needs to be evaluated for each instrument individually; Section 3.3.2 describes the implementation of the wavelength calibration algorithm in the Waltz DRS.

A spectrograph that is stabilized in temperature and pressure can reach roughly 1 m s^{-1} -precision in radial velocities with the use of a ThAr lamp. Nevertheless, the lamp does have some drawbacks: The emission lines are irregularly spaced, and their density greatly decreases in the very blue and very red wavelength regimes. The latter makes it particularly unsuitable for spectrographs working in the near-infrared, which therefore use different sources (such as Uranium-Neon lamps). Additionally, the intensities of the lines vary greatly, which can lead to saturation and bleeding of pixels for the brighter ones. Finally, ThAr lamps degrade over time, which can lead to long-term instrumental shifts.

In the last two decades, the development of alternative calibration sources has made great progress, with the aim to overcome the drawbacks of ThAr lamps: Laser frequency combs can be tuned to produce a dense forest of emission lines with similar intensities, and by synchronizing them with atomic clocks, absolute wavelengths of the lines can be produced with extremely high accuracy and long-term stability, thus enabling radial velocity measurements with a precision of 1 cm s^{-1} (Li et al., 2008). However, they are very expensive and still under development. In contrast, another calibration technique, the Fabry-Perot interferometer (or etalon), is already in use in many modern spectrographs that reach $\text{sub-} \text{m s}^{-1}$ RV precisions, such as CARMENES (Quirrenbach et al., 2016) or ESPRESSO (Pepe et al., 2021). Here, light of a continuum source is coupled between a pair of parallel reflecting surfaces, which leads to constructive and destructive interference and produces a fringed pattern of evenly separated emission lines. Wavelength coverage and separation of the lines can be tuned, and the stability of the calibration depends only on the stability of the Fabry-Perot system (see Halverson et al., 2014; Reiners, Banyal, and Ulbrich, 2014; Stürmer et al., 2016, for more background). Nevertheless, absolute wavelengths of the lines are not known a priori; therefore, Fabry-Perot etalons are usually combined with ThAr lamps. The lamp is then used to regularly calibrate the etalon, which extends the wavelength solution on a wider and finer grid, thus improving the overall result.

Finally, for the computation of radial velocities from wavelength-calibrated spectra, the use of a cross-correlation function (CCF) has been proven as a robust and efficient way to incorporate information from all observed stellar absorption lines simultaneously. Usually this is done by means of a binary mask of a typical stellar spectrum, being 1 at wavelengths of absorption lines and 0 elsewhere. The CCF can be evaluated at steps of different velocities v as

$$CCF(v) = \sum_l \sum_{x,m} p_{l,x,m}(v) f_{x,m} \quad , \quad (1.16)$$

where $f_{x,m}$ is the extracted flux at pixel x and order m , and $p_{l,x,m}$ is the fraction of the l th line of the mask that falls into pixel $\{x, m\}$ at velocity v (Baranne et al., 1996). The radial velocity can then be computed as the mean of a Gaussian fitted to the CCF.

Despite its success, this simple CCF method does not extract the radial velocity information from the spectrum in an optimized way, as deep, sharp absorption lines containing more valuable RV information are not weighted sufficiently with respect to weaker, broader lines. This effect becomes especially important with the precision of today's stabilized instruments being mainly limited by photon noise. Pepe et al. (2002) therefore weighted the lines used in the CCF according to their relative depth, which improved the radial velocity dispersion stemming from photon noise by a factor of 1.3 (evaluated for the spectrograph CORALIE). A completely different approach to improve the radial-velocity computation is used in the projects HARPS-TERRA (Template-Enhanced Radial velocity Re-analysis Application, see Anglada-Escudé and Butler, 2012) and SERVAL (Zechmeister et al., 2018): They employ a direct least-squares modelling of the observation spectra with a high-S/N template spectrum. The template is created in an iterative process from the observation spectra themselves, which are shifted by their relative barycentric and radial velocity offsets and then merged. In comparison to RV data extracted with the standard HARPS DRS (data reduction software), both these algorithms show a slight improvement.

However, when using the CCF method and relying on wavelength calibration spectra obtained only in the beginning and end of observation nights, spectrographs need to be highly stabilized to allow the determination of precise radial velocities. The HARPS spectrograph (Mayor et al., 2003) set new standards in this regard, with temperature fluctuations below 0.01 K, and reaching long-term accuracy at the 1 m s^{-1} -level and below. In less stable instruments, pressure and temperature variations can lead to significant changes of the refractive index of air and expansion and contraction of opto-mechanical components, causing changes of the PSF and wavelength shifts over time. If reference spectra for wavelength calibration are only sparsely obtained, they might not account for these changes. Many fiber-fed spectrographs therefore use a second fiber, sitting right next to the main science fiber, that produces a second set of orders and can be used for simultaneous acquisition of reference spectra, thus being able to also track short-term variations and improve the wavelength accuracy to some degree. Such a design requires larger optical elements though, which can quickly lead to high manufacturing costs. Many smaller and more financially restricted projects, such as the Waltz telescope project, thus rely on a different method for the extraction of radial velocities, which is presented in the following section.

1.1.4 The iodine-cell method

The problems of temporally varying PSFs and wavelength solutions in unstabilized spectrographs can be overcome by simultaneously acquiring stellar and reference spectra — in the same spectrum. To achieve this, a gas-cell containing a gas with known absorption spectrum is placed into the light path between the telescope and spectrograph, thus imprinting the spectral features of that gas onto the stellar spectrum. Any temporal changes of the instrument then affect the spectral content of

both sources exactly the same, whereas the radial velocity manifests itself as a shift of the stellar absorption lines with respect to the reference lines.

After a long search for suitable gases to be used in the absorption cell, molecular iodine (I₂) was chosen as the best overall compromise as it fulfills several important requirements (Marcy and Butler, 1992):

- it allows to use a comparably large wavelength range between 5000 ~ 6200 Å;
- over this wavelength range, its spectrum shows a very dense forest of narrow absorption lines with at least two features per Å;
- due to its strong line absorption coefficient, it requires a path length of just a few centimeters at pressures around 0.001 atm;
- it can be operated at easily achievable temperatures around 50° C;
- it is chemically stable and nonlethal.

The observed spectrum of the star superposed by the I₂ features can be modelled directly; however, many instrumental effects influence the exact shape of the overall spectrum, which would require an extremely complex parametrization. Therefore, the spectrum is split up into many small chunks of typically ~ 2 Å in length, and it is assumed that within these chunks large-scale effects can be neglected. Each of these chunks can then be modelled as a product of the intrinsic stellar spectrum I_s and the spectrum of the I₂ cell T_{I_2} , convolved with the line-spread-function (LSF)⁸ of the spectrograph:

$$I_{\text{obs}}(\lambda) = k [T_{I_2}(\lambda) \cdot I_*(\lambda + \Delta\lambda)] * \text{LSF} \quad , \quad (1.17)$$

where k is a normalization factor (describing the continuum), $\Delta\lambda$ is the radial velocity shift, and $*$ represents convolution (see Figure 1.8, Butler et al., 1996).

The above relation needs to be evaluated in pixel-space (x) in order to be compared to the observed spectrum, and therefore requires a wavelength model, $\lambda = \lambda(x)$. Typically, a simple first-degree polynomial suffices. The wavelength-shift of the stellar template I_* can be described by $(\lambda + \Delta\lambda) = \lambda(x) \cdot (1 + v_{\text{meas}}/c)$ (using the Doppler shift, Equation 1.1).

The exact shape of the LSF strongly depends on the instrument used, and therefore its functional description must be selected with great care: It needs to be flexible enough to mimic the true LSF and incorporate its temporal variability, but should not permit to include noise-induced artifacts (Valenti, Butler, and Marcy, 1995). A standard approach is to express the LSF as the normalized sum of a central Gaussian with several satellite Gaussians on each side at fixed positions and with fixed widths, letting only the amplitudes of the Gaussians vary in the model. This way asymmetries in the wings of the LSF can be followed very well.

The model in Equation 1.17 is evaluated on an oversampled pixel grid (e.g. with an oversampling of 4 in Butler et al., 1996), in order to allow for correct convolution; therefore, the input spectra of the star, I_* , and the pure I₂ transmission, T_{I_2} , both are required at a much higher resolutions than that of the observed spectrum. For the I₂ cell, this can be easily achieved by recording its transmission spectrum in a Fourier

⁸We use LSF instead of PSF from hereon, as the I₂ method is performed on extracted 1D-spectra.




FIGURE 1.8: Basic steps in the I2-cell method, from top to bottom: The reference I2 spectrum T_{I2} , the stellar template spectrum I_s , the model convolved with the PSF (solid line) along with the pixel-to-pixel observation (dots) I_{obs} , and the residuals between the model and the observation (dots). *Credit:* With changes from Butler et al. (1996)

Transform Spectrometer (FTS), thus achieving resolving powers of typically several 10^5 at high S/N .

The stellar template spectrum I_* however is more difficult to obtain. The typical recipe is to basically reverse the above method, in the following way: A stellar spectrum of the target star is obtained with the same instrument as used for later RV observations, but *without* the I2 cell. The features in this spectrum are only those of the star, smeared out by the instrument LSF. The stellar observation is framed by observations of bright, rapidly rotating O- or B-stars *with* the I2 cell; due to their high temperatures and fast rotation rates, these stars essentially do not have any spectral features (in the required wavelength range), and their spectra therefore contain only the I2 transmission spectrum, also broadened by the instrument LSF. The observed B-star spectra are now modelled by Equation 1.17, using a simple continuum function instead of the stellar template I_* , and this way the LSF of the instrument at the time of the observations is found. The stellar template I_* is finally constructed by deconvolving that LSF out of the stellar observation (Butler et al., 1996).

Using the I2 method on the Hamilton spectrograph, Butler et al. (1996) demonstrated that they could achieve RV measurements with long-term errors of $\sim 3 \text{ m s}^{-1}$ for spectra with $S/N \approx 200$; their estimated error budget in the limit of photon noise suggested errors of only 2 m s^{-1} , so some systematic instrumental effects were probably not accounted for yet. More recently, for radial velocity measurements obtained with the Stellar Observations Network Group (SONG) telescope and spectrograph on Tenerife, and using the I2 method, Andersen et al. (2016) reported a short-term precision as low as 1 m s^{-1} for some optimal stars, and a long-term stability on the order of 5 m s^{-1} . In Chapter 3, I will describe the mathematical formalism and code implementation of the I2 method in the Waltz DRS in detail, and present first results which seem to reach the same level of precision as the dedicated reduction pipelines of Lick and SONG.

1.2 Exoplanets around evolved stars

To date, at least 112 planets in 102 systems have been detected around evolved stars⁹ — quite a small fraction of the overall number of known exoplanets, most of which orbit main-sequence (MS) stars. Nevertheless, we can learn a lot from RV surveys of evolved stars and resulting planet discoveries: First, they allow us to probe the effects of stellar evolution on planetary systems, as the host star leaves the main sequence to first become a red-giant-branch (RGB) and later a horizontal-branch (HB) star. Second, in comparison to RV surveys of MS stars, surveys of evolved stars can easily be extended to higher stellar masses: MS stars above $\sim 1.5 M_{\odot}$ show only few and quite broadened absorption lines, due to their high temperatures and fast rotation rates, which makes them unfavourable candidates for precise RV measurements; their evolved counterparts however are cooler, and their rotation has slowed down, resulting in a spectrum with many narrow absorption features. Therefore they are perfectly suited targets for RV surveys.

The growing number of planet detections around giant stars allows us to draw statistical conclusions about the dependence of the planet occurrence rate on stellar and orbital parameters. Due to the different parameter space available to RV surveys of giant stars in contrast to those of MS stars, namely larger stellar masses and later evolutionary stages, these results should complement known planet distributions and contribute to our general understanding of the formation and evolution of planetary systems.

1.2.1 The Lick RV survey of evolved stars

The Lick survey started in June 1999, by measuring radial velocities for a sample of originally 86 evolved K-giant stars. The objective was to demonstrate the capability of these targets to serve as stable grid stars for the Space Interferometry Mission (SIM), which was cancelled in 2010. The selection criteria of the program required the stars to be neither variable nor part of a multiple system (Frink et al., 2001). One year after the survey started, another 96 K-giants were included in the sample with less stringent constraints concerning photometric variability.¹⁰ Finally, in 2004 the survey was extended again by 194 G and K-giant stars with on average higher masses and bluer colors. The complete sample therefore consists of 373 evolved stars, with masses mostly between 0.8 and $5 M_{\odot}$, and magnitudes brighter than 6 mag.

The observations were conducted with the 0.6m Coudé Auxiliary Telescope (CAT) at Lick observatory in conjunction with the Hamilton spectrograph, and precise RVs were determined using the I2 method (compare Section 1.1.4; Vogt, 1987; Butler et al., 1996). Between 1999 and 2012, observations were carried out between six and eight nights per month, with exposure times between 5 and 30 minutes for stars of 3 mag and 6 mag, respectively. Thus, typical S/N of the spectra is 80 – 100, and the resulting RV precision is $3 \sim 8 \text{ m s}^{-1}$. This does not only allow to identify stellar binaries (the original goal of the survey), but also planetary companions with masses $\gtrsim 1 M_{\text{J}}$.

⁹<https://www.lsw.uni-heidelberg.de/users/sreffert/giantplanets/giantplanets.php>, as of May 6th, 2021

¹⁰Three of these stars were later excluded again, as they were found to be visual binaries.




FIGURE 1.9: Planet occurrence rate as a function of metallicity and stellar mass in the Lick survey. The "first sample" comprises the 186 stars from the start of the survey, while the "second sample" denotes the 196 stars added later on. The fraction of stars with a planet and/or planet candidate of all stars within each bin is given by the percentage numbers. *Credit: Reffert et al. (2015)*

The first exoplanet discovery of the survey around ι Draconis (Frink et al., 2002), which also marked the first exoplanet detection around an evolved star, shifted the primary focus to investigate further candidates in the sample and study planet occurrence properties of these giant stars. Since then, more than a dozen additional planets were found around stars from the Lick sample, some of them in multi-planetary systems or stellar binaries, as well as brown dwarfs (Frink et al., 2002; Reffert et al., 2006; Mitchell et al., 2013; Trifonov et al., 2014; Ortiz et al., 2016; Luque et al., 2019; Quirrenbach et al., 2019; Tala Pinto et al., 2020). Additionally, the RV measurements helped to reject proposed planet candidates and to reveal yet not fully understood phenomena (see e.g. Reichert et al., 2019; Heeren et al., 2021).

Unfortunately, in 2012 the I2 cell at Lick observatory was damaged and RV measurements at the site could not continue (Fischer, Marcy, and Spronck, 2014). For some stars with strong planet candidates, additional observations could be performed with the SONG telescope on Tenerife, and in a few cases with the HARPS spectrograph in La Silla/Chile. In the future, a more frequent observing program is planned again, using the Waltz telescope at the Landessternwarte Heidelberg (see Chapter 2).

1.2.2 Occurrence rate of planets as function of stellar mass and metallicity

Using the planet and planet candidate detections from the Lick survey, Reffert et al. (2015) estimated occurrence rates of giant planets around the giant stars in the sample as a function of stellar properties. Figure 1.9 shows the distribution of these stars depending on their stellar mass and metallicity, and stars hosting planets or planet candidates are indicated. As can be seen, the total number of detections strongly increases with metallicity, and the highest detection rates are reached for values larger than solar metallicity. This positive planet-metallicity correlation agrees well with results from RV surveys of both MS and subgiant stars (compare e.g. Fischer and Valenti, 2005; Johnson et al., 2010), and can be well explained by planet formation

theory: Stars with higher metallicity feature metal-richer proto-planetary disks during their birth, which in turn offer more building material to form planet seeds.

From Figure 1.9 it is also obvious that the number of planet detections depends on the stellar mass bin. Reffert et al. (2015) modelled the distribution of confirmed detections (i.e. without candidates) in stellar masses with a Gaussian function, and found them to peak at $M_* = 1.9^{+0.1}_{-0.5} M_\odot$; above masses of $\sim 2.7 M_\odot$, the planet occurrence drops down to zero. Similarly, in their EXPRESS¹¹ survey of 166 evolved stars, Jones et al. (2016) report a peak occurrence rate at $M_* \approx 2.3 M_\odot$ and a sharp drop at stellar masses above $\sim 2.5 M_\odot$. These results extend the findings from RV surveys of MS stars, which are mostly limited to stellar masses $\lesssim 1.5 M_\odot$. In that limited mass range, they generally report an increase of giant planet occurrence rate with mass (Johnson et al., 2010), which agrees with the results from the giant star surveys.

As for the planet-metallicity correlation, the increase of planet occurrence with stellar masses up to the inversion point reflects the planet formation history of the observed systems: The total proto-planetary disk mass generally is correlated with the stellar mass, and higher disk masses lead to an increase of the formation rate of planets, particularly of giant planets. Models (e.g. by Kennedy and Kenyon, 2008) however show that at a certain stellar mass this positive correlation breaks down due to a combination of several effects: As the snow line of the proto-planetary disk gets pushed further out for more massive and thus more luminous stars, planets around these stars need to form at larger orbital separations, where the growth rates are slower. Additionally, the migration time scale increases for more massive stars, and the disk dispersion happens faster. As a result, the formation of giant planets at intermediate orbital separations might be halted.

In a more recent analysis, Wolthoff et al. (in preparation) combine results from the Lick and EXPRESS surveys with the ones from the PPPS¹² survey to analyze a large sample size of 482 evolved stars with 37 planets detected around 32 host stars. To ensure compatibility of stars from different samples, all stellar parameters were newly computed using a Bayesian approach (as in Stock, Reffert, and Quirrenbach, 2018). Additionally, the whole sample was corrected for survey completeness. In their analysis, they confirm the planet-metallicity correlation and find a peak occurrence rate stellar mass of $M_* \approx 1.68 M_\odot$. The discrepancy to the values from Reffert et al. (2015) and Jones et al. (2016) might be explained by the fact that the stellar masses used especially in the latter publication were partially strongly over-estimated; it is however also possible that the differences between the determined peak masses are simply of statistical nature, and therefore not significant (compare the large uncertainties of the value from Reffert et al. (2015)).

1.2.3 Overall occurrence rate and period distribution of planets

From the completeness-corrected sample analyzed by Wolthoff et al. (in preparation), the estimated global occurrence rate for giant planets around evolved stars is $\sim 10.7\%$; similarly, Wittenmyer et al. (2020b) computed a corrected occurrence rate around 7.8% for a much smaller sample of low-luminosity giant stars from the PPPS alone. These numbers are in good agreement with occurrence rates for MS stars: Cumming et al., 2008 for example found that roughly 10.5% of all MS stars harbor

¹¹EXoPlanets aRound Evolved StarS, (Jones et al., 2011)

¹²Pan-Pacific Planet Search, (Wittenmyer et al., 2011)

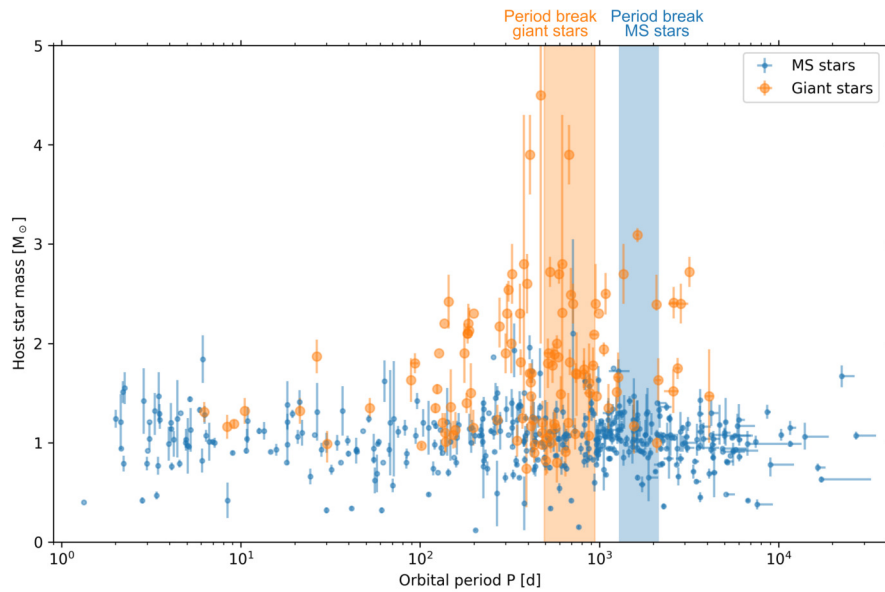


FIGURE 1.10: Distribution of MS (blue) and evolved stars (orange) that host giant planets ($M_p \leq 0.1M_J$) detected by the RV method. The shaded regions denote the period breaks of the occurrence rates, blue for MS stars (Fernandes et al., 2019) and orange for evolved stars (Wolthoff et al., in preparation). Credit: NASA Exoplanet Archive for MS stars, <https://www.lsw.uni-heidelberg.de/users/sreffert/giantplanets/giantplanets.php> for evolved stars.

giant planets with masses between 0.3 and $10 M_{Jup}$ on orbits with periods between 2 and 2000 d.

Whereas the global planet occurrence rates as well as the dependency on stellar mass and metallicity seem to agree well between RV surveys of MS and evolved stars, there is a strong disparity in the orbital configuration space: Figure 1.10 shows giant planet detections from RV surveys for both MS and giant stars, with the orbital period of the planet on the x-axis and the host stellar mass on the y-axis. Around MS stars, quite a number of so-called Warm and Hot Jupiters with orbital periods below ~ 100 to 300 d have been detected; evolved stars in contrast seem to harbor far less of these close-in planets, especially when moving to higher stellar masses. In fact, the completeness-corrected sample mentioned above suggests a giant planet occurrence rate of only roughly 1% at orbital periods between ~ 80 to 300 d, and probably an even further decrease towards lower periods.

The problem of missing close-in planets around evolved stars is known for some time now, and there are several possible explanations:

1. As stars with masses up to $\sim 2 M_\odot$ move upwards on the RGB sequence, they largely inflate their outer layers. Very close-in planets around these stars, such as Hot Jupiters, could then be engulfed directly; planets orbiting a little further out might be forced to migrate inwards due to tidal interaction with the star and eventually also be swallowed. When the stars leave the RGB, only planets with periods $\gtrsim 300$ d might be left over.

However, stars with masses $\gtrsim 2 M_\odot$ do not experience a helium flash and therefore never increase their radii as drastically. Therefore, one would expect

close-in planets to still exist around the evolved counterparts of these stars, if the evolution/engulfment scenario was the only phenomenon dictating the orbital configuration disparity between MS and evolved stars. As for these more massive stars close-in planets are particularly rare (see Figure 1.10), there must be additional effects at play.

2. As mentioned in Section 1.2.2, giant planet occurrence generally drops for stars with masses $\gtrsim 2 M_{\odot}$, at least up to orbital separations of a few AU, most probably due to a decrease of planet formation and migration rates. With RV surveys of evolved stars covering these more massive objects, in contrast to MS RV surveys, it is possible that the fewer numbers of close-in planets simply reflects that stellar-mass-dependency of planet formation rate — or, this effect might complement the stellar evolution impact described in point 1.
3. Finally, the low number of planets detected on close-in orbits around evolved stars might at least partially be due to bias effects: Whereas global occurrence rates of giant planets seem to be quite similar for giant and MS stars, overall far fewer giant planets have been actually found around the former due to a smaller number of targets and observations. Close-in planets around giant stars might be actually there but successfully hiding from RV surveys to date. In fact, Wittenmyer et al. (2020a) found that for MS stars "Cool" Jupiters with orbital periods $\gtrsim 100$ d are about 8 times more frequent than Hot Jupiters; they also found a distinct step in the period distribution around ~ 300 d, below which the giant planet occurrence rate is between 0 to 2%. At least the number of Warm Jupiters around evolved stars therefore might still increase with more observations.

Another disparity between MS and giant stars becomes apparent at longer orbital periods: Fernandes et al. (2019) found an increase of the giant planet occurrence rate around MS stars up to orbital periods between ~ 1200 to 2200 d, beyond which it decreases again. The position of the period break (blue shaded region in Figure 1.10) corresponds roughly to the snow line for a Solar-mass star, and planet formation models indicate that it could be an imprint of the planet formation and migration history. Wolthoff et al. (in preparation) found a similar turnover in planet occurrence around giant stars, but the peak seems to be at smaller periods around ~ 700 to 800 d). This is much closer to the host star than the snow line; similarly as for the missing close-in planets, it is unclear to date whether this apparent shift of the peak occurrence period is an effect of stellar evolution, of higher stellar masses or caused by observational biases.

1.3 RV variations beyond the Keplerian model

In Section 1.1.1, I have discussed the RV signature caused by a planet on a perfect Keplerian orbit, and how to model and interpret it. In reality however the movement of planets can deviate substantially from Keplerian orbits under some circumstances, and completely different astrophysical phenomena can mimic planet-induced RV variations. This can lead to a wrong interpretation of the measured RVs, and it is therefore extremely important to take other effects into account and rely on several confirmation methods before publishing a putative planet detection. In the following sections, I will introduce additional causes of RV variations and explain their impact on the K-giant survey at Lick, and I will introduce the dynamical modelling of planet-induced RVs, which takes deviations from Keplerian orbits into account.

All these effects and appropriate analysis tools are discussed and used in relation to the K-giant star ϵ Cygni in Chapter 4.

1.3.1 Solar-like oscillations

Similar to the sun, most stars exhibit short-period oscillations which are excited by convection processes close to the surface. These so-called pressure- or p-mode oscillations give rise to intrinsic photometric and RV variations of the observed stars, which act as noise source in RV surveys of exoplanets. Both the periods and the amplitudes of the oscillations depend on stellar properties, and scaling relations predicting these quantities have been developed and gradually improved for example by Kjeldsen and Bedding (1995), Kjeldsen et al. (2008), and Kjeldsen and Bedding (2011). The p-mode periods range from several minutes for solar-like stars to hours and few days for giants; as this is much shorter than the typical observing sampling of RV surveys such as the Lick planet search around evolved stars, the oscillation-induced RV variations act as white noise or jitter in the measurements.

While the RV amplitudes of the p-mode oscillations are quite low for solar-like stars (e.g. roughly 0.2 m s^{-1} for the sun, see Kjeldsen et al., 2008; Fredslund Andersen et al., 2019), they can reach several 10 m s^{-1} in late-type evolved stars, thus hindering the detection of giant planets that give rise to RV amplitudes much smaller than the oscillation-induced jitter. From our K-giant survey a clear correlation between observed RV scatter and the $B - V$ color of the observed stars can be deduced, with redder stars showing larger intrinsic RV variability (Frink et al., 2001; Trifonov et al., 2014). Most of the stars in the survey however have fairly stable RVs with scatters around 20 m s^{-1} .

It is important to note that while RV variability stemming from p-mode oscillations need to be taken into account as noise source when searching for planets, due to the short periods of the oscillations they usually cannot be confused with planet-induced signals. Nevertheless, when achieving a sufficient RV sampling and precision of the p-modes, an asteroseismic analysis of the data can extract stellar properties such as the mass, radius or luminosity of the star (e.g. Grundahl et al., 2017). The results deduced from asteroseismology are often much more precise than the ones from traditional approaches such as evolutionary tracks, which can be ambiguous.

1.3.2 Gravity-mode and mixed-mode oscillations

In addition to the p-mode oscillations discussed before, stars exhibit oscillations where, instead of pressure, buoyancy acts as restoring force — so-called g-mode oscillations (as buoyancy is coupled to the gravitational force). These oscillations are constrained to the radiative zones of stars, as they quickly decay in convective environments; therefore, for solar-type MS stars and evolved G- and K-giants that have convective envelopes, g-mode oscillations are generally not observable through photometry or RV measurements.

Saio et al. (2015) however presented a theoretical model where non-radial mixed modes may occur in the deep convective envelopes of RGB stars and have amplitudes large enough to be observed. These so-called oscillatory convective modes are coupled to high-order g-modes in the deep interior of the star, and have periods of several 100 d. They might therefore be an explanation for the long secondary period (LSP) variations that have been detected in photometric observations of many highly luminous RGB stars, and also found in RV measurements for some of these

stars (e.g. Hinkle et al., 2002). As the periods correspond to the orbital periods of many exoplanets discovered around evolved stars, these oscillations could potentially lead to false-positive detections of planets (see Hatzes et al., 2018; Reichert et al., 2019, and Chapter 4 in this thesis).

1.3.3 Stellar spots

Stellar spots are regions in the photosphere of stars with high magnetic field flux, causing strong convection and thus a decrease in temperature and luminosity within the spot region. When taking spectra of a star, due to the stellar rotation a spot thus leads to a periodic modulation of the absorption line shape with the rotation phase, by reducing the amount of blue-shifted light received by the telescope when sitting on the approaching side of the star, and vice-versa. In RV measurements this can appear as a periodic RV signal, and might easily be mistaken as a planetary signature. The period of that signal however corresponds to the rotation period of the star (or is an alias of it). Several techniques exist to test whether a signal is caused by stellar spots: As the total luminosity of a star is also affected by spots the signature should also appear in precise photometry timeseries. Furthermore, absorption lines whose depths are strongly affected by stellar spots can be analyzed (e.g. H α , CaII, see Kürster et al., 2003), and if the resolving power of the spectrograph is high enough bisector spans of the absorption lines can be measured by analyzing the change of line shape over time (Figueira et al., 2013). Additionally, as the strength of a spot-induced signal decays from blue wavelengths towards the infrared, this so-called chromaticity measurement can serve as a tracer (Zechmeister et al., 2018).

The strength of spot-induced signals however strongly depends on the size of the spot (or a group thereof) relative to the stellar surface as seen by the observer. Thus, many M-dwarfs are heavily affected, as they are often very active and have comparably small surface areas; for these stars spot-induced RV variations of several m s^{-1} can easily occur, and observers need to be careful not to confuse them with RV signals from rocky planets (e.g. Tal-Or et al., 2018). For evolved stars such as the ones in our K-giant sample however it would be surprising to find RV modulations with high amplitudes caused by spots: As these stars typically have radii on the order of 10 to 50 R_{\odot} , the spots or spot groups would have to be extremely large to cause noticeable RV signals, and photometry measurements of these stars should then reflect that. For instance, simulations for the K-giant star and member of our sample ϵ Cyg, using the star spot simulator *SOAP 2.0* (Dumusque, Boisse, and Santos, 2014), resulted in a flux variability on the order of 3% in order to achieve RV variations with an amplitude of 25 m s^{-1} (see Section 4.6.2, or Heeren et al., 2021). One of the selection criteria for the K-giant sample however was a low photometric variability, and the HIPPARCOS photometry of ϵ Cyg only varies with 0.34%. We can therefore safely assume that spot-induced RV signals should not play a big role in our survey.

1.3.4 Dynamical analysis of RV measurements

Keplerian orbits describe the movement of two point masses (i.e. a star and a planet in our case) around their common center of mass, with the only force present being their mutual gravitational interaction (and ignoring relativistic effects). Even when adding a third body to the system, the Keplerian model might still give a good approximation of the actual movement of the bodies over some limited time, if the distances between the bodies are always large and their masses are relatively small. For densely packed systems however, for instance multiple planets on close orbits or a

planet in a close stellar binary, significant deviations from Keplerian orbits can occur even on short timescales due to the gravitational forces acting between the bodies. Consequently, RV signatures of these systems will show residuals when modelled with the standard Keplerian description, and the orbital and companion parameters from these models might not give a good representation of the system.

In these cases of multiple-companion systems, fitting the RV measurements with a fully dynamical model based on a numerical N -body integrator which takes inter-companion forces into account can improve the results and give better estimates of the true system properties. Furthermore, this approach can help to distinguish between "true" multi-planetary RV signals, which should be well described by the dynamical model, and false positives stemming from internal stellar sources as discussed in the previous sections; for instance, an RV curve consisting of two signals with nearby periods, one planet-induced and one caused by stellar spots, can look indicative of a multi-planetary system at first glance, but this hypothesis might be disproved by a dynamical model (under the right circumstances, e.g. observations over many orbital periods and high S/N).

A large number of simulators for N -body systems bound by gravitational forces exist today, for instance the Fortran-based *SWIFT* (Duncan, Levison, and Lee, 1998) and *MERCURY* (Chambers, 1999) packages, or the Python-based *REBOUND* (Rein and Liu, 2012), all of which feature symplectic integrators and thus allow to reliably model also close encounters between the system's bodies.

Additionally, when analyzing a putative multiple-companion system with non-negligible gravitational interaction between the companions, one should always test for long-term dynamical stability as we generally expect stellar and planetary systems to be stable over timespans of at least several 10^5 yr. If dynamical tests of a given model parameter combination show a chaotic behaviour of the system, possibly with one or more companions being ejected from the system or collisions occurring, this parameter combination is probably not a true representation of the system; after all, it would be extremely unlikely to observe a system just before its disintegration. As the stability of a simulated system can change abruptly even with only small alterations to the starting parameters, one needs to explore the parameter space around the best-fit model and test the stability for many different parameter combinations (usually on the order of 10^4 to 10^5 sets). This way, regions of stability within the parameter space can be identified, which again helps to gain a better understanding of the true system configuration.

Stability tests can be performed using the same N -body integrators as described above, which also allow to gain an understanding of possible resonant behaviours occurring between companions of a system (see e.g. Trifonov et al., 2018). However, when integrating over long model timespans these simulators require a lot of computation time, and as the stability tests need to be performed for many different parameter combinations around the best fit, the overall computational expense quickly sums up. Analytic stability tests (e.g. Gladman, 1993) can be performed much quicker, but they do not give detailed insights into inter-companion resonances and are usually not safely applicable to all system configurations. Recently, Tamayo et al. (2020) introduced the Stability of Planetary Orbital Configurations Klassifier (SPOCK), which makes use of machine-learning techniques to gain a probability estimate of the long-term stability of multi-planetary systems, and thus combines comparably quick computation with a higher true-positive rate than other analytic or semi-analytic tests.

In my research on the stellar binary system ϵ Cyg, presented in Chapter 4, I applied both dynamical models and stability analysis tests to check whether the short-period RV variations observed for that star might be caused by a planetary companion. For the dynamical models, I used a modified version of the Bulirsch-Stoer N -body integrator of *SWIFT* in combination with a Levenberg-Marquardt minimization algorithm; for the stability tests, mostly a modified version of the Wisdom-Holman integrator in the same package was used. The same basic Fortran routines had been used previously for instance in Trifonov et al. (2018), and they are employed in the background by the Python-based RV and transit data analysis tool Exo-Striker, which thus also allows to perform broad dynamical tests and models (Trifonov, 2019).

THE WALTZ TELESCOPE PROJECT: AN OVERVIEW AND RECENT UPDATES

In the following, I will give an overview of the current status of the Waltz project at the Landessternwarte Heidelberg, including all important hardware and control software, as well as early results. In the build-up and testing of the instrument components, my main contributions were the following: I designed and integrated an addition to the detector and camera stand; I changed the calibration fiber and characterized its throughput on the pinhole mirror; I optimized the design of the FFU, assembled and tested it in the laboratory, aligned it at the telescope and measured its efficiency. This work is described in Section 2.4. Furthermore, I obtained and analyzed the data presented in Section 2.3; the tracking test in Section 2.3.1 was performed together with bachelor student Lars Wolfgramm, who wrote the required algorithms to obtain the raw data. Finally, I also re-aligned the spectrograph and frontend prior to first on-sky observations.

The chapter is organized as follows: Section 2.1 gives an introduction to the Waltz project, and in Section 2.2 I give an overview of the instruments and controls. In Section 2.3 I present some results from tests and first observations, which should indicate the current performance of the whole system. I then describe my own main opto-mechanical contributions to the project in Section 2.4, before I close with a summary in Section 2.5.

2.1 Introduction

Starting in 1999, the exoplanet group at the Landessternwarte (LSW) Heidelberg regularly carried out RV measurements of a sample of evolved stars using the Hamilton spectrograph at Lick observatory. These observations ended in 2012 when the I2 cell used to measure the radial velocity shifts was damaged. Several sub-stellar companions have been published already on the basis of the Lick RV data (e.g. Frink et al., 2002; Trifonov et al., 2014; Ortiz et al., 2016), allowing for statistical analyses of occurrence rates (Reffert et al., 2015). Possibilities of continuing the observations were searched for, in order to extend the baselines for several promising candidates and investigate long trends seen in the data. For some stars, additional RV measurements were performed with the SONG telescope on Tenerife (Andersen et al., 2014; Fredslund Andersen et al., 2019; Grundahl et al., 2017), and in a very few cases with the HARPS spectrograph in La Silla/Chile (Mayor et al., 2003). Generally however, the dependence on the very limited observation times at telescopes available

through public application is unfavourable for a RV survey of a stellar sample with the focus of exploring statistical properties.

Therefore, the decision was made to build an own instrument, located and operated at the Landessternwarte Heidelberg. This would allow to make use of every clear night and have long-term control over the observations. The historic Waltz telescope at the institute on top of the Königstuhl, which is similar to the CAT telescope used originally at Lick, was equipped with a fiber-coupled, high-resolution spectrograph. Due to the limited budget of the project, it was decided to refrain from building a temperature- and pressure-stabilized instrument; instead, precise radial velocities will be extracted using the iodine-cell method, as was done for the Hamilton spectrograph at Lick (Butler et al., 1996). While being mainly dedicated to science, the project also serves an educational purpose, with students contributing to the implementation of software and mechanics, analysis of diagnostic data and operation of observations.

Long-term goal of the project is to allow fully remote observations from a distance, and all sub-systems have been developed accordingly. In the current status, a central computer already controls most parts of the system. First light on sky was achieved on 9 November, 2020.

2.2 Instrument description

2.2.1 Telescope and dome

The Waltz telescope has been constructed as a Newton telescope in 1906, and was later modified to be used in the Nasmyth-focus. Its primary mirror diameter is 72 cm, and in an experiment performed by Tala Pinto (2019), the focal ratio was measured to be $F/\# = 20.4$; the focal length of the telescope therefore is $f_T = 14.7$ m, and the plate scale is $s = 14.04 \frac{\text{arcsec}}{\text{mm}}$. All three telescope mirrors are aluminium-coated, and focussing of the telescope is achieved by repositioning the secondary mirror (M2) in vertical direction along the light path inside the telescope tube.

The telescope is mounted in a fork in an equatorial setup. It was equipped with modern stepper motors for movement in the declination (DEC) and right ascension (RA) axes. The motors are controlled by the publically available *DELTACODE* telescope control system (Schneider, 2013), which allows communication with the main PC via an RS-232 serial connection, using the LX200 protocol. The telescope control system will be upgraded with encoders measuring the telescope movement in the DEC and RA axes in the near future, to get a true feedback of the position of the telescope at any time.

The Waltz telescope sits on the second floor of a building separate from the main LSW building, under a wooden dome that is encased with iron plates. Rotation of the dome is achieved by a belt-system connected to a powerful servomotor, which is controlled through the off-the-shelf dome control system *ScopeDome* (ScopeDome, 2016), connected to the main PC via USB. When in operation, the dome can be slaved to the telescope, to always position the dome slit in the same azimuthal direction. As of now, the slit still needs to be opened and closed manually by hand, using a rope; in the future, it will be automated with a cable winch controlled by the central computer.




FIGURE 2.1: Model of the Waltz frontend, depicted in observation mode with pickup mirror moved in to redirect the light beam coming from the telescope. Pickup mirror and I2 cell can be moved by stepper motors on linear stages. With the pickup mirror moved out, light coming from the calibration fiber can pass towards the fiber feeding unit (FFU).

2.2.2 The frontend

On the side of the telescope, where the light beam exits the tube in the Nasmyth-focus, sits the frontend, which fulfills several tasks: (a) coupling the light coming from the telescope into the optical fiber leading to the spectrograph; (b) coupling light of calibration lamps, injected through the calibration fiber, into the spectrograph fiber; (c) driving the iodine cell into the light beam if needed; (d) allowing to precisely point to and guide on the target star.

Figure 2.1 shows a 3D-model of the frontend: In observation mode a pick-up mirror sits at a 45° -angle in front of the telescope exit hole and redirects the light beam coming from the telescope sideways. The mirror can be moved out of the light-path on a linear stage, enabling light from the calibration fiber to pass; this way, the calibration light takes exactly the same path as stellar light after it exits the telescope tube, ensuring that all instrumental effects impact both science and calibration spectra the same.

Next to the pick-up mirror sits the iodine cell, enclosed by a housing to ensure thermal and mechanical stability, both developed and constructed by the University of Göttingen. The cell itself is cylindrical, with a length of 100 mm and diameter of 50 mm, and made of Quartz glass. The front and back windows are slightly tilted to prevent fringing of the light passing through the cell. It sits in a robust housing, which contains heating elements and temperature sensors; thus, using a closed-loop control system, the iodine cell can be heated to temperatures around 50°C , where the molecular iodine contained in the cell becomes fully gaseous. The long-term thermal stability is roughly 0.1°C . Insulation material in the housing minimizes dissipation of heat into the frontend to ensure stability of the opto-mechanical components. The complete structure is mounted on a linear stage, so that it can be moved into and

out of the light beam. When moved in, as the star light passes through the cell, the absorption features of the iodine are imprinted onto the stellar spectrum. Radial velocities can then be extracted by modelling the relative shifts of the stellar absorption lines with respect to the iodine lines over time.

Behind the iodine cell comes the guiding unit; both pointing and guiding of the telescope is performed on a pinhole mirror, that sits in the focal plane of the telescope. The mirror is tilted at an angle of 15° , thus redirecting a telescope field-of-view of 3×3 arcmin through a lens into the guiding detector. We use a Starlight Xpress Lodestar X2 Autoguider, with 752×580 pixels and $8.6 \times 8.3 \mu\text{m}$ pixel size¹, which is connected to the main computer via USB. This presents live feedback of the target star, and enables precise pointing and guiding of the telescope such that the stellar light falls into the pinhole of the mirror. An optical filter is placed at the entrance of the guiding CCD, with an allowed spectral range of 4500 to 6500 Å, thus ensuring that guiding is performed in the wavelength of interest dictated by the iodine-cell method. The pinhole used for guiding measures 265 μm in diameter, corresponding to an angular size of $3.7''$ on sky; this value is slightly larger than the average seeing conditions at the telescope site on Königstuhl/Heidelberg, and was chosen in order to optimize the light-throughput without introducing too severe guiding errors.

Directly behind the pinhole mirror sits the fiber-feeding unit (FFU), which, if aligned correctly, ensures coupling of the infalling light from the pinhole into the spectrograph fiber. The FFU houses a single doublet achromat lens with a focal length of 3 mm and diameter of 2 mm, which demagnifies the seeing disk of the star by a factor of 0.264. For an average seeing of $3.5''$, this translates to an image size of 66 μm . When using the original spectrograph fiber with a rectangular cross-section of $33 \times 132 \mu\text{m}$, the FFU therefore leads to a light-loss of roughly 40% due to over-filling (in addition to regular coupling losses). However, the spectrograph fiber will be upgraded with a microscopic image slicer 3D-printed onto the tip of the fiber, transforming the circular PSF into a pseudo-slit before injection into the fiber. Early tests in the laboratory deliver a coupling efficiency of the upgraded fiber of $\sim 80\%$ at a wavelength of 780 nm, slightly degrading towards bluer wavelengths. Throughput simulations in contrast suggest a theoretical efficiency of roughly 90%, and the reason for the discrepancy is not fully understood yet. Most probably some light is scattered at the surface of the printed material due to micro-roughness stemming from the printing process (J. Stürmer, private communication).

With the magnification factor above, the light is fed into the fiber at roughly $F/5.4$; according to measurements of our fiber, losses in the spectrograph due to focal ratio degradation (FRD) are expected to be on the order of 5%.

2.2.3 The calibration unit

The calibration unit houses several calibration lamps, whose spectra are needed in the processing of the acquired science spectra:

- a thorium-argon (ThAr) hollow-cathode lamp (HCL), which produces a rich spectrum of emission lines used for wavelength calibration;
- a Thorlabs QTH10/M quartz-tungsten halogen lamp for flatfields/order tracing & extraction;

¹https://www.teleskop-express.de/shop/product_info.php/info/p6984_Starlight-Xpress-Lodestar-X2-Autoguider-und-CCD-Kamera.html.

TABLE 2.1: Properties of the Waltz Spectrograph

Wavelength coverage	4500 – 7000 Å
Collimator	Parabolic mirror, Ag-coating, $f_{\text{coll}} = 900$ mm, $D = 300$ mm
Echelle grating	<i>Richardson Gratings</i> , replicated from Newport Master MR152, area: 100×200 mm, 31.6 grooves/mm, blaze angle: 63.9°
Cross-disperser	Prism, 60° apex angle, height: 130 mm, length: 185 mm, material: Schott F2, surface quality: $\sim \lambda/4$, SolGel anti-reflection coating
Camera	Takahashi FSQ-106ED (Petzval quadruplet apochromat), $f_{\text{cam}} = 530$ mm, $D_{\text{cam}} = 106$ mm
Detector	Andor iKon-L 936, 2048×2048 , pixel size: 13.5×13.5 μm
Resolving power R	60,500

- a white-light LED (Thorlabs MCWHL2), also for flatfields/order tracing & extraction.

The continuum spectra of the halogen lamp and the LED peak at different wavelengths; therefore, depending on the required wavelength range in the science spectra, either one of the lamps can be chosen as continuum source, thus optimizing the S/N — or spectra of both can be acquired and combined. For RV measurements using the iodine cell, we are interested in the wavelength range of ~ 5000 to 6000 Å, and the LED suffices as continuum lamp.

The light of the lamps is coupled directly into the input fibers of a fiberswitcher, a LEONI 2×4 optical switch, which can be controlled from the main computer. From the output of the fiberswitcher, the calibration fiber, a circular multi-mode fiber with a width of 200 μm , injects the light of the calibration lamps into the frontend.

2.2.4 The spectrograph

The Waltz Spectrograph is a fiber-fed Echelle spectrograph in white-pupil configuration, using a prism as cross-disperser and working in quasi-Littrow mode (see Figure 2.2). It covers a wavelength range of roughly $4500 - 7000$ Å, and its efficiency at a wavelength of 633 μm was measured to be nearly 50% from fiber input to CCD focal plane (Heeren, 2016; Tala et al., 2016). The spectrograph is fed by a fiber with rectangular cross-section of size 33×132 μm ; the fiber output thus acts as an evenly illuminated entrance slit to the spectrograph, allowing to abstain from using an image slicer inside the instrument and therefore simplifying the reduction of the spectra.

The spectrograph fiber, injecting the light coming from the telescope, is connected to the FN-system, which refocuses the beam into the focus of the collimator with a magnification of $M_{\text{FN}} = 1.84$; it was designed to accept a $\sim F/5.5$ output beam of the fiber and transform it to the working $F/10$ of the spectrograph. With the focal length of the collimator of $f_{\text{coll}} = 900$ mm, the diameter of the collimated beam then is roughly 90 mm, thus perfectly filling out the free optical area of the Echelle grating. However, due to a fiber input of roughly $F/5.4$ in the frontend FFU as well as slight FRD of the fiber, the output beam is a bit faster than originally expected,




FIGURE 2.2: Model of the Waltz spectrograph, with the light beam indicated in turquoise.

leading to a slight overfilling of the Echelle grating. Light-absorbing aperture masks in front of the grating and the collimator prevent unwanted reflection of light on the aluminium mounts, limiting the effects of straylight.

Owing to the white-pupil design, the beam, reflected and diffracted from the Echelle grating, hits the collimator another time; the outgoing focus is offset by a few millimeters in horizontal direction from the focus of the FN-system by a rotation of the grating by $\sim 0.5^\circ$ with respect to the main optical axis of the instrument. Keeping this angle as small as possible minimizes optical distortions. A small transfer mirror, sitting right in the 2nd focus, directs the light back towards the collimator, from where the collimated beam travels through the cross-dispersing prism and is finally focused by the camera onto the detector. The camera, a Takahashi quadruplet achromat, has a focal length of $f_{\text{cam}} = 530 \text{ mm}$, leading to a total spectrograph magnification factor of $M_{\text{spec}} = 1.08$. The width of the imaged rectangular fiber output in the CCD focal plane therefore is $35.8 \mu\text{m}$, which delivers a theoretical resolving power of $R = 60,500$.

As detector we employ an Andor iKon-L 936 with a back-illuminated CCD and a pixel size of $13.5 \times 13.5 \mu\text{m}$, resulting in a mean spectral sampling of 2.6 pixel. The detector reaches a quantum efficiency of more than 90% over the relevant wavelength range of the I2 cell method, ~ 5000 to 7000 \AA . It uses a 4-stage peltier cooler to reach minimum temperatures of -70° C at an ambient temperature of 25° C when air-cooled. The detector has a low dark current of $4 \cdot 10^{-4} \text{ e}^-/\text{pix}/\text{s}$ at that temperature, and its read-out noise is 2.9 e^- at the slowest horizontal shift speed (HSS) of the pixel charges of 0.05 MHz.

The spectrograph has been designed and constructed with the aim of keeping light losses as low as possible. Both the transfer mirror and the collimator are therefore silver-coated, and all lenses carry anti-reflection coatings. For the determination of the photon-weighted mid-point of the observation, which is required for a precise barycentric velocity correction, an exposure meter needs to simultaneously monitor

the light flux. We use a Hamamatsu Corporations H9319 series photon multiplier tube (PMT), which is mounted underneath the Echelle grating and observes the 0th order of the spectrum; in this setup, no light is lost from the science spectrum.

2.2.5 Control software & electronics

From a central computer, all electronic components can be controlled, and observations are planned and carried out, using the "Waltzcontrol" software. The code is mostly written in Python 3 and features a Graphical User Interface (GUI). The observer can thus perform the following tasks:

- point the telescope, either via sky coordinates, direct movement commands in a specific direction, or through a build-in HIPPARCOS catalogue, and move the dome;
- stream the images from the guiding camera and focus the telescope;
- position the pick-up mirror and I2 cell in the frontend in and out of the telescope light path, and control the temperature of the I2 cell;
- control the calibration unit;
- control the CCD detector and PMT of the spectrograph, and save acquired spectra as FITS files, with all necessary information from the connected components in the FITS headers;
- display weather information, such as temperature and cloud coverage, which is recorded by a weather station on top of the LSW main building.

Electro-mechanical components, such as driving motors for the I2 cell and pick-up mirror in the frontend of the M2 focussing motor in the telescope, are controlled by Arduino microcontrollers², and communication with the central computer is established through Raspberry Pi computers³ linked in a local network. Currently, essentially all components required for observations can thus be controlled through the Waltzcontrol GUI from the central computer, with the exception of opening the dome slit, and removing the protection cover of the telescope tube. As soon as the dome slit is motorized, remote observations will thus be generally possible if the telescope cover is removed prior to the observing night.

The software has been build by former master student Dane Späth (Späth, 2019), who also incorporated most control modules, and it has been maintained and extended by several other master and bachelor students as well as myself. Namely, the dome control module has been added by former master student Kai Roth, who also incorporated an early version of a guiding algorithm in a Beta-version of the Waltzcontrol software (Roth, 2019). Tests of the guiding module however showed that the algorithm did not work as intended; in a deeper analysis performed together with Bachelor student Lars Wolfgramm, it was found that unaccounted backlash of the telescope, and move speeds that depend both on the direction of the movement and the pointing of the telescope are probable causes of the guiding problems. Therefore, observations currently need to be performed without guiding corrections, which decreases the overall efficiency of the system. An updated guiding algorithm is under development by L. Wolfgramm and, if successful, will be added to the Waltzcontrol GUI.

²<https://www.arduino.cc/>

³<https://www.raspberrypi.org/>

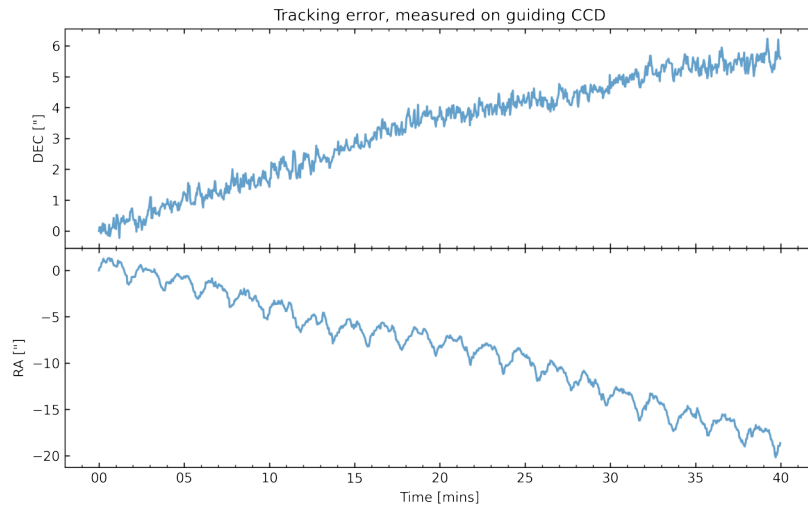


FIGURE 2.3: Measured tracking error of the Waltz telescope over a 40 min time period, both in DEC (top) and RA (bottom) direction.

Further updates of the software included a new control code for the CCD detector in the spectrograph, to allow the usage of a new version of the Andor software development kit (SDK) which enables a manual stopping of an ongoing exposure without losing the acquired data. The basic algorithms and required electrical components were implemented by myself, and the code was optimized for better timing precision by bachelor student Mauritz Wicker, who also adapted the control code for the calibration unit after the transition to a fiberswitcher (Wicker, 2020).

2.3 Performance of the system

2.3.1 Test of the tracking stability

As mentioned in Section 2.2.5, the guiding software of the telescope is currently not functional, and in observations the light throughput thus completely depends on the precision of the tracking of the telescope control system. While a solution for the guiding problem is being developed by bachelor student Lars Wolfgramm at the moment, the stability and precision of the tracking over time was tested to better understand the behaviour of the telescope control system and tailor the new guiding mechanism to our needs.

This test was performed together with L. Wolfgramm, who wrote the code to acquire the relevant data, and was performed as follows: A bright star was found and the telescope was pointed such that the star was positioned near the center of the guiding camera field-of-view in the Waltzcontrol GUI (but not too close to the pin-hole). Then, over a time of roughly 40 min, the position of the star in the guiding images was determined every few seconds, using a center-of-mass algorithm; during that time, no movement commands were sent to the telescope controller, so only the tracking motion influenced the results. Figure 2.3 displays the timeseries of the position changes of the star in DEC and RA directions relative to the start, where the originally measured pixel position changes were transformed to an angular change on sky. Clearly, over the long-term tracking errors become significant, especially in

RA direction: The angular size of the pinhole on sky is $3.7''$, so roughly 5 to 10 min after positioning the star on the pinhole and starting an exposure, no light from the stellar PSF would pass through the pinhole anymore. In addition to the long-term trends, a periodic variation in RA direction can be observed, with an amplitude of roughly $3''$ and period of 2 min. It was found that the period corresponds to the rotation period of the RA motor during tracking, and it can be assumed that it periodically de- and accelerates, for instance due to load changes caused by friction or non-symmetric gears.

The test clearly shows that in the current state, without a working guiding algorithm, observations with exposure times of several minutes will suffer from substantial light-losses. Even short exposures around roughly a minute will be affected due to the periodic tracking error. This underlines the importance of a working guiding algorithm to allow long observation times with good throughput. It is also possible that the application of encoders in the telescope control system will at least partly counteract the long-term errors observed in Figure 2.3.

2.3.2 Spectrograph detector characteristics

To test the specifications of the Andor iKon-L detector employed in the spectrograph, and the impact of the read-out noise and dark counts on the data presented in this thesis, I obtained a number of dark and bias exposures, using the same settings as for all other Waltz spectra: The CCD detector was cooled to a temperature of -30°C , the HSS set to 1 MHz, and the pre-amplifier gain to configuration $4\times$, which corresponds to a gain of $1.2\text{ e}^-/\text{ADU}$ according to the spec sheet. The dark frames were obtained with the shutter closed and an exposure time of 60 s, for the bias frames the exposure time was set to 0 s.

Figure 2.4 displays histograms of the accumulated counts per pixel for the masterdark and masterbias frames in the top and bottom panel, respectively, both of which have been median-averaged over five individual images. The bias frame shows quite uniform pixel counts of roughly 300 ADU, with a standard deviation of 2.7 ADU. With the gain of $1.2\text{ e}^-/\text{ADU}$, this can be converted to a read-out noise of 3.24 e^- . In a report by Andor, summarizing the results from a system test executed prior to delivery of the detector, the read-out noise was measured to be 6.5 e^- , and the base mean level was roughly 900 ADU, using the same gain and HSS settings but a detector temperature of -90°C . It is surprising that my measurement yields smaller numbers both for the base mean level and the read-out noise than the test conducted by Andor, and a further and more extensive characterization of the detector should be executed in the future to examine the characteristics also with other settings. In the analysis of spectra conducted in this thesis, however, I fell back on the values reported by Andor.

The median count level measured in the masterdark frame is $306 \pm 15\text{ ADU}$ over 60 s. By subtracting the base mean level and using the pre-amplifier gain, I can estimate the dark current to be roughly $0.1\text{ e}^-/\text{s}/\text{pix}$. This is a little larger than the Andor measurement of roughly $0.05\text{ e}^-/\text{s}/\text{pix}$ as stated in the test report; however, my dark count measurement has a relatively large standard deviation, and its median might be slightly over-estimated due to hot pixels. As no observation frame shown in this thesis used longer exposure times than 200 s, the dark current at the settings described above can definitely be neglected for these exposures. For longer

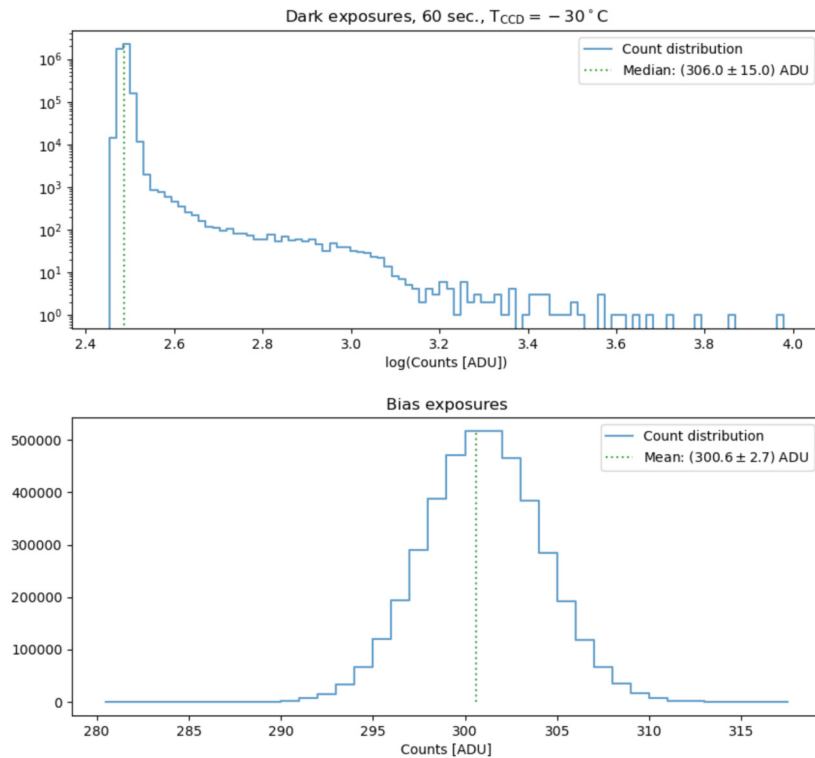


FIGURE 2.4: Histograms of average pixel flux values in 60s dark exposures (top), and bias frames with minimum exposure time (bottom), both at a CCD temperature of -30°C .

observations in the future, the detector can be cooled to lower temperatures, where the dark current drops below $10^{-4} \text{ e}^{-}/\text{s}/\text{pix}$ as tested by Andor.

2.3.3 Analysis from calibration spectra

First tests of the system were performed on calibration spectra, and allow to evaluate properties of the system such as resolving power and light throughput. Figure 2.5, top left, shows a sub-frame of a flat continuum spectrum using light from the LED, and the corresponding section of a ThAr spectrum on the bottom left. In both spectra, the advantage of abstaining from the use of an image slicer in the spectrograph itself becomes apparent, as the resulting lines are undisturbed images of the fiber output.

The ThAr spectrum allows to deduce the actual resolving power and sampling of the spectrograph from the widths of the individual lines. Figure 2.5, right, shows a histogram of the FWHM of all lines in the full-frame spectrum stronger than a certain threshold, measured by fitting 1D Gaussian functions along the pixel rows (i.e. in dispersion direction) to each line; the theoretical FWHM by design is indicated by the green dotted line. The median measured FWHM is 2.58 pixels, and the distribution shows a very distinct peak around that value. Using Equation 1.11 and the properties of the spectrograph, we can compute the effective resolving power for each line from the FWHM, and thus obtain a median resolving power of roughly 62,200 over the whole detector. The average real sampling, corresponding to the

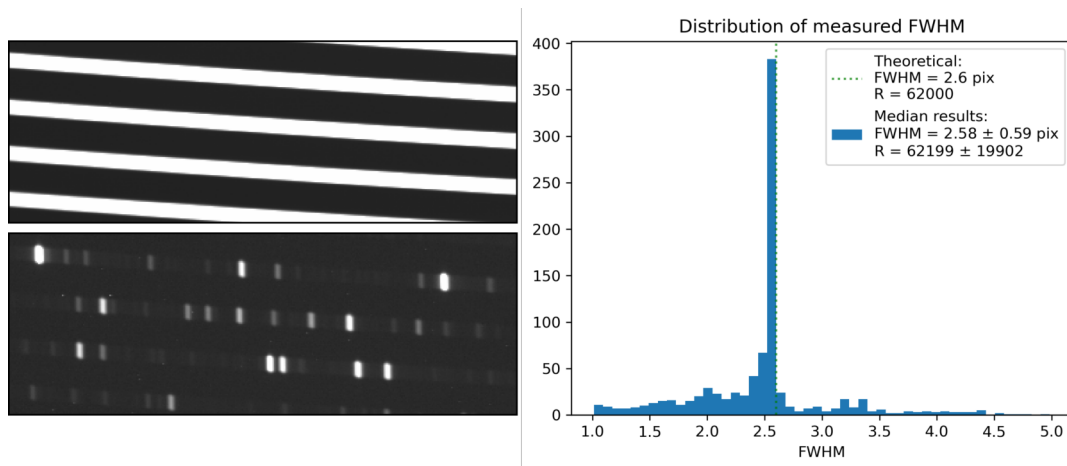


FIGURE 2.5: *Left*: Sub-frames of a LED flat spectrum (top) and a ThAr spectrum (bottom). *Right*: Distribution of the FWHM of lines found within a ThAr spectrum. The median FWHM and corresponding resolution are given, along with the theoretical values.

FWHM, and resolving power are thus a little smaller and higher than the design values, respectively. This might be explained by the left tail of lines in the FWHM distribution, which are probably CCD artefacts wrongly classified as lines, thus skewing the overall average to a smaller value. These artefacts also strongly increase the estimated uncertainty of the FWHM and resolving power. Still we can conclude that the resolution and sampling of the spectrograph meet the design requirements.

To test the stability, or lack thereof, of the spectrograph, a series of ThAr spectra was reduced and the relative shifts of the wavelength solutions deduced from the spectra were compared, using the capabilities of the Waltz DRS (see Chapter 3). The analysis showed short-term drifts of the spectral format of several 100 m s^{-1} over less than an hour, and a long-term drift on the order of 1000 m s^{-1} over few days. Additionally, the Echelle orders moved in cross-dispersion direction over the long term, with a shift of several pixels over the timespan of 14 d. This behaviour could be caused by temperature changes affecting opto-mechanical components, particularly the cross-dispersion prism, as the refractive index of glass, and thus its dispersion, is a function of temperature. Currently, the temperature of the spectrograph room is not controlled at all, and can thus undergo large variations, reflecting the changes of ambient temperature (with some time delay). Furthermore, the CCD detector is air-cooled, and the air from the detector fans is simply blown into the room.

We do not have any temperature sensors in the spectrograph room at the moment, but during the days at which the ThAr spectra were obtained in fact the ambient temperature on the Königstuhl underwent variations of roughly 20°C . To better understand and monitor the stability of the spectrograph in the future, I recommend equipping the spectrograph room and the optical bench with temperature (and pressure) sensors and save the data on the central computer. Furthermore, it might be worth investing into a temperature control system for the spectrograph room, to keep the temperature variations within a few centigrade, or at least to slow down the rate of temperature changes to achieve better short-term stability.

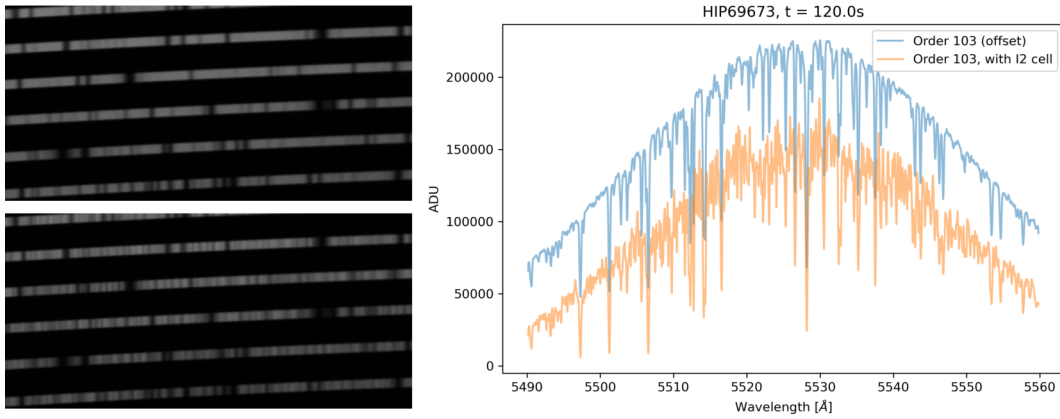


FIGURE 2.6: Observation spectra of the star HIP 69673 (Arcturus), obtained March 30th, 2021, with exposure times of 120 s. *Left*: Sub-frames of the raw spectra, without I2 cell (top) and with I2 cell (bottom). *Right*: One extracted order of the two spectra. The non-I2 spectrum has been offset in vertical direction.

2.3.4 Stellar spectra

Figure 2.6, left, shows subsets of two raw spectra of the star Arcturus (HIP 69673), obtained with an exposure time of 120 s with and without I2 cell, respectively. The I2 absorption lines imprinted onto the stellar spectrum is clearly visible. On the right of Figure 2.6, one order of the extracted spectra is shown; extraction was performed with the Waltz DRS, and the spectrum without I2 cell was offset in vertical direction to allow better distinction. The number of counts at the peak of the blaze function in the middle of the order is above 150,000 ADU, and the S/N as computed by the Waltz DRS is higher than 400.

As the guiding system of the telescope is not working yet, the light throughput of the FFU is still compromised by tracking errors, leading to the star not being centered on the pinhole at all times. Due to a known periodic tracking error, which leads to changes in the position of the stellar PSF on the pinhole mirror with an amplitude of roughly 100 μm and a period of 120 s, even the short observations shown in Figure 2.6 are affected by light losses.

To better understand the current throughput of the Waltz telescope without guiding corrections, I compared an observation of the star Pollux (HIP 37826) obtained with the Waltz spectrograph to archival Lick spectra of the same star. In a Waltz spectrum of that star, without I2 cell in the light path and using an exposure time of 200 s, the peak S/N as computed by the Waltz DRS was around 500 in the extracted spectrum. The Lick observation which is used in the stellar template creation for the I2 method was obtained with an exposure time of 90 s, and the peak flux reached in the extracted spectrum is roughly 32,000 ADU. When adopting a pre-amplifier gain of $2.8 \text{ e}^-/\text{ADU}$ and a read-out noise of 10.1 e^{-1} for the detector specifications⁴, and using Equation 3.9, this corresponds to a S/N of roughly 300, which is considerably less than in the Waltz spectrum of that star. Assuming a linear behaviour of the acquired flux with time and interpolating the result of the Waltz spectrum of Pollux, the same S/N could be reached with the Waltz telescope after 70 to 80 s, so a little faster than with the CAT used in the Lick survey of giant stars.

⁴Dewar #6, under: <http://mthamilton.ucolick.org/techdocs/detectors/>




FIGURE 2.7: Modified detector mount of the Waltz spectrograph.

This is not completely surprising, as the Waltz telescope's primary mirror is larger than the one of the CAT at Lick, and the modern Andor iKon-L detector of the Waltz spectrograph has a higher quantum efficiency than the detectors used with the Hamilton spectrograph. Specifically, the S/N achieved for that Pollux spectrum proves that the FFU is well-aligned on sky and allows to reach good efficiencies. However, with the guiding algorithm of the Waltz telescope not working yet, this performance can only be achieved in short exposures, and it requires a precise positioning of the stellar PSF on the pinhole at the start of the observation. Therefore, other stellar spectra obtained with the Waltz telescope in the same night as the Pollux spectrum described above show worse results. For instance, a second spectrum of Pollux obtained right after the first one and with the same exposure time just achieves a peak S/N of roughly 240. While for this spectrum the I2 cell was positioned in the light path, the extreme drop in the received maximum flux can only partially be explained by the absorption through the I2 gas as the intensity peak of Pollux lies outside the I2 absorption region. Furthermore, spectra of other stars showed strong variations in the received flux independently of whether the I2 cell was inserted or not.

2.4 Modifications of the system

2.4.1 Spectrograph detector mount

The spectrograph is mounted on a rectangular optical bench of size 1200×600 mm; in the design of the spectrograph, it became obvious that the Takahashi camera would extend over the edge of the bench, and the detector therefore had to be connected directly to the back of the camera, using a mechanical mount that allows precise alignment. The camera is principally able to carry the detector weight, and to ensure safety and stability, it sits in a solid stand which counteracts the leverage force introduced by the detector. However, calibration spectra showed a small upwards shift of the spectral format on the CCD after some time, suggesting a slight sinking of the detector. While this shift was only of the order of several pixels, corresponding to several dozen micrometers, any time-dependent mechanical variation

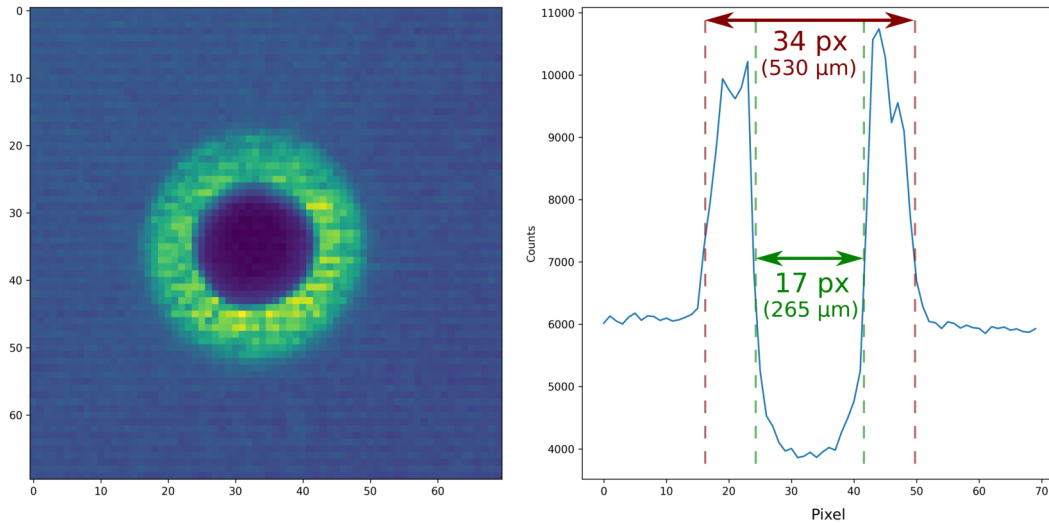


FIGURE 2.8: *Left*: Sub-frame of an image of the guiding camera, showing the pinhole and the focus of the calibration fiber, with light from the LED. *Right*: A horizontal cut through the sub-frame, with the imaged width of the pinhole (green) and of the calibration light (red) indicated.

affects the stability and thus the precision of the instrument; furthermore, it was possible that over longer times the camera would be impaired at least optically, maybe even mechanically.

I therefore designed an upgrade to the existing camera and detector stand, which is shown in Figure 2.7: Two aluminium legs are attached to the camera stand and extend further out underneath the detector. There, a horizontal aluminium plate connects the two legs, and a large, fine-threaded screw presses from the aluminium plate upwards against the detector mount. The screw has a ball-bearing head to spread the stress, and allow slight rotation of the detector, thus ensuring a correct orientation of the spectral format on the CCD chip. With the help of the fine-threaded screw, the detector can now be positioned in vertical direction. When the alignment is completed, the screw can be locked with a counter nut. This modification of the detector mount thus should reduce the stress exerted by the detector on the camera, and stabilize the imaged spectral format.

2.4.2 Calibration fiber

The calibration fiber guides the light from the calibration unit to the frontend, where the output beam is redirected by a mirror to take the same path as the light coming from the telescope. A lens transforms the beam to resemble the telescope focal ratio of $F/\# \approx 20$ and focus it onto the pinhole in the guiding unit. In a former setup, a multi-mode fiber with a core diameter of $300 \mu\text{m}$ was used as calibration fiber; this resulted in an image size of the fiber output of $\sim 800 \mu\text{m}$ on the pinhole mirror, thus strongly overfilling the pinhole and leading to a throughput of only around 10%. However, in this setup the calibration unit did not use a fiberswitcher to choose what light to send into the frontend; instead, the light of the lamps was fed directly into the calibration fiber using an optical system of lenses, a beam splitter and a movable mirror. Any throughput gains at the pinhole of the guiding unit by reducing the fiber core size would therefore have been counteracted by decreases in coupling efficiency in the calibration unit.

This changed with the upgrade to the fiberswitcher: The light of each lamp is now injected directly into the 200 μm -core size input fibers of the switch, and the output fiber, which has the same core size, is connected to the calibration fiber using a fiber connector. The core size of the calibration fiber could therefore be reduced to 200 μm , matching the fiberswitcher output, without significant light losses. The imaged fiber output size on the pinhole mirror then shrinks to $\sim 530 \mu\text{m}$, improving the throughput by a factor of ~ 2.5 to roughly 25% and thus increasing the overall light received in the spectrograph (see also Section 2.4.4). Figure 2.8 shows an image of the guiding camera on the left-hand side, with the focus of the calibration fiber centered on the pinhole; the right-hand side displays a horizontal cut through the center of that image, and the rough widths of the pinhole and the calibration light spot are indicated both in pixels of the guiding detector and actual size in the focal plane.

2.4.3 Design modification and assembly of the FFU

The FFU injects the light passing through the pinhole of the mirror in the guiding unit into the spectrograph fiber. In the original frontend design, a Bowen-Walraven type image slicer was positioned between the guiding unit and the FFU, transforming the circular PSF of the telescope to a pseudo-slit and thus optimizing the coupling efficiency into the rectangular fiber. This setup required two lenses, sitting in front of and behind the image slicer, and the total length of the optical path from the pinhole to the fiber tip was on the order of 10 cm. Early tests showed that this configuration was highly unstable: The frontend is exposed to strong thermal variations as well as vibrations and a constantly changing gravity vector caused by the telescope movement, and due to the long distance between pinhole and fiber and the many optical components involved a constant coupling of the light into the spectrograph fiber was impossible under these circumstances.

It was therefore decided to replace the macroscopic image slicer with a microscopic slicer which is 3D-printed directly onto the tip of the fiber. This modification allowed to greatly reduce the complexity of the FFU, by placing the fiber directly behind the pinhole, with only one lens in-between to couple the light into the fiber. An opto-mechanical design of the modified FFU was developed by Marcelo Tala (Tala Pinto, 2019), containing a doublet achromat lens (84-127, Edmund Optics) with a focal length of 3 mm and a diameter of 2 mm. The lens images the stellar PSF (or rather, the pinhole) with a magnification factor of $M = 0.264$ onto the fiber tip, leading to an input focal ratio of $F/\# = 5.38$. In the design, the lens is contained by an off-the-shelf lens adapter mount (Thorlabs LMRA2), sitting inside the fiber mount, which in turn is mounted in a XY precision translation stage (Thorlabs CP1XY). This allows precise alignment of the lens focus, and thus efficient coupling into the fiber.

However, in the final design review of the FFU the employment of the Thorlabs lens adapter mount proved to be unfeasible, mainly due to its outer diameter being too large. Furthermore, for an efficient coupling the relative distances between pinhole and lens, and between lens and fiber tip are quite sensitive. Whereas the former can be controlled by screwing the whole fiber mount (and thus also the contained lens) further into the XY-translation stage or farther out, the lens-fiber distance was fixed in the mechanical design. As that distance is only 2.7 mm and the input $F/\#$ quite fast, small imprecisions in the production of the mechanical mounts could greatly degrade the final efficiency.




FIGURE 2.9: Design of the modified FFU.

Before construction, I therefore modified the FFU draft again, taking these problems into account. Instead of the Thorlabs lens adapter mount, I designed a custom lens mount with a smaller outer diameter (Figure 2.9 bottom right, indicated in red). The lens was glued into that mount. Additionally, to allow control of the effective lens-fiber distance, I included a cylindrical spacer that is positioned between the lens mount and the back wall of the fiber mount (Figure 2.9 bottom right, indicated in blue). We produced a number of spacers with lengths between 2.7 and 3.6 mm in steps of 0.1 mm; thus, the lens-fiber distance and correspondingly the efficiency can be optimized by exchanging the spacer used in the FFU. Just as in the original design, the fiber mount is then screwed into the XY-translation stage, which is mounted behind the guiding unit.

I first tested the FFU in the laboratory, to determine which spacer delivers a magnification of the system closest to the optical design. To achieve this, I connected a fiber to the FFU and back-illuminated it with a LED; in the back-focus of the FFU, which will coincide with the pinhole in the frontend, a detector was positioned to record the image of the fiber. Thus, the width of the fiber image s_{im} could be measured, and with the actual fiber width s_{fi} known, the forward-magnification of the FFU is simply

$$M_{fw} = M_{bw}^{-1} = \frac{s_{fi}}{s_{im}}, \quad (2.1)$$

where M_{bw} is the magnification of the back-illuminated system. This was repeated for all spacers as to find the optimum magnification.

For the measurements, the spectrograph fiber with a rectangular cross-section was used to eliminate any measurement errors arising from differences in fiber dimensions. The top-left side of Figure 2.10 shows the imaged back-illuminated fiber, obtained with a spacer of length 3.3 mm in the FFU. The rectangular image was rotated such that it roughly aligned with the pixel rows of the detector. The length of the rectangle was then determined on a median-average of the middle ten columns of

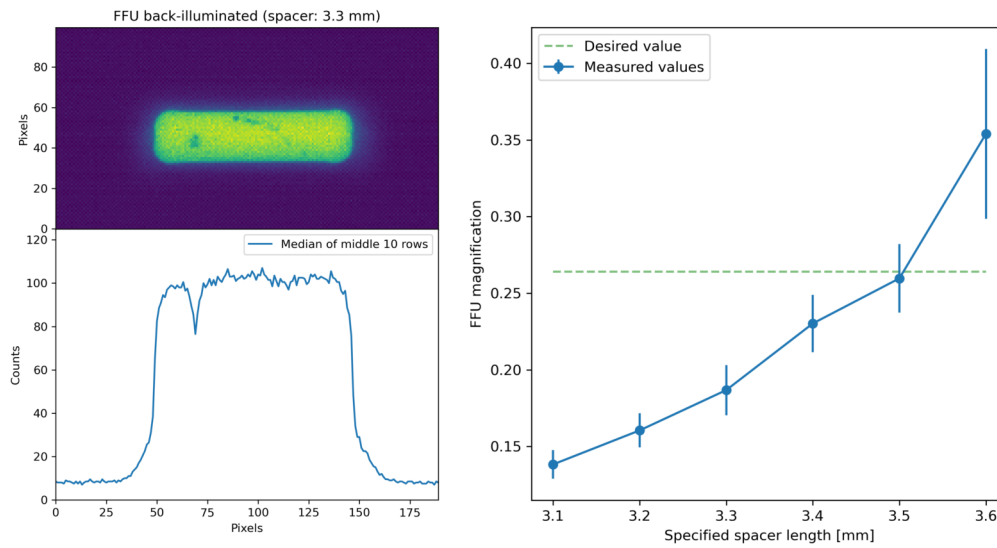


FIGURE 2.10: *Left*: Sub-frame of an image of the back-illuminated FFU in the laboratory, using the rectangular spectrograph fiber and a spacer of length 3.3 mm (top), and the median of the middle rows used for measuring the size (bottom). *Right*: Effective magnifications of the FFU for differing spacer lengths, measured on the back-illuminated images.

the image, using roughly half its maximum counts as mean measurement for s_{im} , and the lengths close to the peak counts and close to the background of the image as upper and lower bounds of the measurement, respectively (see Figure 2.10, bottom left). The actual length of the fiber cross-section is $s_{fi} = 132 \mu\text{m}$. Figure 2.10, right, displays the resulting forward-magnifications as measured with the different spacers, as well as the desired magnification from the optical design. The spacer with length 3.5 mm clearly delivers the closest result, and was therefore used during final integration of the FFU in the frontend.

2.4.4 FFU alignment and efficiency measurements

The FFU is the most critical element concerning the overall light efficiency of the whole telescope and spectrograph system, as even small errors in the alignment can lead to a complete loss of light in the spectrograph. In order to gain an understanding of the actual throughput of light in the system, I performed efficiency measurements at various points of the light path while inserting the FFU into the frontend. For the tests and the alignment of the FFU, light from the calibration LED, injected into the frontend via the calibration fiber, was used as source, and intensities were measured with a Thorlabs S120C photodiode⁵. The first measurement right in front of the pinhole on the nearly focussed beam coming from the calibration fiber served as reference. Then, before mounting the FFU, I measured the light intensity behind the pinhole to quantify the losses on the pinhole mirror.

Next, I inserted the FFU, connected the spectrograph fiber and measured the intensity on the fiber output; alignment of the FFU was achieved by varying its three degrees of freedom (X- and Y-translation by the positioner, Z-movement by screwing the fiber mount inside the positioner) until the intensity was maximized. Finally, I

⁵<https://www.thorlabs.com/thorproduct.cfm?partnumber=S120C>

TABLE 2.2: Cumulative efficiencies measured at various points throughout the FFU, with respect to the full light intensity falling on the pinhole mirror

	Calibration light [%] ¹	Star, 2'' seeing [%] ²	Star, 3'' seeing [%] ²
behind pinhole	25.5 ± 0.7	88.6 ± 3.4	63.2 ± 2.5
behind FFU unit	22.5 ± 0.7	78.2 ± 3.4	55.7 ± 2.5
behind fiber	11.4 ± 0.4	39.6 ± 2.8	28.2 ± 2.0

¹ Measured using LED light, and Thorlabs S120C photodiodes.

² Values based on theoretical throughput for given seeing at the pinhole, and measured calibration light efficiencies.

unplugged the fiber from the FFU and measured the intensity on its output, to learn about the internal losses on the lens. The measured intensities were divided by the reference value in front of the pinhole to get cumulative efficiencies, and uncertainties were estimated from several iterations of each measurement.

The first column of Table 2.2 gives an overview of the efficiency results. Clearly, the largest loss happens on the pinhole; however, the measured value agrees very well with the expected throughput from theoretical considerations of ~25%. Losses within the lens-system of the FFU are roughly 12%, which is a bit larger than expected. One possible explanation for the losses is that some adhesive material might have been pushed into the optically active area of the lens when it was glued into the lens mount, thus blocking parts of the light. In fact, under a microscope some smears of glue were visible on the backside of the lens, but the material is constrained close to the edges. Still, the possibility of it affecting the light beam cannot be completely excluded.

Finally, on the output of the connected spectrograph fiber a throughput of roughly 50% with respect to the intensity behind the FFU is measured, while the theoretical efficiency based on geometrical calculations of the circular PSF and the rectangular fiber cross-section should be around 60%. At least part of the discrepancy is due to coupling losses on the surface of the fiber tip (e.g. reflected light), which can easily be on the order of 10% even for well-aligned systems.

Overall, a little more than 11% of the calibration light falling onto the pinhole mirror reach the spectrograph in the end. Most of the loss is due to the strong overfilling of the pinhole, and for average and good seeing conditions the FFU will perform much better on sky. As the light gathered even from bright stars is too weak to be measured with the photodiodes, I approximated the on-sky efficiency of the FFU by computing theoretical throughputs at the pinhole for stellar PSFs of different seeing conditions, using Gaussian profiles with varying *FWHM*; the measured efficiency values from the calibration light were then referred to that new pinhole throughput. Figure 2.11 shows the curves for the theoretical pinhole throughput and the predicted overall efficiency of the FFU for a wide range of seeing conditions, and the throughputs for seeing widths of 2 and 3'' are given in Table 2.2. As can be seen, in the current setup we are losing more than 50% of the light even under best conditions, and towards larger-than-average seeing values the losses quickly increase to 80%.

For good conditions however, the coupling losses on the fiber dominate over the light lost at the pinhole, due to the non-matching geometry of the PSF and the fiber cross-section. This will change as soon as the microscopic image slicer is printed onto the fiber tip: As described in Section 2.2.2, in recent tests a prototype reached

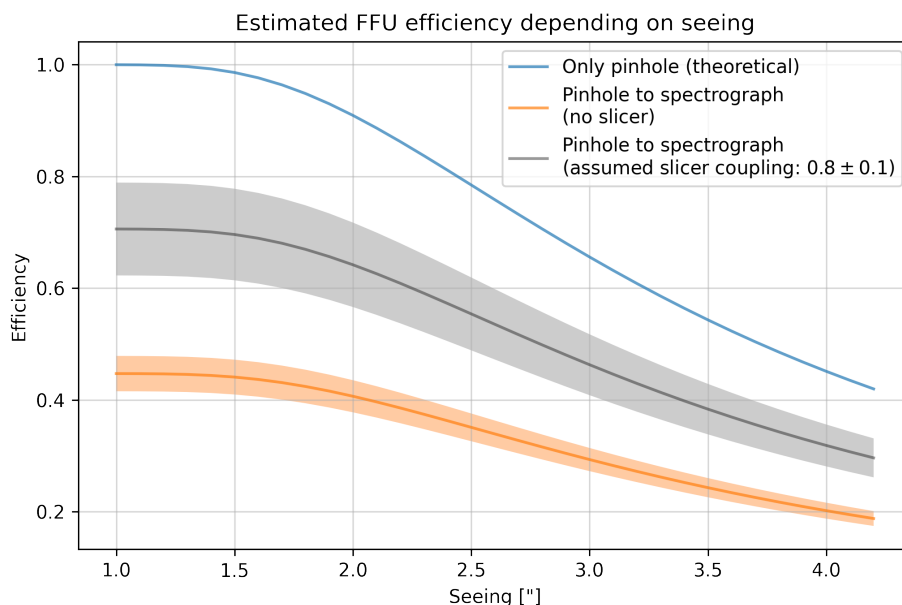


FIGURE 2.11: Estimated efficiencies of the FFU for different seeing conditions. The throughput at the pinhole was computed using a Gaussian-shaped PSF (blue curve), the FFU efficiency was then approximated using the measurement results from the calibration light (orange curve). In gray an efficiency estimate after the future upgrade to the printed image slicer is shown, assuming a slicer coupling efficiency of $(80 \pm 10)\%$.

a coupling efficiency of roughly 80%. Using this value, with an uncertainty of 10%, instead of the measured coupling efficiency from the calibration light tests and redoing the calculations for different seeing conditions results in the gray curve depicted in Figure 2.11. For average to good seeing conditions the FFU will most probably deliver overall efficiencies of 50 to 70% from the telescope focal plane into the spectrograph. This is roughly in line with conservative expectations from the start of the telescope project; with the spectrograph efficiency being nearly 50% at wavelengths around $600 \mu\text{m}$, and using a pessimistic guess of combined losses of 40% on the telescope mirrors and the I2 cell, the complete system efficiency should then be above 10% in the wavelength region of the I2 method even for bad seeing conditions. For a seeing of $2''$ and better, the system efficiency might even reach values of 20% and more.

2.5 Summary

After the I2 cell at Lick observatory was damaged in 2012, the Waltz telescope project was initiated to provide a new workhorse instrument and resume observations of stars with planet candidates in the survey of G-/K-giant stars. As of writing this dissertation, the building of the telescope is close to completion, and on-sky tests have started. Arguably, the most important step in the last few years was to enable a stable and efficient coupling of the stellar light coming from the telescope into the spectrograph fiber. In this chapter, I have presented my modifications to the fiber feeding unit (FFU), its assembly and first tests in the laboratory, and its alignment and measurements of the acquired efficiency at the telescope. For calibration light, the throughput from the telescope focal plane into the spectrograph is roughly

11%; the high losses are due to the strongly overfilled pinhole, a consequence of the calibration fiber width, but the reached efficiency is highly adequate thanks to the bright calibration lamps. Thorough on-sky throughput tests of stellar light could not be performed as the guiding algorithm is not functional yet, but the projected efficiency from the calibration light tests suggest values between roughly 28 and 40% for seeing conditions between 2 and 3'', respectively. Once the microscopic image slicer is printed on the spectrograph fiber tip, the throughput should increase by a factor of roughly 1.5.

Early stellar spectra acquired with the Waltz telescope and spectrograph show a large variance in S/N , due to displacements of the stellar PSF away from the pinhole over time, which are mostly caused by tracking errors. However, using an exemplary spectrum of Pollux, for which a high throughput was achieved, I could show that under good conditions the Waltz telescope will perform at least as well as the CAT telescope at Lick observatory concerning the S/N reached within a given time. Furthermore, the resolving power and sampling required for precise RV measurements are met by the spectrograph. It is therefore of highest priority to quickly solve the issues that occurred with the guiding algorithm, after which all important components will allow to start carrying out science observations. Shortly before finishing this dissertation, first tests of an updated guiding algorithm performed by L. Wolfgramm showed a significant improvement to earlier results, as the stellar PSF was held centered on the pinhole within few pixels on the guiding camera over a time of 10 min.

As mentioned in Section 2.3.3, further future upgrades should include the implementation of temperature and pressure sensors in the spectrograph room, so as to track environmental changes and reach a better understanding of the instrumental stability of the spectrograph. The data thus retrieved could also be used to automatically switch on a dehumidifier once the ambient temperature comes close to the dew point, in order to prevent condensation of water vapor in the air on the optics of the spectrograph.

Ultimately, the Waltz telescope is planned to allow completely remote observations, and to reach this goal the following components are still required:

- installation of an all-sky camera on the roof of the main institute building, to monitor the sky and possible weather changes. The camera, a OMEA 2M Mono Camera from Alcor-Systems⁶, has already been purchased and tested, and a control software has been written by M. Wicker (Wicker, 2020);
- integration of the cable winch to open and close the dome slit automatically. The winch along with a battery pack and a wireless charging system through magnetic induction have been purchased, and required electronics and mechanics for mounting on the dome are currently being prepared;
- proper routing of cables inside the telescope room, so as to prevent any damages.

The results presented in this chapter however prove that science observations can begin even before these upgrades.

⁶http://www.alcor-system.com/new/AllSky/Omea_camera.html

A DATA REDUCTION SOFTWARE (DRS) FOR THE WALTZ TELESCOPE

3.1 Introduction

As presented in Chapter 2, the Waltz project was initiated with the aim of measuring precise radial velocities of giant stars, and in recent months early on-sky observations have started with promising results concerning stability of the light throughput and spectral performance of the instrument. Extracting the radial velocity (RV) information from Echelle spectra however is not an easy task and requires a complex analysis of the data: a) pre-processing of the spectra, tracing and extraction of the Echelle orders; b) wavelength calibration of the extracted spectra using reference spectra; c) computation of barycentric velocities at the time of each observation; d) analysis of the RV content in each observation.

In the case of the Waltz spectrograph, as the instrument is not stabilized, the last point (analysis of the RV content) involves the use of the iodine cell method, which is considerably more complicated than the RV analysis of spectra from stabilized spectrographs using the cross-correlation function (CCF) method. Additionally, for a RV survey such as will be performed with the Waltz telescope, a highly automated data reduction code (or pipeline) is desirable: First because of the large amount of data that needs to be analysed, making a manual extraction of every single spectrum too time consuming; and second to ensure consistency and reproducibility of the results over long time spans.

A large number of code packages specifically developed for the analysis of Echelle spectra do exist; some of them are private, especially those written as default data reduction pipelines for specific instruments such as the HARPS data reduction software (DRS), while others have been published as open-source packages, usually designed to be easily applicable to different instruments. The latter can be further categorized into "full" data reduction pipelines, which offer a complete analysis from the raw spectra to radial velocity output (e.g. CERES; Brahm, Jordán, and Espinoza, 2017), and "partial" code packages, which only perform certain parts of the analysis, such as the extraction of the raw spectra (e.g. REDUCE; Piskunov and Valenti, 2002), or the RV analysis of already extracted spectra (e.g. HARPS-TERRA, SERVAL or WOBBLE; Anglada-Escudé and Butler, 2012; Zechmeister et al., 2018; Bedell et al., 2019, respectively).

In this chapter, I present the development of a data reduction software (DRS) for the Waltz telescope project, which I refer to as Waltz DRS. It builds upon existing

code packages, but I implemented substantial changes and upgrades to these packages to merge them into the Waltz DRS. The original sources and an overview to the project is given in Section 3.2. Section 3.3 describes the spectrum reduction routines employed by the Waltz DRS and presents some early results on Waltz spectra. Section 3.4 then focuses on the methods for the extraction of radial velocities (RVs) and their code implementation, and in Section 3.5 I present first test results of the RV extraction code. Finally, in Section 3.6 I analyze the computational performance of the Waltz DRS, before I close with a discussion and outlook in Section 3.7.

3.2 General structure of the DRS

For the development of the DRS of the Waltz project, it was decided to make use of existing packages wherever possible in order to utilize the benefits of already successfully tested code. Python 3 was chosen as the programming language for the DRS because it has long-time support, is license-free, and has a virtually endless number of publicly available resources, including a large number of packages specifically for astronomy, modelling and statistics. Also some of the spectral and RV packages mentioned in the introduction are implemented in Python. Furthermore, the control software of the Waltz Telescope is written in Python 3.

The Waltz DRS is therefore based on three major Python packages (also summarized in Figure 3.1):

- For the reduction of the raw spectra, including pre-processing and wavelength calibration, the open-source Echelle spectra reduction package CERES (Brahm, Jordán, and Espinoza, 2017) delivers the foundation. As CERES is largely written in Python 2, for which support has been halted, I translated the relevant parts into Python 3. CERES originally also uses a strict functional programming approach; in my translated version, I included a number of classes to ease handling of the data and code.
- For the computation of precise barycentric velocities, the open-source package barycorrpy (Kanodia and Wright, 2018), which is a Python 3 implementation of the IDL code BaryCorr by Wright and Eastman (2014), is used.
- The RV analysis by means of the iodine cell method is based on the package pyiodine, developed by René Tronsgaard Rasmussen in 2015/16 for the SONG project (Rasmussen, 2016). It is written in Python 3 using a mainly object-oriented approach, and was meant as an alternative for the dedicated SONG pipeline iSONG (Antoci et al., 2013; Grundahl et al., 2017), but never completely finished and tested, and therefore never published. I extended its functionalities and adapted it to our needs.

The three individual packages offer all necessary functionalities to perform the many different steps required for extracting RVs from the raw spectra. The exact workflows for the analysis are defined in main routines, that import the required modules from the underlying packages. The main analysis routines are therefore comparably compact and easy to understand, thus allowing quick changes to be made to the general reduction steps.

Additionally, all important instrument-, machine- and user-specific parameters have been bundled in parameter input files, where they can be easily defined and changed using Python classes and dictionaries. These parameter files are imported from the

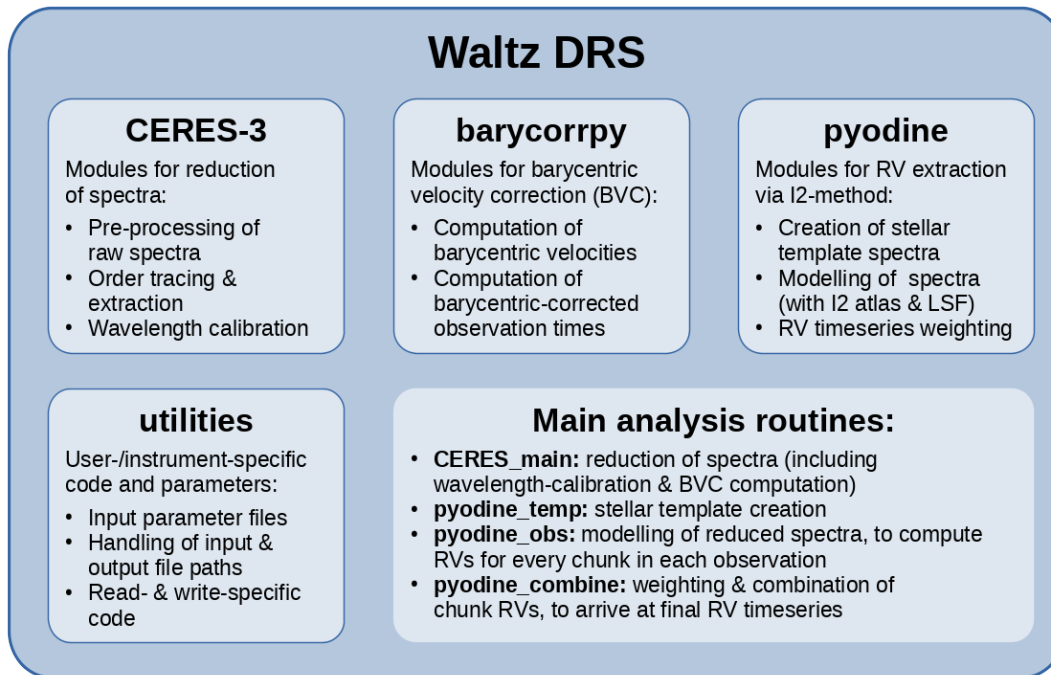


FIGURE 3.1: Structure of the Waltz DRS; bordered rectangles denote modules that are imported into the main analysis routines.

`utilities` module along with filehandling routines, that manage input and output file paths, and Python classes defining the reading, storing and writing of data that also depends on the machine or instrument used. Thus, when testing different parameter combinations or adapting the complete software to a different instrument, only the routines in the `utilities` module need to be changed.

All three (partially modified and extended) packages reside in a common Waltz DRS repository, along with the `utilities` module and the main analysis routines, strictly separated from the data. The repository is managed through the version-control software Git and hosted on the web-based DevOps lifecycle tool GitLab¹ to allow easy access and development from various computers.

3.3 Reduction of the spectra

The reduction routines of the Waltz DRS are based on the Collection of Elemental Routines for Echelle Spectra (CERES), which was published in 2016 under the free-software MIT license by Brahm, Jordán, and Espinoza (2017), and is built in a modular approach in order to be easily applicable to many different Echelle spectrographs. At the moment, the software’s GitHub repository² contains reduction pipelines for a total of 15 instruments. Namely, for the stabilized fiber-fed spectrographs CORALIE and FEROS the RV precision achieved by their respective CERES pipelines is on the order of 7 m s^{-1} , which is similar to the precision of the official CORALIE pipeline, and much better than the roughly 30 m s^{-1} precision achieved

¹https://gitlab.com/Heeren/waltzy_full

²<https://github.com/rabrahm/ceres>

by the dedicated FEROS pipeline. These CERES pipelines have been used to confirm transit detections of exoplanets in the HATSouth survey, and in different other studies (see Brahm, Jordán, and Espinoza, 2017).

Due to its high flexibility, demonstrated by the large number of CERES-based reduction pipelines developed for different instruments, and its good performance, CERES was a natural choice to build the Waltz DRS on. However, I made some structural changes to merge the CERES code into the Waltz DRS. In Section 3.3.1 I present an overview of these changes and of the general code structure. Section 3.3.2 then moves through the typical workflow for the reduction of spectra and describes the mathematical methods and code implementation step-by-step. In Section 3.3.3 I explain the implementation of the barycorrpy package in the reduction routines, and in Section 3.3.4 I finally present some test results from Waltz spectra.

3.3.1 Overview of the modified CERES code

As mentioned above, most parts of the CERES code are written in Python 2, with the exception of computationally heavy tasks such as order extraction and computation of cross-correlation functions (CCF), which are implemented in C and Fortran, respectively. I therefore had to translate the code into Python 3, ensuring correct wrapping of the C and Fortran routines. Furthermore, in its current form CERES only offers RV computation using the CCF method (see Section 1.1.3), which produces large systematic errors when used for spectra from non-stabilized spectrographs. Thus I only employ CERES for the reduction of the raw spectra, that is for the pre-processing of the images using bias frames, the tracing, extraction and normalization of the Echelle orders, and the wavelength calibration with ThAr spectra. Only CERES functions relevant for these tasks were translated into Python 3. As in the original package, all these functions sit in a common module called GLOBALutils, from where they can be imported. During translation of the code I preserved the internal logic of the functions as far as possible, in order to allow comparability with the original CERES code.

Additionally, I integrated some Python classes for storing relevant data internally, thus benefiting from the capabilities of Python objects. The basic classes are defined in the components module:

- A `RawSpectrum` class acts as basic container for spectra, both raw 2D image spectra and reduced spectra. A subclass `ScienceSpec` builds upon this class, adding a method for saving the data to FITS (or possibly also HDF5³) files.
- A `HDF5_Dict` class serves as a wrapper for Python dictionaries, providing capabilities to easily save and load data to and from HDF5 files. It's child class `Meta_data` adds more specific properties and methods and is used for storing trace and wavelength data.
- `Instrument` and `Star` classes store relevant data about the instrument used, such as geographical longitude and latitude, and the star whose spectra are being reduced, such as its coordinates.

Classes and functions that provide user- and instrument-specific capabilities are defined in the `utilities.load_CERES` module: The `WaltzSpectrum`, a subclass of `ScienceSpec`, introduces methods to load Waltz spectra from FITS files, along with

³<http://hdfgroup.org/>

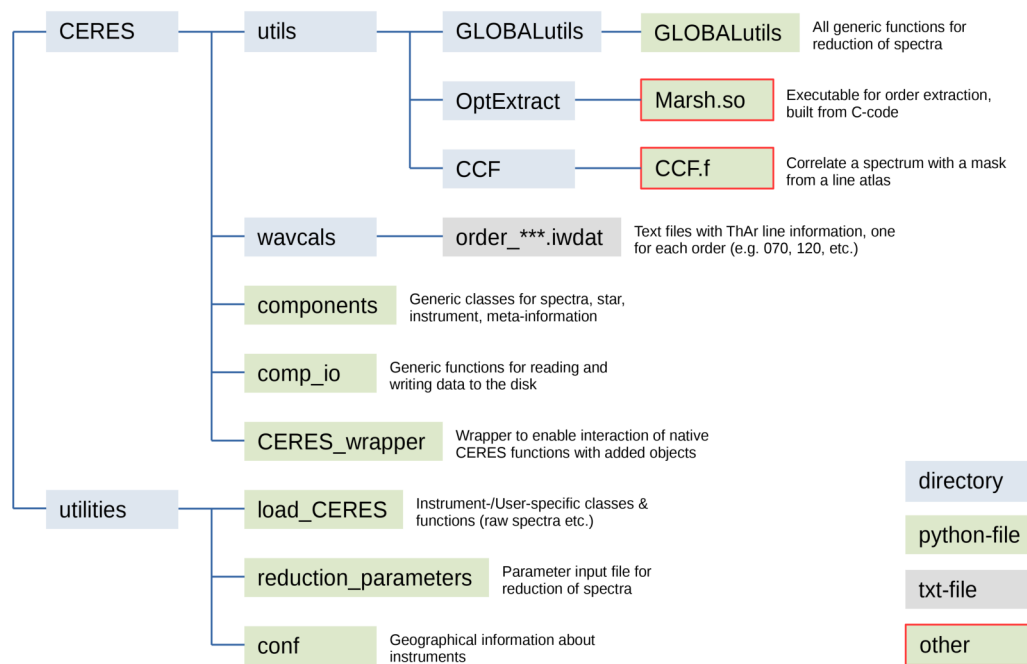


FIGURE 3.2: Organisation of the modified CERES package and its relevant modules within the Waltz DRS.

relevant information from the FITS headers (such as observation times, observed targets, etc.). Furthermore, a `NightLog` class is defined here, which is used in the reduction pipelines to classify the input spectra by type (i.e. science or calibration spectrum) and store meta-data that is relevant for the general reduction process. When changes happen to how the raw spectra are saved from the instrument, or the pipeline is adapted to another instrument, only the `utilities.load_CERES` module needs to be modified.

Due to the function-based approach of the original CERES code, the implementation of the Python classes required a slight restructuring, as data from Python objects need to be unpacked before being handed to the functions. I therefore included a `CERES_wrapper` module, which establishes communication between the high-level, object-based main routines and the low-level, function-based computation code.

Finally, to reduce spectra from an instrument not yet contained in the CERES repository, a ThAr line-atlas for that instrument is needed for the computation of wavelength solutions of the spectra; the atlas is contained in a `wavcals` directory and consists of txt-files, one for each order, that contain approximate pixel positions and laboratory wavelengths of all lines used in the wavelength solution. I created an atlas for the spectral format of the Waltz spectrograph by analyzing the line positions in several Waltz ThAr spectra; laboratory wavelengths for the lines were taken from atlases of other instruments within the original CERES package, and controlled by wavelength information from the Scientific Report booklet about the ESO Echelle Spectrograph (D’Odorico et al., 1984). The atlas covers physical order numbers 71 up to 124 of the Waltz spectrograph, which corresponds to wavelengths between roughly 4550 and 8050 Å.

Figure 3.2 gives an overview of the structure of the modified CERES package, including all important modules. The workflow for the reduction of spectra is defined by

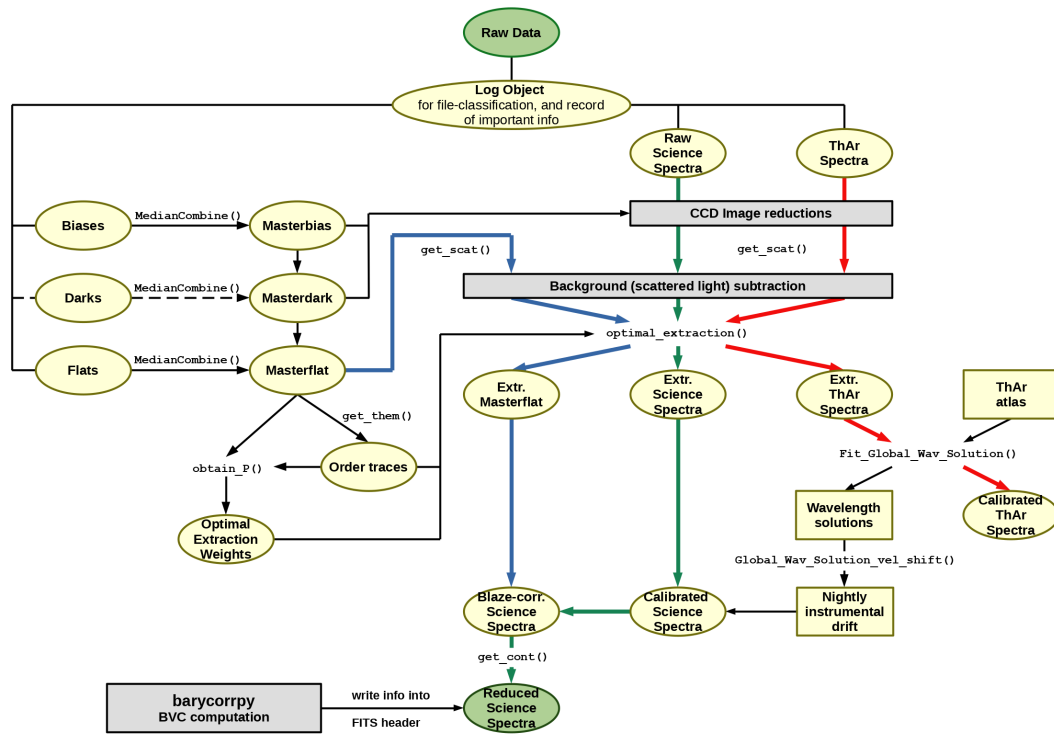


FIGURE 3.3: Flowchart of the reduction of science spectra, using the modules of the modified CERES package, and the package barycorrpy for the barycentric velocity correction.

the user in main routines that sit outside the CERES package and import all required modules. For the reduction of Waltz spectra I created the routine `CERES_waltz`. In the following section, I will briefly explain the different steps of the reduction process employed for the Waltz instrument, which generally follow the description given in Brahm, Jordán, and Espinoza (2017). Differences arise mainly due to my modifications of the package, and will be mentioned wherever they occur. A flowchart of the reduction process is given in Figure 3.3.

3.3.2 Workflow of the reduction process

Pre-processing using calibration frames

After the input images have been classified and their meta-data (including filenames and observation times) stored in the `NightLog` object, the basic image processing steps are performed: From the calibration spectra, master frames are created by median combination. First, all available bias images are combined to a masterbias frame, which is a measure of the count level created during read-out of the CCD array. Usually, in a second step masterdark frames are created from individual dark images, to correct for the count level produced by thermally freed electrons. However, the detector used in the Waltz spectrograph, the Andor iKon-L, has a dark current of only $0.0004 e^- / \text{pixel} / \text{s}$ at a temperature of -70°C ; dark counts for typical exposure times up to half an hour therefore remain below $1 e^- / \text{pixel}$, corresponding to less than $0.25 \text{ADU} / \text{pixel}$ at a gain of 4, and the use of dark frames is not required.

Finally, flatfield (flat) correction in spectrographs is performed differently than for standard imaging devices: As the pixel sensitivities depend on the wavelength of

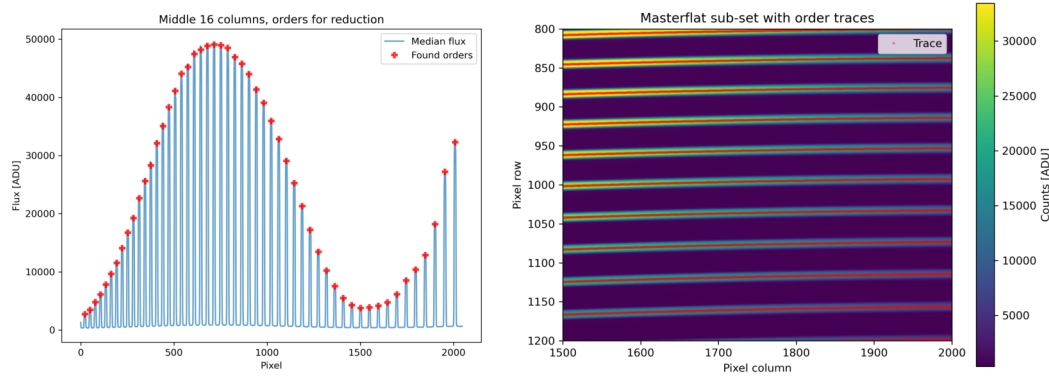


FIGURE 3.4: *Left*: Median flux of the middle pixel columns of the masterflat frame (blue), with the identified order positions (red crosses). *Right*: Sub-set of the masterflat frame, with the order trace polynomials (red dots).

the in-falling light, the light of the continuum lamp used for flatfielding also needs to be dispersed by the spectrograph. In slit spectrographs usually a longer slit is used for the flat spectra, which illuminates the orders uniformly over a larger distance in cross-dispersion direction than with the slit used for science observations. The science spectra can then be corrected by dividing through a masterflat spectrum created from the individual long-slit flat spectra. In a fiber-fed instrument such as the Waltz spectrograph however, where the input cannot be easily altered, flat spectra recorded through the same fiber as the science spectra show very low pixel values at the order edges and are therefore unsuitable for a conventional flat correction.

Nevertheless, we still require a number of flat spectra from a continuum source, that are median-combined to a masterflat frame (after subtraction of the master-bias frame). The masterflat is then used for the tracing of the orders, as well as for the calculation of pixel weights in the optimal extraction algorithm; here, pixels with extreme outlier sensitivities ("hot" or "dead" pixels) may be partially corrected for, as the algorithm estimates the weights from smooth polynomials in dispersion-direction. Finally, the masterflat frame is also needed for the estimation of the blaze function of the spectrograph (see Section 3.3.2).

Tracing the Echelle orders & scattered light subtraction

Before the extraction of the Echelle orders in a spectrum, the exact order positions on the CCD pixel array need to be determined. For this we use the masterflat frame, and through the `CERES_wrapper` we employ the CERES function `get_them()`, which works as follows: The orders are first identified by a cut along several central columns of the image (i.e. in cross-dispersion direction), which are median-combined along the dispersion direction in order to avoid the influence of artefacts caused by cosmic rays or the detector itself. After the reference cut is smoothed with a Gaussian kernel, all peaks are determined; an iterative routine then selects only those peaks above a certain threshold, defined by N times the standard deviation of the inter-order background, where N is defined in the input parameter file (see Figure 3.4, left).

These identified order positions serve as starting points for the tracing in dispersion direction: Moving along the rows of the image in steps of four pixels in both horizontal directions, the position of each order in each new column is fitted by a

Gaussian function. Each fit is performed on a subset of the column, centered around the last identified position of that order with a width corresponding to the height of the order in cross-dispersion direction (also defined in the parameter input file). Any outlier fit results are rejected. The positions are finally modelled with a high-order polynomial for each order, and the coefficients of the order traces are returned to the `CERES_wrapper`, which stores them in a `Meta_data` object for further usage in the extraction of the spectra. In a final step, the function `good_orders()` sorts out all orders that do not lie fully on the detector. Figure 3.4, right, displays a sub-set of a masterflat frame along with the evaluated polynomials which determine the order traces.

After the order traces are known, the scattered light background can be determined for each Echelle spectrum by using the CERES function `get_scatter()`: Median flux values for each inter-order region are computed for each column of the frame, and the scattered light frame is then created by a linear interpolation of these values within each column to the whole detector height, using capabilities of the Python package `scipy.interpolate`. The frame is finally smoothed by a 2D median filter and returned to the `CERES_wrapper`, which packs it into a `WaltzSpectrum` object and returns it to the main routine. The scattered light frames are then subtracted from the respective spectra used for their estimation.

The scattered light contamination stems from the Echelle grating, caused by inhomogeneities in its surface, from the roughness of optical surfaces, and it can also include light scattered from mechanical mounts if the light-path is intersected by the mounts at some point in the spectrograph. As the latter contamination however can be quite inhomogeneous, the scattered light routine `get_scatter()` might produce erroneous results if the scattered light from mechanical mounts dominates over the other sources. This would then greatly impair the resulting science spectra, and therefore it is important to reduce the light scattered from mounts as much as possible.

Optimally extracting the orders

CERES implements both a simple extraction and the optimal extraction algorithm in the modified version for curved spectra (Horne, 1986; Marsh, 1989), where the flux values of pixels i along a detector column at pixel position λ within an order are combined through

$$\hat{F}_\lambda = \sum_i W_{i,\lambda} F_{i,\lambda} \quad . \quad (3.1)$$

The weights $W_{i,\lambda}$ are calculated through an estimate of the fraction of the true flux of the object falling into each pixel i along the spatial direction, normalized by the total estimated flux:

$$P_{i,\lambda} = \frac{\hat{F}_{i,\lambda}}{\sum_i \hat{F}_{i,\lambda}} \quad \Rightarrow \quad \sum_i P_{i,\lambda} = 1 \quad . \quad (3.2)$$

The flux fractions $P_{i,\lambda}$ are modelled as low order polynomials running parallel to the traces of the spectra. Pixels with outlier flux values, caused by cosmics or instrumental defects, are ignored in an iterative process. By imposing the condition

that the flux estimate is unbiased, and minimizing its variance, the weights are then computed by

$$W_{i,\lambda} = \frac{P_{i,\lambda}/V_{i,\lambda}}{\sum_i P_{i,\lambda}^2/V_{i,\lambda}} \quad , \quad (3.3)$$

where $V_{i,\lambda}$ is the variance of the flux in pixel $\{i, \lambda\}$. By using the assumption that the S/N of the spectra is dominated by read-out noise (RON) and Poisson noise, the variance can be calculated by

$$V_{i,\lambda} = \frac{\hat{F}_\lambda P_{i,\lambda}}{G} + \frac{RON}{G^2} \quad , \quad (3.4)$$

where G is the detector gain.

For a correct computation of the flux fractions $P_{i,\lambda}$, the polynomials need to be spaced with subpixel offsets, and the flux values need to be evaluated on a finer pixel grid. Due to the interpolation of the flux values, the normalization of the polynomials and the iterative process required to ignore cosmics, the computation of the pixel weights is computationally expensive. Both the weights calculation as well as the extraction of spectra is performed in an executable module `Marsh.so` generated from C-code, which offers higher efficiency than Python code. The module is accessed by the `GlobalUtils` module. To further accelerate the computation, it can be performed in parallel processes, the number of which are set by the user and should not exceed the number of processor cores of the computer.

The following parameters need to be specified by the user in the parameter input file:

- `min_extract_col, max_extract_col`: first and last detector column to be extracted (0 and -1 correspond to the full detector width);
- `order_width`: the width of the orders in spatial direction;
- `NSigma_Marsh`: the rejection limit for flux outliers such as cosmics in the computation of extraction weights, in number of standard deviations;
- `NCosmic_Marsh`: the same as above, but for the optimal extraction;
- `S_Marsh`: the offset between the polynomials in pixels (needs to be between 0 and 1);
- `N_Marsh`: the degree of the polynomials.

In the main routine for the extraction of Waltz spectra, the weights are computed once from the masterflat spectrum, and then used in the extraction of all other spectra of the same night. The masterflat spectrum and all ThAr spectra are extracted before wavelength calibration. Extraction of the science spectra is performed after wavelength calibration.

Wavelength calibration of extracted spectra

After extraction of the ThAr spectra, a wavelength solution needs to be found, assigning each pixel of the reduced spectrum a corresponding wavelength. CERES

offers functions for the wavelength calibration both from ThAr and Fabry-Perot reference spectra. As we are using a ThAr lamp in the Waltz project, I only included the functions relevant for analysis of these spectra in the translated version.

As the Waltz spectrograph is not stabilized, the spectral format on the detector may vary by a few pixels in cross-dispersion direction over long time scales; thus, orders very close to the upper and lower edges of the CCD might at least partially migrate out of the frame and be rejected in the order tracing, which will then result in a different "physical order" $m_j = m_0$ being the first extracted "code order" $j = 0$. Physical order numbers m correspond here to the ones from the grating equation, while code order numbers j count through the extracted orders from 0 to $N_{\text{ords}} - 1$ ordered from the reddest to the bluest, where N_{ords} is the total number of extracted orders. A correct mapping of physical to code orders is essential in the wavelength calibration in order to load the proper ThAr line-atlas files. For each set of spectra of a given night, we therefore need to first find the conversion factor between the two, which by definition is simply m_0 . This is achieved in the following way: A number of ThAr orders in the range $[j^{\text{guess}} - l, j^{\text{guess}} + l]$ are cross-correlated with a binary mask, using the function `cor_thar()`. The mask is constructed from the ThAr line-atlas file of the physical order $m = m_0^{\text{guess}} + j^{\text{guess}}$, where m_0^{guess} is the expected conversion factor. The cross-correlation with the largest maximum defines which code order j corresponds to the used physical order m and thus returns m_0 . Additionally, the position of the maximum of that cross-correlation function corresponds to the rough pixel shift in dispersion direction Δ_p , caused by long-term instrumental drift.

The parameters j^{guess} , l and m_0^{guess} are defined in the parameter input file, and need to be determined by the user. For j^{guess} an order should be selected that is usually free of any blends and provides a dense forest of ThAr lines; for the Waltz spectra, order 113 (containing wavelengths $\sim 5007 - 5070 \text{ \AA}$) was chosen. Furthermore, a maximum pixel shift Δ_p^{max} is defined in the parameter input file and used in the cross-correlation of the orders as to constrain its range.

Next, the software creates initial wavelength solutions for each order, using the function `Initial_Wav_Calibration()`: The rough pixel position of ThAr lines within each order, x_i^{app} , is approximated by correcting the pixel positions defined in the line-atlas by the estimated pixel shift Δ_p from the last step. The exact pixel position of each emission line x_i is then computed by fitting Gaussians to small zones of the ThAr spectrum around x_i^{app} and using the mean of the fit. Thus, each line is assigned a pixel position x_i , an Echelle order j_i (or m_j), and a corresponding laboratory wavelength λ_i from the line-atlas; within each order, a polynomial is fitted in an iterative process between x_i and λ_i , where outliers are being rejected, to find an initial order-for-order wavelength solution. CERES offers the possibility of using either regular polynomials from the Python package `scipy.polyfit`, or Chebyshev polynomials of the first kind from the package `scipy.special.chebyt`, which are designed to be the solution to

$$(1 - x^2) \frac{d^2}{dx^2} T_n(x) - x \frac{d}{dx} T_n(x) + n^2 T_n(x) = 0 \quad , \quad (3.5)$$

where $T_n(x)$ are polynomials of degree n (e.g. $T_0(x) = 1$, $T_1(x) = x$, $T_2(x) = 2x^2 - 1$, etc.). These polynomials have some distinct properties, which make them robust when interpolating over a grid; in my implementation of CERES I use Chebyshev polynomials for the computation of all wavelength solutions.

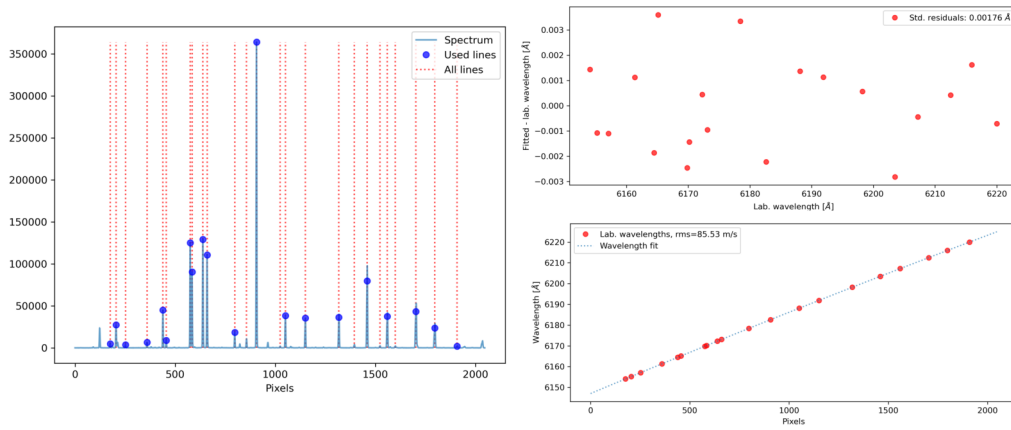


FIGURE 3.5: Analysis plots of the initial wavelength calibration for order 92 of a ThAr spectrum. *Left*: The spectrum of that order, with all lines from the corresponding ThAr atlas file (red dotted), and the lines used in the wavelength fit (blue dots). *Right top*: Residuals between the laboratory and fitted wavelengths for these lines. *Right bottom*: The fitted wavelength solution for that order (blue dotted) and laboratory wavelengths (red dots).

In the parameter input file, the minimum number of lines required for the solution of each order can be defined by the user, as well as the maximum root-mean-square (rms) of the residuals between the wavelength solution and the used lines in velocity space (i.e. in $[\text{m s}^{-1}]$). Iteration will continue until the minimum line limit or the rms limit of the line residuals is reached. Additionally, the degree of the polynomials n_x and several other parameters (e.g. the width of the window around each line for the Gaussian fit) are defined in the parameter input file. In my modification to CERES, the polynomial coefficients of the initial wavelength solutions to all orders and fit information for each used line, such as determined pixel position, wavelength, and residual to the polynomial solution, are returned by the CERES_wrapper as a Meta_data object.

Figure 3.5, left, shows an exemplary spectrum of physical order 92 of a ThAr frame; the pixel positions of all lines from the corresponding ThAr atlas file, corrected by the estimated pixel shift, are indicated by red dotted lines, and the lines used for the polynomial fit are marked with blue dots. On the right, the residuals between the laboratory wavelengths of used lines and the fit wavelengths are shown in the top, and the fitted wavelength solution along with the laboratory wavelengths in the bottom. The rms of the wavelength residuals within this order is 86 m s^{-1} .

Finally, the information from all lines found in the previous step is bundled to create a global wavelength solution: In the function `Fit_Global_Wav_Solution()`, an expansion of the grating equation 1.9 for small deviation angles is fitted to the line data using combined polynomials $P(x, m)$ (compare also Baranne et al., 1996):

$$\lambda(x, m) = \frac{d}{m} \cdot (\sin \alpha + \sin \beta) = \frac{1}{m} P(x, m) = \frac{1}{m} \sum_{i=0}^{n_m} \sum_{j=0}^{n_x} a_{ij} c^i(m) c^j(x) \quad , \quad (3.6)$$

where c^n denotes the Chebyshev polynomial of order n , a_{ij} are fitted polynomial coefficients, and x and m is the pixel position and Echelle order of each line. The polynomial degrees in cross-dispersion (order) and dispersion (pixel) direction, n_m

and n_x , respectively, are defined in the parameter input file and need to be determined by the user. In the main routines, I included plots that display the residuals between the lines and the global wavelength solution, thus allowing to judge the goodness of the fit and the impact of varying the degrees of the polynomials. The fitting employs an iterative process to reject outliers with residuals between the global wavelength solution and the individual lines larger than 3σ ; similarly to the initial wavelength solutions, the minimum total number of lines used and the maximum rms of the residuals are defined in the parameter input file and used as additional limits to stop the iteration. Polynomial coefficients and line data are again returned as a `Meta_data` object by the `CERES_wrapper`.

The above steps allow to compute wavelength solutions at the exact times of the ThAr exposures. For a highly stabilized spectrograph such as HARPS, the wavelength solution barely changes throughout one observing night, and using just one wavelength solution from a ThAr spectrum obtained before the start of the night might suffice to achieve m s^{-1} -precision of RVs for all observations. However, even for stabilized instruments wavelength calibration frames are usually obtained both before the start and after the end of the night, so that the instrumental drift throughout the science observations can be monitored. Unstabilized spectrographs might even require wavelength calibration exposures between all individual science observations so that any instrumental drifts on short timescales can be tracked and a high RV precision reached.

CERES offers to compute the velocity drift of a ThAr spectrum with respect to a reference wavelength solution using the function `Global_Wav_Solution_vel_shift()`: As input, it takes the polynomial coefficients of a reference wavelength solution a_{ij} , determined through `Fit_Global_Wav_Solution()`, and the line data of the ThAr spectrum to be analyzed, namely pixel positions x , order numbers m and wavelengths λ of all lines, as returned by `Initial_Wav_Calibration()` for that spectrum. Then, a modified version of Equation 3.6 is fitted to that line data, using the reference wavelength coefficients a_{ij} at fixed values, and varying only the relative instrumental velocity drift v_{inst} between the two spectra:

$$\lambda(x, m) = \frac{1}{m} \left(1 + \frac{v_{\text{inst}}}{c} \right) \sum_{i=0}^{n_m} \sum_{j=0}^{n_x} a_{ij} c^i(m) c^j(x) \quad . \quad (3.7)$$

The factor in front of the polynomials thus describes an instrumental Doppler drift of the whole spectrum, where c is the speed of light. As for the global wavelength solution, the fitting is done in an iterative process to reject outliers, using the same arguments as explained above. The CERES function models the velocity drift as a coefficient of the form $p = 10^6 \cdot v_{\text{inst}}/c$, and that coefficient is returned through the `CERES_wrapper` to the main routine.

In my main routine for the reduction of Waltz spectra, initial and global wavelength solutions are first determined for every single ThAr spectrum of a given observing night. Then, the first ThAr spectrum in time is selected as a reference, and the velocity drifts of all other spectra with respect to that reference are determined. The overall drift throughout the night is modelled using a linear interpolation from the Python package `scipy.interpolate`, and an analysis plot is created displaying the drift over time (see Section 3.10). The drifts of all spectra are finally saved, both to a HDF5 file using a `Meta_data` object as well as to a text file for easier access. If

only one ThAr spectrum is handed to the routine, the wavelength solution of that observation is extended over the whole night.

The interpolated drift allows shifted wavelength solutions to be created with respect to the reference wavelength solution at any time in the night, thus obtaining more precise wavelength solutions for the science observations. The overall precision obviously depends on the density of ThAr frames obtained. For the Waltz spectrograph, we do not aim to reach highest precision in this step, as the wavelength scale is mainly needed to serve as starting point for the later RV computation using the I2 cell method.

Creation of science output spectra

In my main reduction routine for the Waltz project, first all the steps described above are performed for the calibration spectra, that is: pre-processing of the ThAr and flat spectra; order tracing, scattered light subtraction, and order extraction of the masterflat, using pixel weights determined from the same spectrum; scattered light subtraction and order extraction of the ThAr spectra, using the trace information and pixel weights determined with the masterflat; and finally computation of the global wavelength solutions and the instrumental velocity drift, using the extracted ThAr spectra. Having done all this, we can proceed to the reduction of the science spectra.

My main routine offers two different output formats for the reduced science spectra, both of which are saved as FITS files: The first is oriented on the output of the original CERES routines and includes a wide variety of processed data, in a cube of size $[9 \times n_m \times n_x]$, where n_m and n_x are the number of extracted orders and number of pixels per order, respectively. The second output resembles the reduced spectra of the SONG project, with only the data necessary for the further I2 analysis in a data cube of size $[4 \times n_m \times n_x]$. The data arrays contained in each of the two outputs are summarized in Table 3.1. The output of the routine can be defined in the parameter input file by setting the `output_science` keyword either to "CERES" or to "SONG" (or "both" to create both output spectra). In the following, I will first describe the reduction steps as required for the more comprehensive "CERES" output, and later detail the differences in the "SONG" output.

Each raw science spectrum is loaded from file into a `WaltzSpectrum` object. The barycentric velocity correction (BVC) at the time of the observation v_{BVC} is computed, using the package `barycorrpy` (see Section 3.3.3). As timestamp the weighted midpoint of the observation is used, that is the time at which half of the summed photons in the exposure had arrived, which was determined using the exposure meter of the spectrograph. The output of the `barycorrpy` routine is written into the headers of the FITS output files in the end.

As described in the previous section, the instrumental velocity drift at the weighted midpoint time of observation with respect to the reference wavelength solution, $v_{\text{inst}}(t_{\text{obs}})$, is evaluated from the interpolated instrumental drift, and the wavelength solution is computed taking that drift into account (see Equation 3.7). Additionally, if required by the user, the barycentric velocity Doppler shift at the time of observation v_{BVC} may be used to correct the wavelength solution and shift it into the barycentric frame, by multiplying Equation 3.7 with a Doppler factor:

TABLE 3.1: FITS output formats of reduced science spectra in the Waltz main reduction routine

"CERES" data [$9 \times n_m \times n_x$]	"SONG" data [$4 \times n_m \times n_x$]
1. Wavelength solution	1. Simply extracted flux values
2. Optimally extracted flux values	2. Optimally extracted flux values
3. Errors of extracted flux	3. Continuum flux
4. De-blazed flux values	4. Wavelength solution
5. Errors of de-blazed flux	
6. Continuum-normalized flux values	
7. Errors of cont.-norm. flux	
8. Continuum flux	
9. Estimated S/N in continuum	

FITS header

- Original header of raw spectrum
- BVC, weighted midpoint JD, barycenter-corrected JD from barycorrpy
- Min. and max. extracted physical order number
- Estimated S/N at wavelength 5130 Å
- Information about the enclosed data frames

$$\lambda_{\text{bary}}(x, m) = \left(1 + \frac{v_{\text{BVC}}}{c}\right) \lambda(x, m) \quad . \quad (3.8)$$

In the case of Waltz spectra, which will be analyzed later using the I2 cell method, we stay in the uncorrected wavelength frame by setting the `bary_wave_scale` keyword in the parameter input file to "False". Furthermore, CERES offers functions to convert the wavelengths from air to vacuum, `ToVacuum()`, or vice versa, `ToAir()`, using the formalism published in Edlén (1953). Any of these options can be used in my implementation of CERES by setting the parameter file keyword `wavelength_type` to "air_to_vac" or "vac_to_air" (or "None" for no conversion). In the Waltz project, the wavelengths in the ThAr line atlas are given in conditions with air, and as the later I2 analysis will also use that wavelength scale, no conversion is applied. The wavelength solution is finally stored in the 1st array of the "CERES" output spectrum.

From the flux values of the raw spectrum, the scattered light in the spectrum is computed and subtracted, as described for the calibration frames. Then the orders of the spectrum are extracted, using the trace information and pixel weights determined with the masterflat. I use both the simple and the optimal extraction routine for the science spectra, thus the results can be compared and, if the optimal extraction fails, one can fall back on the simply extracted flux values. The "CERES" output spectra however only contain the optimally extracted flux at the moment. The optimal extraction algorithm also returns the inverse variance of the extracted flux values, which are contained in the "CERES" output as errors of the extracted flux.

Each order of the extracted spectrum shows a large-scale modulation due to the blaze function of the spectrograph. CERES corrects that effect with the help of the

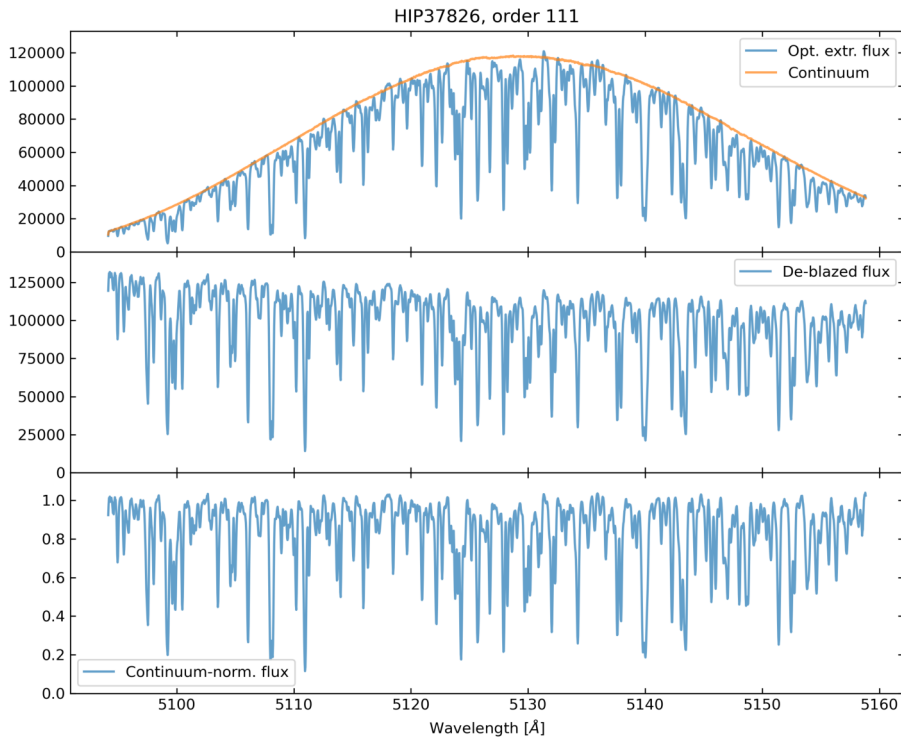


FIGURE 3.6: Reduction result of an observation of HIP 37826, order 111. *Top*: Optimally extracted flux, and fitted continuum. *Middle*: De-blazed spectrum, obtained by division of the extracted flux through the masterflat spectrum. *Bottom*: Continuum-normalized spectrum.

extracted and normalized masterflat spectrum, which is assumed to roughly represent the blaze function. The de-blazed science spectrum is therefore created by dividing each of its extracted orders by the corresponding order of the normalized masterflat. This operation also partly corrects for column-averaged variations of the pixel sensitivities, and thus is similar to a traditional flatfield correction. The de-blazed spectrum is stored into the 4th output array, and the respective errors, which are again the inverse variance of the de-blazed flux values, sit in the 5th array of the output data cube.

Additionally, CERES offers a function `get_cont()` to compute the continuum of each order of a spectrum: To the de-blazed flux values, a polynomial of low degree is fitted, using an iterative process to ignore wavelengths with absorption lines from the fit. The relevant parameters for the degree of the polynomial and limits for the outlier rejection are defined in the parameter input file. The polynomial coefficients are then returned to the main routine, and the continuum of the science observation can be computed by evaluating the polynomial coefficients and multiplying the orders with the respective normalization factors of the masterflat. The continuum flux is stored in the 8th array of the output. The 6th and 7th output array hold the continuum-normalized flux values, computed by dividing the de-blazed spectrum by the continuum, and the respective errors (again the inverse variance).

Finally, in the 9th frame an estimate of the S/N in the continuum for each pixel

is given. As the S/N depends on the overall flux registered by the detector, the "blazed" continuum needs to be used, which is simply the continuum F_C multiplied by the normalized masterflat $F_{\text{norm}}^{\text{flat}}$. The S/N is then the combination of the Poisson error of that blazed continuum, and the noise of the detector:

$$S/N \approx \frac{F_C \cdot F_{\text{norm}}^{\text{flat}}}{\sqrt{F_C \cdot F_{\text{norm}}^{\text{flat}}/G + \left(\frac{RON}{G}\right)^2}}, \quad (3.9)$$

where G and RON are the gain and the read-out noise of the detector, respectively.

All the data discussed above is stored in the data cube and written to the FITS file format, along with a header containing the original header of the raw science spectrum and additional information about the reduction process (see Table 3.1).

The "SONG" output in contrast contains four different frames in the data cube: 1st the flux from the simple extraction, 2nd the flux from the optimal extraction, 3rd the continuum flux, and 4th the wavelength solution. The data is obtained as described above for the corresponding frames, and the information written into the FITS headers is the same as for the "CERES" output.

3.3.3 Barycentric correction

The measured Doppler velocity shift of a stellar spectrum is a combination of the "true" RV of the target star and the so-called barycentric velocity, which is caused by several effects, most importantly the movement of the observer around the Solar System barycenter. Therefore, in order to extract the stellar RV with a given precision, the barycentric velocities at the times of observations need to be known with at least the same precision and accuracy. Wright and Eastman (2014) introduced algorithms to compute barycentric velocities (BVs) at the 1 cm s^{-1} -level, which is small enough to allow the detection of Earth analogues, and implemented them in the publically available IDL code BaryCorr. For the Waltz DRS, I make use of the Python package barycorrpy, which was published by Kanodia and Wright (2018) and builds on the same concepts developed by Wright and Eastman (2014).

The correction algorithm considers the following effects which influence the barycentric velocity:

- movement of the Earth's geocenter around the Solar System barycenter;
- rotation, precession, nutation and polar motion of the Earth, to calculate the position and velocity of the observer with respect to the geocenter;
- gravitational time dilation caused by objects inside the Solar System;
- leap second offsets;
- proper motion and systemic radial velocity of the star, and its parallax;
- Shapiro delay.

From all these effects, a barycentric redshift z_B is computed. Wright and Eastman (2014) show that the "true" RV of the target star v_* can then be derived by applying a barycentric correction (`bary()`) on the measured redshift of the stellar spectrum z_{meas} :

$$\frac{v_*}{c} = z_* = \text{bary}(z_{\text{meas}}) = (1 + z_{\text{meas}})(1 + z_B) - 1 \quad . \quad (3.10)$$

Due to a multiplicative cross-term between z_{meas} and z_B in that equation, which can grow as large as $\sim 3 \text{ m s}^{-1}$ for large BVs, the barycentric correction cannot be safely determined to a higher accuracy before knowing z_{meas} . Still, in the exoplanet community a simple sum of the barycentric and measured redshifts or velocities is often used to arrive at the RV of the target star:

$$\frac{v_*}{c} = z_* = z_{\text{meas}} + z_B \quad , \quad (3.11)$$

which suffices when working at RV precisions larger than roughly 3 m s^{-1} .

The barycentric correction algorithm in `barycorrpy` is implemented in a function `get_BC_vel()`, which requires as input arguments the Julian Date (JD), in UTC scale, at which to compute the velocity correction, and information about the observed target and the observatory. Target information can either be handed manually, with coordinates, proper motion, systemic RV and parallax of the target; alternatively, a HIPPARCOS identifier can be passed to the function, in which case the code extracts all required information from a built-in HIPPARCOS catalogue. For the observatory, geographical latitude and longitude as well as altitude are required. Finally, when supplied with a measured redshift z_{meas} , the function returns the true target star redshift as given by Equation 3.10. If no measured redshift is given, it uses $z_{\text{meas}} = 0$ and thus returns the BV as given by $v_B = z_B \times c$.

In the Waltz DRS, the BV at the weighted midpoint of the observation is determined during the reduction of the raw spectra in the modified CERES package and written into the output FITS header. To merge `barycorrpy` into the Waltz DRS, I wrote a wrapper module `bc_ceres` with a function `get_bcvel()`, which extracts the information required by `barycorrpy` from the `Star` and `Instrument` objects used in the modified CERES package. Usage of the built-in HIPPARCOS catalogue can be activated through a keyword. Furthermore, the wrapper accesses a function of `barycorrpy` to compute and return Barycentric Julian Dates (BJDs) of the observations (Eastman, Siverd, and Gaudi, 2010).

The `barycorrpy` package is generally well tested and is being employed in other data reduction pipelines (e.g. `SERVAL`, Zechmeister et al., 2018). Still, to test the correct integration into the Waltz DRS and the functionality of the `bc_ceres` wrapper module, I used it to analyze all observations from the Lick survey of G-/K-giant stars in our archive, and compared the resulting BVs to the values from the original Lick reduction routine. For that, I used the built-in HIPPARCOS catalogue of the `barycorrpy` package. Figure 3.7, left, displays the residuals between the `barycorrpy` BVs and the corresponding ones from Lick for each individual observation. The median difference falls at 0.1 m s^{-1} , and for most stars the residuals remain stable over time within the 0.01 m s^{-1} -level. For some stars however the differences are quite large and even evolve strongly over time. To better understand this behaviour, I correlated the median residual BVs of each star with several metrics, and found a strong correlation with the absolute value of the total proper motion of the star, displayed in the right panel in Figure 3.7. It is possible that the proper motion of the star was not considered in the Lick BV computation algorithm, or its contribution was computed differently; another explanation could be that the Lick algorithm might have




FIGURE 3.7: Test results of the BVs computed by barycorrpy on observations from the Lick survey. *Left*: Residuals of the original Lick BVs and the results from barycorrpy, for every single observation from the G-/K-giant star RV survey. *Right*: Correlation between the absolute value of the total proper motion of the stars in the Lick sample, and the absolute median of the residuals between Lick and barycorrpy BVs for observations of the same stars.

used different values stemming from an earlier HIPPARCOS reduction than the one incorporated in the barycorrpy package.

3.3.4 Test results of the spectrum reduction

The workflow described in the previous section is coded in the `CERES_waltz.py` main routine, which has been tested and optimized with calibration and stellar spectra from the Waltz instrument. As opto-mechanical work on the Waltz project has continued during the development and testing of the Waltz DRS, the parameters of the reduction routine had to be adapted again and again to account for changes of the spectral format. With the fiber-link between the telescope and spectrograph being established now, no further large-scale modifications are expected. The data and results shown in this section represent the latest status of the Waltz project and DRS.

All parameters that are used in the reduction of spectra presented here are bundled in the parameter input file `reduction_parameters.py` in the `utilities` module, and are presented in Table A.1 in Appendix A. To derive the numbers used there, I first started with parameter values as used in other instrument implementations of the original CERES package and then gradually altered them until reaching good results.

Simple versus optimal extraction

As described in Section 3.3.2, the CERES package can perform both simple and optimal extraction of the Echelle orders, and in the "SONG" output format both extracted spectra are stored. This allows us to assess whether the optimal extraction works as desired, by comparing the extracted flux values from both methods.

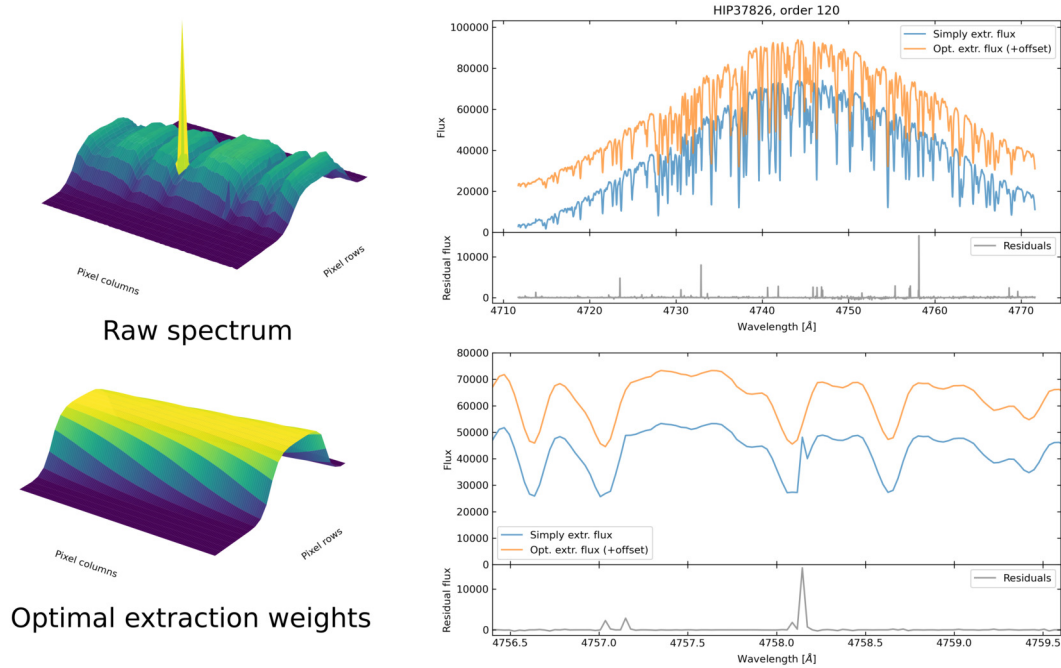


FIGURE 3.8: Reduction result of an observation of HIP 37826, order 111. *Top*: Optimally extracted flux, and fitted continuum. *Middle*: De-blazed spectrum, by division through the masterflat frame. *Bottom*: Continuum-normalized spectrum.

Figure 3.8, top left, shows a 3D plot of a subset of a raw stellar spectrum (HIP 37826, Pollux): A part of an Echelle order is visible, running nearly parallel to the detector rows from top left to bottom right, and absorption lines of the star can be seen. At the wing of one absorption line, a sharp spike caused by a cosmic event is recorded. In the bottom left of Figure 3.8, the optimal extraction weights for the same subset of the detector are shown, which have been computed on the masterflat frame of the same night. The tilt of that part of the order with respect to the detector rows becomes more obvious.

The right-hand side of Figure 3.8 shows the simply and optimally extracted flux values for that same order in blue and orange, respectively, over the whole order in the top panel, and a subset around the cosmic event in the bottom panel. The optimal flux values have been offset in vertical direction for clear distinction, and the residuals between the optimal and simple flux values are plotted in gray at the bottom of the plots. As expected, the cosmic event leads to a clear peak in the simple flux values at a wavelength around 4758 \AA in the right wing of an absorption line. In the optimally extracted spectrum in contrast, the absorption line is unaffected by the impaired pixel. The iterative routine for rejection of cosmics in the optimal extraction algorithm thus seems to work as expected.

Global wavelength solution

Figure 3.9 shows results from the global wavelength calibration, using a ThAr frame obtained with an exposure time of 4 s. The wavelength solution for all orders, given by the Chebyshev polynomials fitted to ThAr lines, is depicted by the gray lines in the top panel, and the blue dots are the pixel positions and wavelengths of all ThAr lines used in the final model. The bottom panel displays the residuals of the ThAr lines with respect to the wavelength solution. All used lines fall within $\pm 0.008 \text{ \AA}$ of




FIGURE 3.9: Result from the global wavelength calibration computed on a ThAr frame. *Top*: The wavelength solution for each order is depicted in gray, the ThAr lines used in the modelling are plotted in blue. *Bottom*: Residuals between the used lines and the wavelength model.

the wavelength model, and the rms of the residuals is 0.0028 \AA . In velocity space, this corresponds to an rms of roughly 140 m s^{-1} . The total number of lines used from all orders in that solution is 663. By treating them as N_{lines} independent measurements, the theoretical precision of the wavelength solution therefore is $140 \text{ m s}^{-1} / \sqrt{N_{\text{lines}}} = 5.4 \text{ m s}^{-1}$.

However, in reality this precision cannot be reached when applying the wavelengths to science frames, as the estimate above only reflects a momentary precision at the time of obtaining the ThAr frame; due to the instability of the spectrograph, even within a few minutes the spectral format might drift by several 10 to 100 m s^{-1} (see next Section). The actual wavelength precision that can be achieved is therefore strongly limited by instrumental effects.

As mentioned above, the total number of ThAr lines used in the determination of the global wavelength solution is 663, and in Figure 3.9 it can be noticed that the density of lines decreases towards redder wavelengths, that is lower order numbers. To a large part this can be explained simply by the fact that ThAr lamps have fewer appropriate emission lines at redder wavelengths, and hence the ThAr line atlas contains fewer entries in the lower orders. Furthermore, at wavelengths around 7000 \AA and longer some very bright Argon emission lines are present, which are heavily saturated when exposing long enough for the weaker lines to be clearly visible. Some of these Argon lines bleed out along detector columns, and they can interfere with emission lines in neighbouring orders.

Nevertheless, in the parameter input file the minimum number of ThAr lines that need to be used in the initial wavelength calibration of each order has been set to 10,

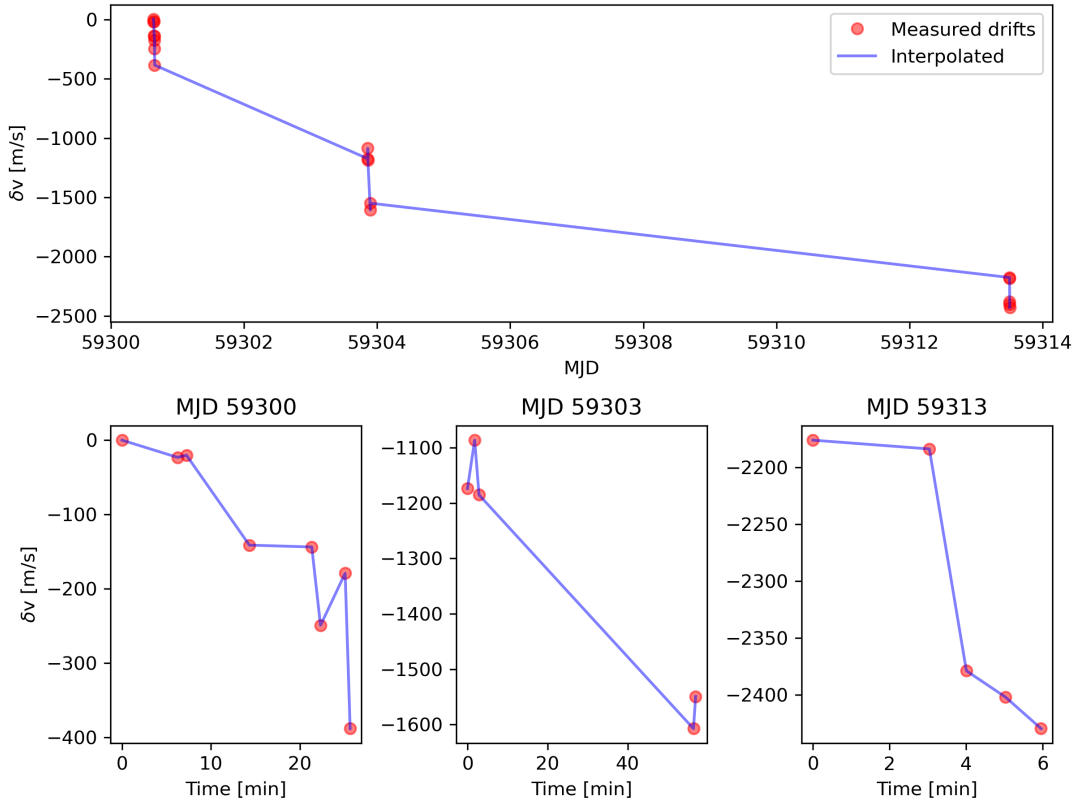


FIGURE 3.10: Measured velocity drifts of wavelength solutions of ThAr spectra, from three nights of observations (top) and within each night (bottom), relative to the first wavelength solution.

so at least that many lines are carried into the global wavelength calibration even by the red orders. While outlier lines are discarded then, leading to fewer than 10 used lines in the final wavelength solution for some orders, the fact that the scatter of the residuals in Figure 3.9 does not increase drastically at redder wavelengths shows that the overall solution satisfies our needs. Furthermore, in the spectral regime that is used in the later I2 analysis between roughly 5000 and 6000 Å, the density of used ThAr lines is much higher.

Velocity drifts of ThAr spectra

As described in Section 3.3.2, when supplied with multiple ThAr spectra, the reduction code computes the velocity drift of spectra throughout the night to interpolate the wavelength solution to the exact observation times of the science spectra and thus account for instrumental drifts of the wavelengths over time. The same CERES function used for the nightly velocity drift computation can be used to estimate the long-term drift when supplied with wavelength solutions of a number of ThAr spectra from several nights.

To test the long-term instrumental behaviour, I wrote an interactive analysis module, implemented as a Jupyter Notebook⁴, which allows wavelength solutions from selected ThAr frames to be loaded, and computes their velocity drifts relative to the first loaded solution in time. Figure 3.10 shows the results for a total of 18 ThAr observations, obtained at three different runs roughly four and ten days apart. Red

⁴<https://jupyter.org/>

dots denote the estimated velocity drifts at the times of observation, the blue line is the linearly interpolated drift. The top panel displays the long-term evolution, and the three bottom panels show the short-term drift within each run, where the time on the x-axis is relative to the first observation of that run.

It is obvious that significant instrumental drifts occur both over long and short timescales: Within the nearly 14 days between the first and the last observation, the spectral format drifted by roughly 2.5 km s^{-1} , which corresponds to a shift of more than one pixel in the dispersion direction. In fact, when comparing the raw ThAr spectra from which the first and the last wavelength solution were obtained, a shift of roughly that order of magnitude is visible by eye. Within each run, instrumental drifts of several 100 m s^{-1} occur over timespans of only a few minutes. The short- and long-term drifts might be explained by temperature changes inside the spectrograph room, affecting the opto-mechanical components (see Section 2.3.3).

3.4 Radial velocity analysis

From the reduced spectra, radial velocities need to be extracted by means of the iodine (I2) cell method, using the general approach presented in Butler et al. (1996). The Landessternwarte (LSW) Heidelberg actually has a copy of the Doppler code package "dop" used by the exoplanet group at Yale University, which is written in IDL, based on the original RV reduction code by Butler et al. (1996) (called Butler code hereafter) and has been greatly extended namely by Debra Fischer. As the Butler code was used for reduction of the RVs in the Lick giant stars sample and has proven to produce very trustworthy results, it was initially considered to adapt the dop code to the Waltz instrument. However, there are some drawbacks:

- the dop code is very complex, and includes a number of hard-coded parameters customized to specific instruments;
- the dop code produced worse results than the original Butler code and showed some systematics that were never completely solved;
- differentiating between the original parts in the code (the Butler version) and the later extensions is not always possible;
- the dop code is not written in Python, which hinders easy compatibility with the telescope software and other parts of the reduction code.

Therefore, I instead based the radial velocity analysis of the Waltz DRS on the pyodine package, which was kindly shared with me by its author R.T. Rasmussen. The pyodine code, developed as part of a master's thesis as a possible future RV reduction software for the SONG project, is completely written in Python 3 and built in a modular and object-oriented approach. It contains all necessary basic modules for performing a simple RV analysis of reduced spectra using the I2 cell method, and in first tests it seemed to match the short-term precision of the dedicated iSONG data reduction code (see Section 3.5). Development of the pyodine package was however halted before more rigorous and broader tests could be performed, and the code was still missing some important features (Rasmussen, 2016).

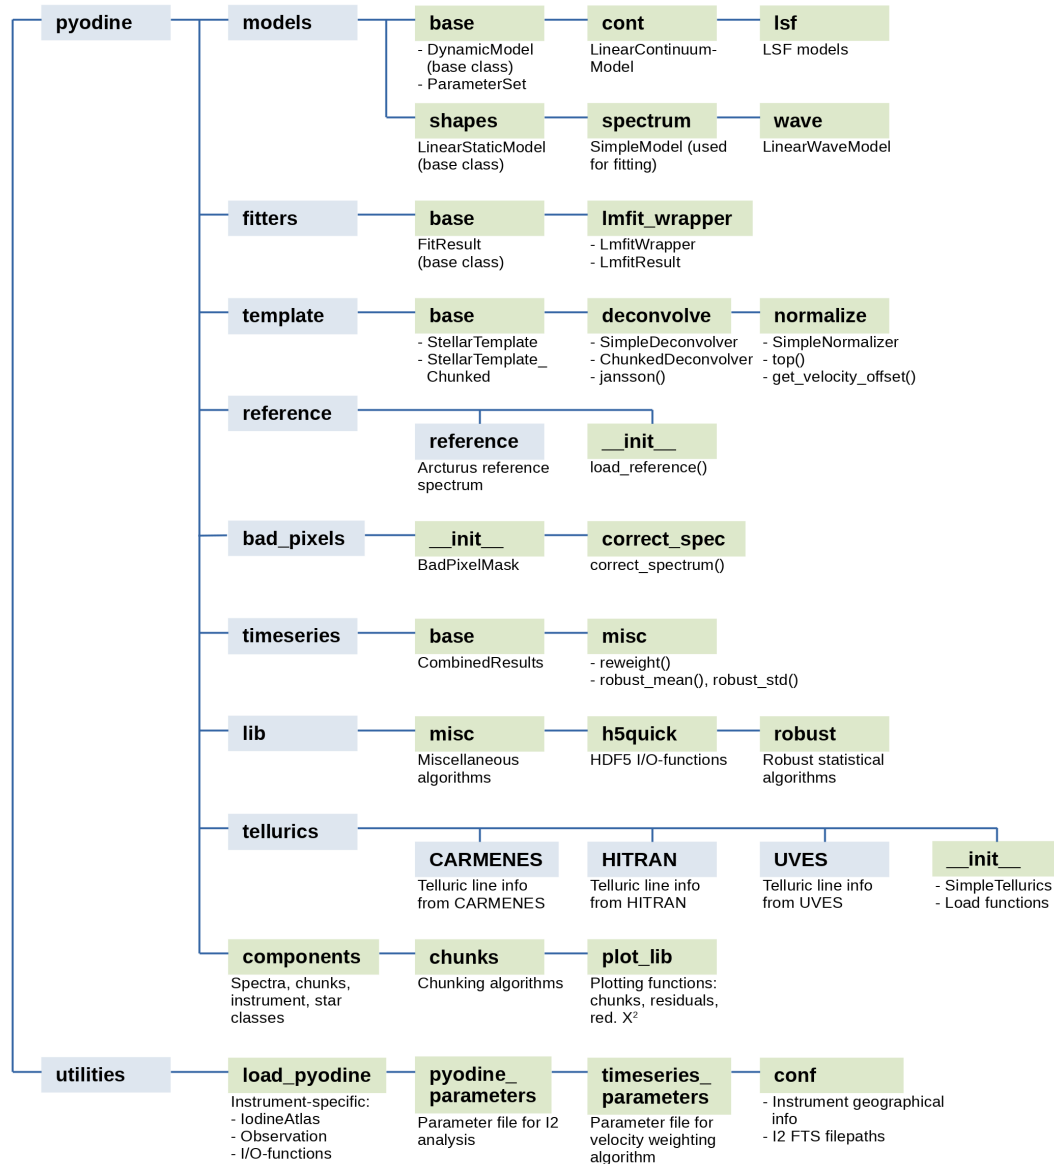


FIGURE 3.11: Structure of the modified pyodine package within the Waltz DRS.

3.4.1 Overview of the upgraded pyodine code

I extended the functionalities of the pyodine package and changed some of its existing modules, merged it into the overall Waltz DRS and adapted it to work with Lick and SONG data (for testing, see Section 3.5). Focusing on usability of the code, I followed the same philosophy as for the modified CERES package and bundled all relevant instrument-specific parameters in a parameter input file in the utilities module, thus allowing to easily test different parameter combinations or adapt the code to a new instrument. Additionally, a `load_pyodine` module again provides the instrument-specific read- and write-capabilities of spectra: Namely, a `WaltzObservation` class is used to load and store reduced science observations from FITS files, and an `IodineAtlas` class stores the I2 FTS (Fourier-Transform Spectrum) required in the modelling process.

Both of these high-level classes are sub-classes of more basic implementations in the components module, which contains general classes for data storing: A `Spectrum`

class is the most low-level container for spectra, accepting only 1D vectors of flux, wavelength and continuum values. A `MultiOrderSpectrum` class allows 2D spectra to be stored as a list of `Spectrum` objects, and its capabilities are extended with various properties by its sub-class `Observation` (which in turn is the parent-class for the instrument-specific `WaltzObservation`). Furthermore, defined within the `components` module are the `Chunk` class, which acts as basic representation for observation spectrum chunks, and the `TemplateChunk`, which is used to store stellar template chunks; both are sub-classes of `Spectrum`. All classes that store spectral information possess two methods that are important for the functionality of the code package: first `check_wavelength_range()`, used to check whether a range of wavelengths is contained in the stored spectrum, and second `get_wavelength_range()`, used to return a wavelength-defined spectral range from the spectrum.

Most of the other classes and functions of the `pyodine` package are contained within directories, which correspond to their respective application, in order to serve the overall clarity of the code package (see Figure 3.11). In the main routines that define the exact workflow to follow, the whole `pyodine` package or individual directories can be imported as code modules, thus allowing all contained functionalities to be accessible.

The following sections will present the mathematical descriptions employed in the I2 method and explain their implementation in the code package: First, in Section 3.4.2 I describe the formalism of the model and its different components, and the algorithms used to create chunks from an observation spectrum. Next, Section 3.4.3 deals with the fitting procedure, and in Section 3.4.4 I explain how the model parameters can be constrained. In Section 3.4.5 I introduce the usage and implementation of pixel weights, and Section 3.4.6 describes the template creation algorithms. In Section 3.4.7 I detail the formalism used to weight all chunk velocity results and combine them into a RV timeseries. Finally, Sections 3.4.8 and 3.4.9 present the workflows of the main analysis routines that are part of the Waltz DRS.

3.4.2 Description and implementation of the model

As described in Section 1.1.4, in the I2 cell method the radial velocities are extracted through a model of the observed spectra, taking the line-spread-function (LSF) of the instrument into account. Each spectral order is first split into a number of narrow chunks, each containing only $\sim 2 \text{ \AA}$ of the overall spectrum. Depending on the spectrograph, this typically corresponds to chunk sizes between 40 and 90 pixels. The spectrum of each chunk i , that is flux values $I_i^{\text{obs}}(x)$ over chunk pixels x , is then modelled by

$$\hat{I}_i(x) = k(x) \left[T_{I2}(\lambda(x)) \cdot I_s(\lambda(x)(1 + v_i/c)) \right] * L(x) \quad , \quad (3.12)$$

which incorporates the following sub-models:

- the LSF model $L(x)$;
- the continuum model $k(x)$;
- the wavelength model $\lambda(x)$;
- the part of the I2 FTS (Fourier-Transform Spectrum) corresponding to the modelled chunk, $T_{I2}(\lambda(x))$;

- the part of the stellar template spectrum corresponding to the modelled chunk, Doppler-shifted by the radial velocity of chunk i (v_i), $I_s(\lambda(x)(1 + v_i/c))$.

This model is evaluated on an oversampled pixel grid, in order to ensure correct convolution, and then re-binned to the pixels of the observation to be compared to the chunk flux. In my modified pyodine version, the oversampling factor is defined in the parameter input file.

In the following, I will outline the mathematical description and code implementation of the sub-models mentioned above, their combination and evaluation, and the algorithms used to chunk the observations.

LSF models

Three different LSF models are incorporated as static object classes in the original pyodine code. Each class includes a `eval()` method which, given a set of variable input parameters and an array of pixel values, returns the evaluated function result over these pixels. Additionally, a `guess_params()` method returns a first guess of the variable parameters, which at the moment are simply hard-coded values. The starting values for the variable parameters can however be altered prior to the fitting (see Section 3.4.3).

The `SingleGaussian` LSF is described by

$$L_{\text{Single}}(x) = \exp\left(-\frac{4 \ln(2)x^2}{FWHM^2}\right) , \quad (3.13)$$

with the only free parameter being the full-width at half-maximum $FWHM$. For most instruments, this model will not deliver a very good description of the actual spectrograph LSF, as asymmetries cannot be accounted for. However, it is useful in order to quickly determine an estimate of other free parameters (e.g. for the continuum and wavelength models), which can then be used as input parameters for a more complex modelling.

More complexity is allowed by the `SuperGaussian` LSF, consisting of a central Gaussian with variable exponent n , and one satellite Gaussian positioned on either side:

$$L_{\text{Super}}(x) = \exp\left(-\frac{|x|^\beta}{2c_0^2}\right) + \sum_{i=1}^2 a_i \exp\left(-\frac{(x - b_i)^2}{2c_i^2}\right) . \quad (3.14)$$

The positions of the satellites b_i and widths c_i are fixed; the four free parameters are the central exponent β , the central width c_0 , and the amplitudes of the satellites a_1 and a_2 . This model has been implemented in pyodine as it has been tested in the extraction of SONG RVs with iSONG.

Finally, the `MultiGaussian` LSF is built from a central Gaussian and five satellite Gaussians on either side:

$$L_{\text{Multi}}(x) = a_0 \exp\left(-\frac{x^2}{2c_0^2}\right) + \sum_{i=1}^{10} a_i \exp\left(-\frac{(x - b_i)^2}{2c_i^2}\right) . \quad (3.15)$$

Here, the positions b_i and widths c_i of all Gaussians as well as the central amplitude a_0 are fixed, leaving ten free parameters, namely the amplitudes of the satellite Gaussians $a_{i \neq 0}$.

All three models are normalized by the sum of their function values, $L_{\text{norm}}(x) = L(x) / \sum_x L(x)$. Furthermore, the two models that allow for asymmetries in the LSF, SuperGaussian and MultiGaussian, may be re-centered after a first evaluation: The offset is computed as the barycenter of the LSF, $\sum_x L_{\text{norm}}(x) \cdot x$, and the LSF is then evaluated and normalized again taking that offset into account. The re-centering is supposed to counteract possible interplay between the wavelength model and the LSF, where a highly asymmetrically positioned LSF could lead to a shift of the wavelengths and thus an incorrect velocity determination.

In addition to these LSF models incorporated in the original pyodine package, I included two more models that provide functionalities closer to the dop code. First, the MultiGaussian_Lick module, which is a direct translation of the dop-code LSF from IDL into Python, corresponds in its general arithmetic form to the previously defined MultiGaussian model; however, there are some small differences in the code implementation: First, the individual Gaussians in the MultiGaussian_Lick model are only evaluated over a limited pixel range of 10 times their respective Gaussian widths (c_i), centered around their respective positions (b_i); all three pyodine LSF models above in contrast are evaluated over the whole pixel range defined by the user (usually of the order of 16 to 20 pixels). The difference between these two implementations is insignificant however, as Gaussian functions quickly drop off to the sides: For a Gaussian with unity amplitude, the function value at 5 times the width is of the order of 10^{-6} . As can be seen in Figure 3.12, for a given set of parameters the un-centered pyodine and dop LSFs (MultiGaussian un-centered, MultiGaussian_Lick uncentered) are basically identical.

The more important difference between the pyodine and dop LSF implementations lies in the re-centering algorithm: While the MultiGaussian LSF is offset to the barycenter of the function, the MultiGaussian_Lick LSF is shifted so that the highest peak of the function falls into the center (at $x = 0$). For small LSF asymmetries these two approaches produce similar results, but when the LSF becomes very asymmetric the difference between the re-centering shifts can be of the order of several 10^{-1} pixel. The lower part of Figure 3.12 shows a very drastic example, where the LSFs are even re-centered in different directions.

The second LSF model that I added is the FixedLSF, which takes a pre-evaluated LSF array as an input argument instead of computing it from a functional form for a given set of parameters. The LSF needs to be supplied through a LSF_Array object, containing the LSF description in the form of a 1D-array for each chunk of the observation, along with the respective order and pixel indices identifying the chunks. While this setup allows the user to hand any fixed LSF description into the model of the observation, it has been incorporated specifically with the thought of using a smoothed LSF from a previous modelling run, as it is done in the dop code. An appropriate function has been added (`smooth_lsf()`), which produces a smoothed LSF for each chunk i , $L_{\text{smooth},i}(x)$, by median-averaging over neighbouring chunk LSFs $L_{k(i)}(x)$:

$$L_{\text{smooth},i}(x) = \frac{\text{med}(w_{k(i)} L_{k(i)}(x))}{\sum_x \text{med}(w_{k(i)} L_{k(i)}(x))} \quad , \quad (3.16)$$

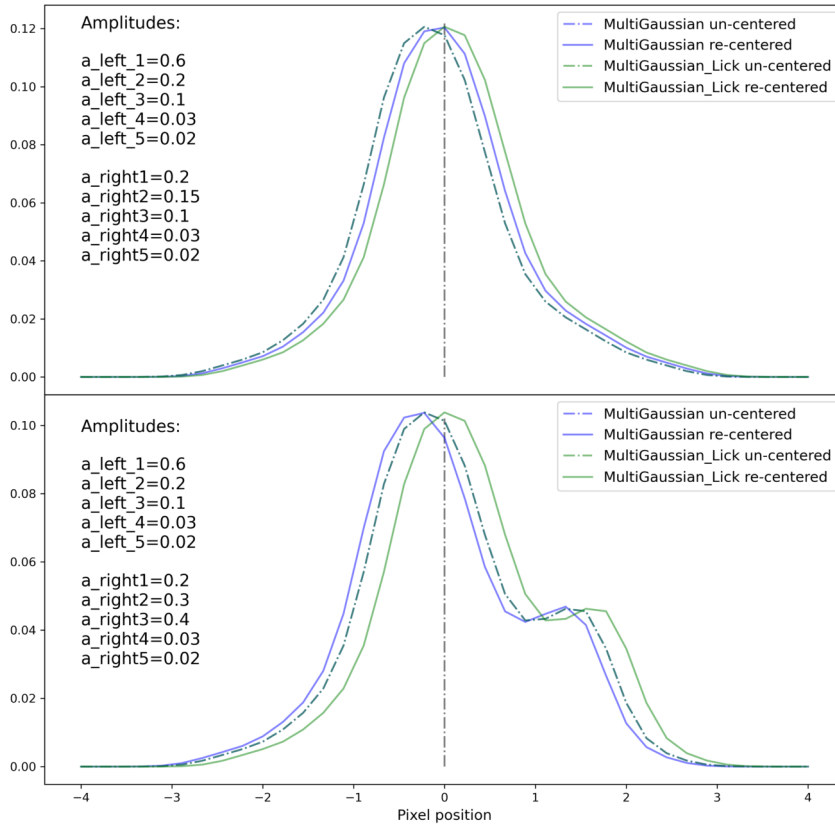


FIGURE 3.12: Exemplary LSFs, computed with the MultiGaussian code stemming from pyodine and the MultiGaussian_Lick algorithm translated from the dop code, both evaluated for satellite amplitudes a_i as indicated.

where $\text{med}()$ denotes the median. The LSFs that are used in the smoothing for chunk i are specified by a function $k(i)$, which returns the indices of neighbouring chunks both in order (cross-dispersion) and pixel (dispersion) direction. The number of orders and pixels to smooth over is defined by the user. The average is weighted by the distance of each neighbouring chunk in pixels on the CCD chip, $w_{k(i)}$.

In the FixedLSF model, the LSF array supplied by the user can be scaled by an amplitude parameter, which can be a variable in the fitting process; if the LSF should be completely fixed, that amplitude parameter needs to be fixed to 1.

Wavelength & continuum models

The currently embedded wavelength and continuum models, `LinearWaveModel` and `LinearContinuumModel`, both use the `LinearStaticModel` as a base class, which incorporates a simple linear relation of the form

$$y(x) = p_{\text{intercept}} + p_{\text{slope}} \cdot x \quad (3.17)$$

In the case of the wavelength model, the variable parameters $p_{\text{intercept}}$ and p_{slope} describe the wavelength intercept (defined at the central pixel of the chunk) and the dispersion within the chunk, respectively. The wavelengths are therefore assumed to vary linearly over the pixel range x of each chunk, which is usually an adequate description of the spectral wavelength solution. In the dop code, there is indication

of a 2nd-degree polynomial having been tested at some point, but the relevant lines of code are commented out now. The disadvantages of an even more complex model by adding another free parameter probably outweigh any gains won by the more flexible description.

The continuum model allows both the amplitude and slope of the continuum to be fitted within each chunk. In its simplest form, reduced Echelle spectra still show a large-scale roughly parabolic variation over each order due to the blaze function of the spectrograph. Over short pixel ranges, such as defined by the chunks, a linear description as given by the `LinearContinuumModel` however delivers a sufficient approximation. When de-blazed, continuum-normalized spectra are used as input to the I2 cell method, the slope parameter of the continuum model can be fixed to $p_{\text{slope}} = 0$, thus reducing the number of free parameters in the whole model.

Both the continuum and wavelength models incorporate `guess_params()` methods to estimate good starting values for the variable parameters $p_{\text{intercept}}$ and p_{slope} . For $p_{\text{intercept}}$, the wavelength (or flux) value at the central pixel of the observation chunk is used as guess, while for p_{slope} , the difference between the wavelength (or flux) values of the first and last chunk pixel, divided by the number of chunk pixels, serves as approximation.

I2 FTS and stellar template spectrum

As mentioned in Section 3.4.1, the I2 FTS is contained in the `IodineAtlas` class and builds on the more basic spectrum classes. Therefore it is equipped with the aforementioned methods that allow any wavelength-defined section of the I2 FTS to be extracted (`check_wavelength_range()` and `get_wavelength_range()`), and the relevant I2 spectrum corresponding to the modelled chunk can thus be retrieved.

Similarly, the stellar template is stored in a dedicated class building on basic spectrum classes, namely the `StellarTemplate` class in the original pyodine package. It is saved on disk as an HDF5-file, and loaded from that file into the class-object upon initialization; spectral data consists of two arrays for each template order: the normalized flux values and the wavelengths. Furthermore, the class contains additional information such as the name of the template star, the oversampling of the template spectrum and the relative velocity of the template observation with respect to a reference spectrum as well as its barycentric velocity (see Section 3.4.6). It also provides methods to extract the wavelength range which corresponds to the modelled chunk spectrum.

To extend the capabilities of the pyodine package, I added a second template class, `StellarTemplate_Chunked`. Here, the flux and wavelength arrays are stored in template chunks that were originally used in the template creation, thus allowing observation chunks that contain exactly the same spectral content to be modelled. The differences are explained in more detail in Sections 3.4.2 and 3.4.6. The modelling process described in this section remains untouched from which template is used.

The combined model and parameter sets

Before the fitting process, all the sub-models described above need to be supplied to a `SimpleModel` object upon its initialization, along with parameters concerning the evaluation process (desired oversampling, pixel range to evaluate the LSF over). This object class builds on the underlying base class `DynamicModel`, and offers a

TABLE 3.2: Overview of free parameters of the sub-models, and overall number of free model parameters n

LSF model	Wavelength	Continuum	Other	n
SingleGaussian Width: $FWHM$	LinearWave Model	LinearContinuum Model	Velocity: v Iodine depth: d_{I2} Template depth: d_s	8
SuperGaussian Center exponent: β Center width: c_0 Side amplitudes: $a_{1,2}$	$p_{intercept}$ p_{slope}	$p_{intercept}$ p_{slope}		11
MultiGaussian / MultiGaussian_Lick Left amplitudes: $a_{1,\dots,5}$ Right amplitudes: $a_{6,\dots,10}$				17
FixedLSF (Amplitude: a_0)				7/8

method `guess_params()` which returns a guess of the starting values for all variable parameters of the model if supplied with an observation chunk. For the LSF, wavelength and continuum sub-models, their respective `guess_params()` methods are executed in this step. Furthermore, three additional parameters are incorporated by the combined model object, namely the Doppler velocity of the chunk, and the relative strengths (or depths of the spectral features) of the I2 FTS and stellar template spectrum. The starting guesses for these three parameters are hard-coded as 0, 1 and 1 for the velocity, I2 FTS and stellar template, respectively, but can be altered before fitting the model. The model parameters are returned as a `ParameterSet` object, which is a Python dictionary with additional methods to easily store and retrieve parameters for each of the sub-models individually. Table 3.2 summarizes all model parameters, and their total number can be between 8 and 17, depending on the LSF model used.

The `SimpleModel` object is used to evaluate the whole model as described by Equation 3.12: When handed a `ParameterSet` object and an observation chunk, it first computes an oversampled wavelength grid by evaluating the wavelength model, using the oversampling factor supplied at initialization and the pixel width of the observation chunk. It then fetches the part of the I2 FTS corresponding in wavelength range to the oversampled wavelength grid, normalizes its flux values, scales it by the I2 depth parameter and interpolates it to the oversampled grid. The same is done with the stellar template spectrum, but correcting the required wavelength range by the Doppler factor computed from the velocity parameter. Next, the LSF model is evaluated with the supplied LSF parameters, and the oversampled I2 FTS and stellar template are multiplied and convolved with the LSF. Finally, the resulting spectrum is rebinned to the observation grid, and it is multiplied with the evaluated continuum model. In the fitting procedure, the `SimpleModel` object thus delivers the evaluated model spectrum which is then compared to the observation chunk spectrum.

Chunking algorithms

As described in the previous section, each observation is fitted in chunks, with usually equal lengths in pixels. In `pyodine`, chunks are created as `Chunk` objects, which

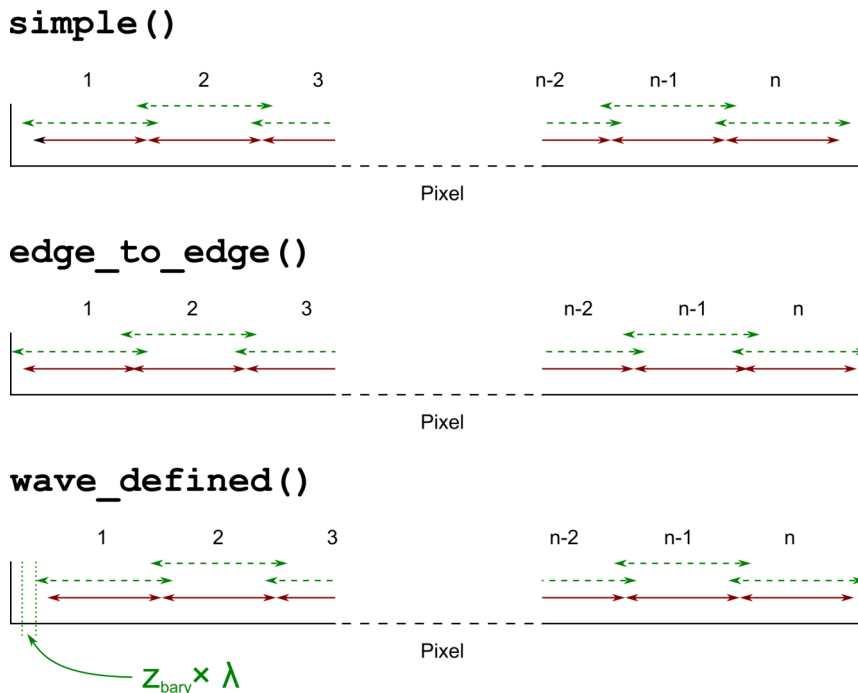


FIGURE 3.13: Schematic representation of the three different chunking algorithms for an order. Red arrows denote each chunk content, green dashed arrows include the padded regions of the chunks.

are defined in the `components` module, by supplying it with a 2D observation spectrum along with the order and pixel-range that define the chunk position within that spectrum. Furthermore, a padding argument allows to add pixels to each side of the chunk; these padded regions are used only during convolution to prevent edge effects degrading the resulting model in the actual chunk region. The `Chunk` class is based on the `Spectrum` class and thus inherits its wavelength-retrieval methods. All `Chunk` objects of an observation are bundled in a `ChunkArray` object, which is a Python list with additional methods (e.g. to return all chunks from one order of the observation).

In the original `pyodine` package, two chunking algorithms are implemented in the module `chunks`: The `simple()` function splits the orders of an observation into a number of fixed-size chunks in pixel space, centered around the middle of the orders (see Figure 3.13, top); the number of chunks per order can be defined as an input argument, if none is given the maximum possible number of chunks that fit into the order will be created. The orders to be used as well as the chunk size are also supplied by the user, as well as the padding of the chunks. This chunking algorithm may leave unused pixels at both edges of the orders, as chunks are not allowed to overlap each other (with exception of the padding regions). When modelling observations, this function is the dedicated chunking algorithm of the original `pyodine` package.

The `edge_to_edge()` function in contrast makes use of all pixels within an order by allowing chunks to overlap (see Figure 3.13, middle). The arguments that need to be supplied are the same as for the `simple()` algorithm. This function can be used when creating templates, which will then correspond in width to the observation spectra.

Both these functions work in pixel space, and when modelling a timeseries of observation spectra of a star, each chunk i created by the dedicated `simple()` algorithm will cover exactly the same pixels x_i within each observation. Due to the large variations of the barycentric velocity over time (which are on the order of several $\pm 10 \text{ km s}^{-1}$, depending on the star), the spectra of the observations can be Doppler-shifted by several pixels from observation to observation and the spectral content within each chunk will then change accordingly over time (i.e. the spectrum will move back and forth through the chunks). This is problematic when determining correct estimates of the RVs of the star with the chunk velocity combination and weighting algorithm, as the contribution of each chunk i to the final estimate is weighted by the median deviation of the timeseries of chunk i from the observation medians (see Section 3.4.7). These weights however only represent true chunk timeseries uncertainties if the spectral content of the chunks remains at least roughly the same over time.

Therefore, I added a third chunking algorithm, the `wave_defined()` function, which is based on the chunking process of the `dop` code and creates "co-moving" chunks. It requires the usage of a `StellarTemplateChunked` object as stellar template, where the original deconvolved template chunks are not stitched back together to full orders, but are stored as `TemplateChunk` objects (see Sections 3.4.2 and 3.4.6). The observation chunks are then created such that they correspond in pixel width to the original template chunks, and for each observation chunk i within an order m_i the starting wavelength $\lambda_{i,0}$ of the chunk is computed from the starting wavelength of the corresponding template chunk $\lambda_{i,0}^s$, shifted by the relative barycentric-velocity Doppler shift of the observation to the template δz_{bvc} :

$$\lambda_{i,0} = \lambda_{i,0}^s \cdot (1 + \delta z_{\text{bvc}}) = \lambda_{i,0}^s \cdot \left(1 + \left(\frac{v_{\text{bvc}}^s - v_{\text{bvc}}^{\text{obs}}}{c} \right) \right) , \quad (3.18)$$

where v_{bvc}^s and $v_{\text{bvc}}^{\text{obs}}$ are the barycentric velocities of the stellar template and the observation. The starting pixel position of the chunk $x_{i,0}$ is then the pixel within the order whose wavelength solution value $\lambda(x)$ is closest to $\lambda_{i,0}$ (see also Figure 3.13, bottom). Thus, the spectral content of each chunk i will only vary by the Doppler-shift of the RVs of the observed star from observation to observation, which for most stars corresponds to shifts of less than a pixel.

3.4.3 The fitting procedure of the chunk models

For each chunk, the model described in Equation 3.12 and evaluated by the modules presented in Section 3.4.2 is fitted to the observed spectrum flux within that chunk. As a fitting algorithm, `pyodine` uses the Python package `lmfit`, which offers a framework for Levenberg-Marquardt non-linear least-squares minimization⁵. It supplies a parameter class which allows properties of the employed model parameters, such as setting starting values, defining limits and assigning fixed values to parameters, to be modified. In `pyodine`, a `LmfitWrapper` class bundles relevant methods and functions of `lmfit` and is initialized with the combined model object `SimpleModel`; when supplied with an observation chunk and a set of `lmfit` parameters, it fits the model to the chunk. For conversion between the `pyodine` `ParameterSet` objects and `lmfit` parameter objects, a method `convert_params` is used. After fitting a chunk,

⁵<https://lmfit.github.io/lmfit-py/>

the fit result is stored in a `LmfitResult` class, which is equipped with methods to explore the fit result (e.g. return the best-fit parameter results and their errors).

In the original implementation in `pyodine`, the `lmfit` minimizer optimizes the spectrum model for each chunk i , $\hat{I}_i(x)$, by minimizing its χ^2 -statistic defined by

$$\chi_i^2 = \sum_x \left[I_i^{\text{obs}}(x) - \hat{I}_i(x) \right]^2, \quad (3.19)$$

where x are the pixels of the chunk, and $I_i^{\text{obs}}(x)$ is the observation chunk spectrum. In this framework, all pixels of the chunk always contribute the same weight in the modelling process.

However, it can be useful at times to assign non-uniform weights to pixels: For example, pixels with too high or low intensity caused by CCD inhomogeneities could then be ignored by the fitting routine by setting their respective weights to zero; also spectral features stemming from sources not considered in the model could be masked in that way (such as telluric lines, see Section 3.4.5).

If we use a more general definition of the χ^2 -statistic, including measurement uncertainties $\sigma_i^2(x)$, we can write

$$\chi_i^2 = \sum_x \frac{\left[I_i^{\text{obs}}(x) - \hat{I}_i(x) \right]^2}{\sigma_i^2(x)} \equiv \sum_x w_i(x) \left[I_i^{\text{obs}}(x) - \hat{I}_i(x) \right]^2, \quad (3.20)$$

where we have defined weights as $w_i(x) = \sigma_i^{-2}(x)$. Pixels with $w_i(x) = 0$ will now not contribute to the χ^2 -statistic and thus not influence the optimization routine. I used the approach from Equation 3.20 to replace the original minimization function in the `LmfitWrapper` class, and included the option of handing a pixel weights array $w_i(x)$ as argument along with the chunk array and parameter set when running the minimizer. When no weights are supplied, the fitter falls back to the original minimization function with uniform weights.

Finally, I added a `fit_lsfs()` method to the fitter class, which accepts an LSF model and a corresponding parameter set as arguments. It then evaluates that LSF, fits it with the LSF model incorporated in the `SimpleModel` which was used upon initialization of the fitter object, and returns the best-fit parameter set. This method can be used when fitting an observation in multiple runs with different LSF models, to find good starting values for the LSF parameters based on the LSF results from a previous run.

3.4.4 Constraining and altering starting parameters

As described in Section 3.4.2, when executing the `guess_params()` method of the combined model object, it returns a starting guess for each model parameter. The starting values for the wavelength and continuum parameters are estimated from the observation chunks themselves, and will in most cases be good estimates of the best-fit results. All other parameter guesses however cannot be easily evaluated from the observation chunks and are therefore hard-coded. Specifically the velocity parameter, which is always returned as 0, will at most times be far from the actual Doppler velocity and therefore a poor guess. The Levenberg-Marquardt fitter however requires to be supplied with starting values for the variable parameters that are

already in the vicinity of the best fit, as it is not designed to explore large areas of the parameter space. Thus, the starting value of the velocity parameter always needs to be modified to a better estimate of the true Doppler velocity before starting the modelling process.

The pyodine package makes use of a reference stellar spectrum to approximate the radial velocity of an observation: Upon creation of the stellar template, the template observation spectrum is correlated to the reference spectrum, and the relative Doppler shift between the two spectra is stored in the stellar template object as the reference velocity. When modelling an observation, the observation spectrum is correlated to the same reference spectrum to estimate its relative Doppler velocity with respect to the reference. The difference between that velocity and the template velocity then serves as a first guess for the true radial velocity of the observation with respect to the stellar template, and that value can be used for the velocity parameter of the model.

Two different reference spectra are implemented in pyodine: one of Arcturus, and one of the Sun. It is obviously important to use the same reference for the observations as was used in the creation of the stellar template. The desired reference spectrum is loaded using the `SimpleNormalizer` object in the `template.normalize` module. When supplied with a spectrum, the object's method `guess_velocity()` returns the relative velocity shift between the supplied spectrum and the reference.

In the parameter input file, I included a `constrain_parameters()` method which can be used to alter the starting values of all model parameters. It is supplied with a dictionary from the main routine which contains the relative velocity shift computed by the method described above, and sets the starting value of the model velocity parameter to that value. Also other parameter guesses can be altered in this method, constrained within certain limits or fixed at set values. For example, when modelling an observation in multiple runs, starting values can be set to the best-fit results from a previous run.

3.4.5 Pixel weights

While the original pyodine code assumed uniform weights for all pixels in the fitting process, my modified version allows differing pixel weights to be assigned. This can be used for instance to compute weights that reflect the relative flux uncertainties. The dop code uses pixel weights defined by

$$w(x) = \left[I^{\text{obs}}(x) + (nI^{\text{obs}}(x))^2 \right]^{-1} , \quad (3.21)$$

where n is the typical relative noise in a flat spectrum (in the dop code: $n = 0.008$). As I wanted to test the influence of using different pixel weights in the modelling process, I added a method `compute_weight()` to all spectrum classes. When supplied with the keyword "lick", it calculates and returns the weights of the flux values contained in the object according to Equation 3.21; with the keyword "flat" or no keyword given, it returns an array with uniform pixel weights. These weights arrays can be further modified by the user (e.g. setting weights of somehow impaired pixels to zero) to finally be used in the fitting routine.

Telluric mask

Spectra that have been recorded in ground-based facilities are affected by telluric absorption and emission features: Molecules in the Earth's atmosphere imprint spectral lines onto the spectrum of the target object, caused by rotational-vibrational transitions for example in water (H₂O), oxygen (O₂), carbon-dioxide (CO₂) and methane (CH₄). Whereas some of these lines are very weak (so-called micro-tellurics), others may absorb all the light in their line cores. The depth, shape and exact wavelength of a given telluric line is influenced by the optical depth along the light path, turbulent wind motions, atmospheric temperature and partial-pressure structure and the chemical composition of the atmosphere. Therefore, the telluric absorption varies with time, position of the observatory on the Earth's surface (especially altitude), and the airmass of the observed target.

In precise RV measurements, telluric absorption can lead to larger measurement uncertainties or even to incorrect results if not accounted for: With the CCF method, correlation between telluric features and the lines from the used stellar binary mask may lead to a broadened and even shifted CCF; with the I2 cell method, telluric lines generally increase the unaccounted noise in the model, and if they lie close to a stellar feature they will lead to an asymmetric broadening of that line, possibly resulting in a shifted template and thus a wrong RV estimate for the respective chunk.

The influence of telluric absorption however greatly depends on the wavelength range used for RV extraction: In the near-infrared regime around and just above $\sim 1 \mu\text{m}$, where M-dwarfs reach their highest luminosity, many wide and strong telluric features caused mostly by O₂ and H₂O molecules lead to a severe contamination of spectra; this greatly complicates the determination of precise RVs using instruments that are specialized for this wavelength range. At optical wavelengths however, telluric lines are generally fewer and weaker, and some large wavelength ranges exist with basically no significant features present.

Different methods about how to deal with telluric lines are used, and they can be divided into two general approaches: The first is to simply mask all pixels in a spectrum whose wavelengths are affected by telluric absorption features above a certain strength threshold before computing radial velocities. This has been applied successfully in many RV extraction pipelines, especially those working at optical wavelengths. The second and more elaborate approach is to actually correct the observed spectra for telluric lines, either from empirical data or through forward-modelling. Several techniques have been developed with differing success; their great advantage of not losing any spectral information at wavelengths affected by tellurics comes with the cost of heavier computation or the need of additional observations to gather the necessary telluric information.

In the wavelength range accessible to the I2 cell method, ~ 5000 to 6200 \AA , telluric features are particularly rare, with only three broad absorption bands between $5034 - 5080 \text{ \AA}$, $5683 - 5810 \text{ \AA}$ and $5859 - 6020 \text{ \AA}$ (as seen in Figure 3.14, top). As the depth of the iodine lines already strongly decreases at wavelengths longer than $\sim 5900 \text{ \AA}$, and the strongest telluric band falls into that wavelength regime, orders covering that spectral range are usually omitted in the extraction of the RVs. This only leaves the telluric bands just above 5000 \AA and around 5700 \AA , along with two narrow features just above 5400 \AA , most of which are usually quite weak.

Nevertheless, I implemented a `SimpleTellurics` class into the `pyodine` package, which allows one of three telluric masks to be loaded:

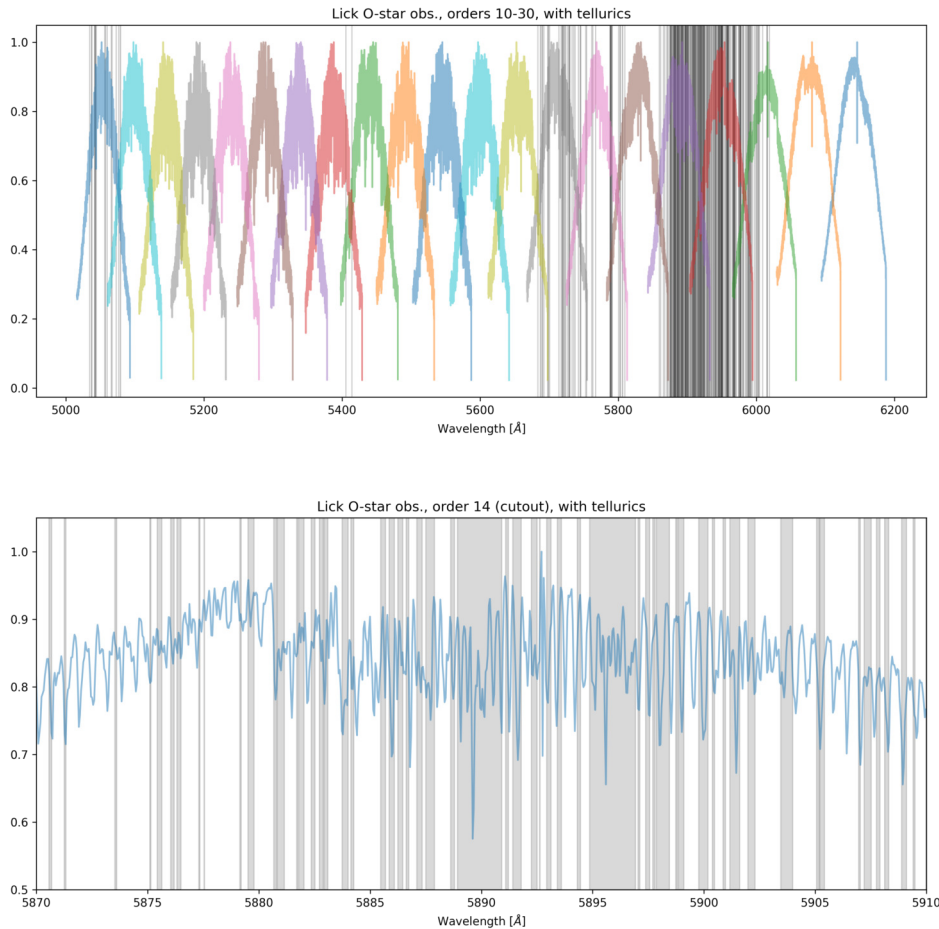


FIGURE 3.14: *Top*: A normalized O-star spectrum from the Lick survey, obtained with I2 cell in the light path, orders 10 to 30, with wavelengths possibly affected by telluric features shaded in grey. The telluric wavelength information stems from the mask used in SERVAL (see main text). *Bottom*: A cutout of order 14, showing strong telluric absorption features from H₂O molecules.

- a mask created for telluric correction of CARMENES spectra, contained in the SERVAL pipeline (Zechmeister et al., 2018) and based on a telluric atlas from Wallace et al. (2011);
- a telluric line atlas created through the online High-Resolution Transmission Molecular Absorption Database (HITRAN)⁶;
- a telluric line atlas called UVES from Keck Observatory, created for the HIRES spectrograph⁷.

Once initialized, the class object offers a method `is_affected()`, which can be used to check whether a certain wavelength or a list of wavelengths is affected by telluric features. Thus, in the RV analysis, observation pixels containing telluric contamination can be identified and their respective weights set to zero in order to omit them in the modelling process.

⁶<https://hitran.org/>

⁷<https://www2.keck.hawaii.edu/inst/common/makeewww/Atmosphere/index.html>

Bad pixel mask

Extracted spectra may still contain pixels which do not represent the true spectrum but are impaired, for instance due to dead pixels of the detector. These appear as regions of low flux, similar to actual stellar absorption lines, and may greatly affect the model results of the chunks containing these pixels. Furthermore, the spectra might still contain traces of cosmics that have not been corrected for completely in the reduction process. For better results it can be helpful to mask the affected pixels out by setting their corresponding weights to zero.

I included a `BadPixelMask` class in `pyodine`, to automatically scan an observation spectrum for pixels affected by instrumental defects and cosmics and return a bad-pixel mask. The scanning algorithm is based on the IDL routine `find_cosmic.pro` written by Sabine Reffert, which makes use of the fact that fake spectral features caused by bad pixels (and cosmics) usually have much steeper wings and lead to quicker intensity changes than true absorption (and emission) lines. The code walks through the pixels x_i of each order and evaluates the local absolute value of the second-order derivative of the continuum-normalized intensity $I^n(x)$ through

$$\frac{d^2 I^n(x)}{dx^2} = |[I^n(x_{i+1}) - I^n(x_i)] - [I^n(x_i) - I^n(x_{i-1})]| \quad . \quad (3.22)$$

These values are compared to a threshold number, and pixel regions with higher $\frac{d^2 I^n(x)}{dx^2}$ than that threshold are identified as impaired pixels. The threshold value can be supplied as an optional argument and defaults to 0.18. The bad-pixel mask is of same shape as the evaluated spectrum, with ones at the found positions of bad pixels and zeros elsewhere, and can be used to set the weights of bad pixels to zero.

Additionally, I implemented a `correct_spectrum()` function in `pyodine`, to correct the regions with zero pixel weights in a spectrum: When supplied with an observation spectrum, corresponding pixel weights array and the orders that are to be corrected, it linearly interpolates flux values around zero-weighted pixels over the affected regions. This certainly does not restore the true spectrum, and when modelling observations one should refrain from correcting the spectra and instead just mask out bad pixels in the fitting process. In the creation of stellar template spectra however, this function can be useful to correct bad pixel regions spanning only few pixels of the template observation spectrum before being deconvolved, thus preventing "fake" absorption features caused by instrumental defects to appear in the stellar template.

3.4.6 Template creation

For the modelling process of the observations described above, a stellar template spectrum is required for every single target star, which should ideally have a very high S/N and be completely free of any instrumental effects. This is impossible to achieve directly with the same spectrograph as used for the RV observations: Even when keeping the I2 cell out of the light path, the stellar spectra will still be contaminated by the instrument LSF. Using a different instrument with a much higher resolving power, such as a Fourier-Transform Spectrometer, however is not possible for most stars as it will be impossible to reach a high S/N then.

Therefore, the stellar template is created in a deconvolution process from an observation spectrum of the desired star, using the same spectrograph as for the observations but obtained without the I2 cell in the light path. To determine the LSF of the instrument, which is needed in the deconvolution, spectra of hot, rapidly rotating O- or B-stars are taken with the I2 cell in the light path right before and after observing the template star. As the spectra basically do not show any stellar features, they represent the I2 spectrum modulated by the LSF of the instrument. These spectra, $I^{O/B}(x)$, are then modelled essentially the same way as the observations described in Section 3.4.2, only the stellar template is replaced by a constant 1 (as ideally there are no stellar spectral features present), and the radial velocity is fixed at 0. The arithmetic description of the model therefore reduces to

$$\hat{I}_i(x) = k(x) T_{I2}(\lambda(x)) * L(x) \quad . \quad (3.23)$$

The resulting LSF $L(x)$ can now be used in a deconvolution of the stellar template observation: If we denote the stellar observation as i and the LSF determined through the O-/B-star observations as s , we are interested in the "true" stellar object spectrum o which satisfies

$$i = s * o + n \quad , \quad (3.24)$$

where n is the noise in the observation. Butler et al. (1996) used a modified version of the Jansson technique for deconvolving the stellar template, involving an iterative process (as described in Gilliland et al., 1992); pyodine incorporates a slightly different recipe developed by Crilly et al. (2002) on the basis of the algorithm by Agard et al. (1989), which uses the same iterative steps, but prefilters the input data: An estimate \hat{o} of the true stellar spectrum before its contamination by the LSF can be found starting in a first step $k = 1$ with

$$\hat{o}^1 = i * s + r(\hat{o}^1) \cdot [i - s * i * s] \quad , \quad (3.25)$$

and continuing with

$$\hat{o}^{k+1} = \hat{o}^k + r(\hat{o}^k) \cdot [i * s - s * s * \hat{o}^k] \quad , \quad k \geq 1 \quad . \quad (3.26)$$

The relaxation function $r(\hat{o}^k)$ is defined by

$$r(\hat{o}^k) = r_0 \left[1 - \frac{2}{b-a} \left| \hat{o}^k - \frac{a+b}{2} \right| \right] \quad , \quad (3.27)$$

where r_0 is a constant controlling the step size, and b and a are the allowed upper and lower bounds of the estimated stellar spectrum. In our case, b corresponds to the continuum of the spectrum and a to the maximum line depth. This deconvolution version allows for effective noise control and reaches faster convergence than other methods (compare Crilly et al., 2002).

In pyodine, when modelling the O-/B-star observations, the same modules and routines are used as for the fitting of the later target spectra; the only difference is that no template spectrum is passed to the SimpleModel class upon initialization, and

the parameters for the model velocity and template depth, contained in the `lmfit` parameter set, are fixed at 0 and 1, respectively. If multiple O-/B-star spectra are accessible, they are summed up in a `SummedObservation` object which is used in the fitting process, thus allowing a better estimate of the LSF by modelling a spectrum with higher S/N .

Deconvolution of a template observation is carried out in a `SimpleDeconvolver` object, which upon initialization takes as input arguments a list of all chunks of the O-/B-star observation, the used `SimpleModel` object, and the model parameter results. In the original `pyodine` package, the method `deconvolve_obs()` then deconvolves the normalized template observation using hard-coded values for the deconvolution parameters r_0 , a , b , and the maximum number of iterations k_{\max} . The template observation is deconvolved in chunks corresponding to the O-/B-star chunks, and using the corresponding best-fit LSF of the chunks; the deconvolution is carried out on an oversampled pixel grid, with the oversampling factor also hard-coded to 10. The resulting deconvolved flux and wavelength values of all chunks within an order are finally stitched back together to an oversampled, deconvolved stellar template order. In the end, all orders are stored in a `StellarTemplate` object, which is saved on disk as a HDF5 file.

In my modified version of `pyodine`, the deconvolution parameters and the oversampling factor of the stellar template can be passed as a dictionary to the method `deconvolve_obs()`, so that they can be easily defined and altered by the user. Additionally, the velocity offset of the template observation to the reference spectrum and the barycentric velocity at time of the observation are required as arguments, and this information is stored in the resulting `StellarTemplate` object. Also, inspired by the `dop` code, I implemented a short routine in the deconvolution, which computes the relative change of the deconvolved flux estimates δ^k from step to step, and if the change become smaller than a certain threshold before k_{\max} is reached the iteration is stopped. The threshold can also be supplied by the user in the dictionary passed to `deconvolve_obs()`.

Furthermore, I incorporated the `ChunkedDeconvolver`, a second deconvolution class. This one works similarly to the routines describes above, except that it does not stitch the deconvolved oversampled template chunks back together, but stores them as `TemplateChunk` objects in a `StellarTemplate_Chunked` template object. When modelling observations with that template, it allows the `wave_defined()` function to be used to create the observation chunks, which will consequently cover the same spectral features as the respective template chunks (see Section 3.4.2).

3.4.7 Combination of chunk velocities

In the modelling process of an observation, each of the n_{ch} chunks is modelled individually and independently from all others, and for the RV of the observation v_{RV} we therefore receive n_{ch} independent estimates, namely the Doppler velocities used in the chunk models, v_{meas} . The easiest possibility to compute v_{RV} would therefore be a simple mean or median of all individual chunk velocities.

However, the individual chunks cover different parts of the spectrum with differing numbers and depths of stellar absorption lines, and thus non-uniform RV content. Butler et al. (1996) derive an estimate of the photon-limited error on the RV of a chunk $\sigma_{\bar{v}}$:

$$\sigma_{\bar{v}} = \left[\sum_i \left(\frac{dI/dV}{\epsilon_i} \right)^2 \right]^{-\frac{1}{2}}, \quad (3.28)$$

where the sum is over all pixels in the chunk i , ϵ_i is the uncertainty in the residual intensity at pixel i (e.g. the inverse square-root of the flux), and dI/dV is the local slope of the spectrum flux. Chunks with more spectral features will have more pixels with large dI/dV , and thus a smaller $\sigma_{\bar{v}}$ — they contain more Doppler information. Therefore, when combining the results from all chunks, it will be beneficial to introduce some form of weighting, as the chunks with more Doppler information will give a more trustworthy estimate of the true RV than chunks with few or no spectral features. Additionally, the weighting procedure will allow the rejection of chunks for which the modelling failed due to some reason, or weight down chunks that are somehow impaired, for example by instrumental effects.

The original pyodine package includes an interactive code with a combination algorithm that was largely inspired by the procedure used in iSONG. Furthermore, in the dop code a complex weighting algorithm specifically designed for the use with Lick spectra is used, which requires some instrument-/observation-specific input from the user. Both codes follow a common philosophy, where the chunk weights are calculated both on the basis of the relative deviation of individual chunks within one observation, as well as of the average deviation of whole chunk timeseries over all given observations. It is then important that chunks of the same index i cover at least roughly the same spectral content over all observations, otherwise the weights would not reflect the true Doppler content of the chunk timeseries.

In order to comply with the overall structure of the Waltz DRS, I created a new module called `timeseries` within the modified pyodine package, and included a `CombinedResults` class which can be used to collect all important model results from a list of observations and save them as a common data cube. Namely, this then includes several 1D vectors of shape $[n_{\text{obs}}]$, and 2D matrices of shape $[n_{\text{obs}} \times n_{\text{ch}}]$, where n_{obs} and n_{ch} are the number of observations and chunks, respectively:

- chunk velocities v_{ij} , and velocity errors from the fit Δv_{ij} ;
- goodness of the chunk fits $(\chi_{\text{red}}^2)_{ij}$;
- median counts of the chunks \bar{f}_{ij} ;
- barycentric velocities v_i^{bc} , respective dates t_i , and filenames of the original spectra and of the individual results of the observations.

Furthermore, general information like the I2 FTS, stellar template, oversampling and LSF model used in the analysis of all observations are included, as well as all other best-fit parameters for each chunk and their respective errors.

I modified the existing velocity combination algorithm from pyodine to work with this newly created class and to merge it into the overall structure of the Waltz DRS. Additionally, I implemented a second combination algorithm which is based on the dop-code routine used for the combination of chunk velocities. Both accept the input of a rejection mask, consisting of a text-file with a list of observation names to be ignored in the combination of chunk velocities, thus allowing the user to exclude bad observations from the computation of the RV timeseries. In the following, I will briefly outline the SONG- and dop-based weighting algorithms.

SONG velocity weighting algorithm

This weighting algorithm was implemented in the original pyodine package by R.T. Rasmussen, and is based on the routine used in iSONG. For mean and standard deviation calculations, it uses outlier-resistant functions, called R_{mean} and R_{std} hereafter. They are implemented as a Python version of the respective functions in the IDL package ROBUST⁸.

First, all chunk velocities v_{ij} are corrected by the barycentric velocities of each observation. Next, the chunk velocities are averaged over the observations i , using a resistant mean $R_{\text{mean}}()$, which yields the average chunk series velocities \bar{v}_j . If we subtract their resistant mean over the chunks j , we receive the average relative velocity offsets of the chunk series from the observation means, \bar{v}_j^0 :

$$\bar{v}_j = R_{\text{mean}}(v_{ij}, i) \quad \Rightarrow \quad \bar{v}_j^0 = \bar{v}_j - R_{\text{mean}}(\bar{v}_j, j) \quad . \quad (3.29)$$

By subtracting these chunk series velocity offsets from the individual chunk velocities v_{ij} , we can compute zeropoint-corrected chunk velocities:

$$v'_{ij} = v_{ij} - \bar{v}_j^0 \quad . \quad (3.30)$$

Now, we take the resistant mean of the corrected chunk velocities along the chunks j within each observation, yielding average observation velocities \bar{v}'_i . We center these around zero by subtracting their resistant mean over the observations i , yielding an unweighted RV timeseries RV_i^{RV} :

$$\bar{v}'_i = R_{\text{mean}}(v'_{ij}, j) \quad \Rightarrow \quad RV_i^{w=0} = \bar{v}'_i - R_{\text{mean}}(\bar{v}'_i, i) \quad . \quad (3.31)$$

In an ideal world with perfect models and free of any noise sources, $RV_i^{w=0}$ would be the signal. By subtracting it from the individual chunk velocities v_{ij} , we can thus find a representation of the velocity noise contribution of each chunk:

$$\eta_{ij} = v_{ij} - RV_i^{w=0} \quad . \quad (3.32)$$

A measure of the noise of each chunk series is then computed by taking the resistant standard deviation $R_{\text{std}}()$ along the observations i :

$$\sigma_j = R_{\text{std}}(\eta_{ij}, i) \quad . \quad (3.33)$$

Now, weights w_j can be computed as the inverse square of these σ_j :

$$w_j = \sigma_j^{-2} \quad , \quad (3.34)$$

and a weighted RV timeseries can be derived by summing over the weighted zeropoint-corrected chunk velocities:

⁸<https://idlastro.gsfc.nasa.gov/contents.html#C17>

$$RV_i^w = \frac{\sum_j v'_{ij} w_j}{\sum_j w_j} . \quad (3.35)$$

The corresponding velocity errors are then the inverse square-root of the sum of the weights:

$$\sigma_i^{RV} = \frac{1}{\sqrt{\sum_j w_j}} . \quad (3.36)$$

However, when using this simple weighting by the average chunk series noise σ_j , large outlier chunk velocities within an observation will not be weighted down particularly strongly and still greatly influence the final result. Furthermore, the velocity errors σ_i^{RV} are the same for each observation, as the weights w_j are always the same.

Therefore, before computing the weights we can derive a relative deviation of each chunk velocity in terms of its respective chunk series deviation, δ_{ij} , by subtracting the resistant mean of the noise over the chunks j in an observation from the noise of each chunk within that observation, and dividing it by σ_j . These chunk-for-chunk deviations are centered around zero by subtracting their resistant mean over the chunks j :

$$\delta_{ij} = \frac{\eta_{ij} - R_{\text{mean}}(\eta_{ij}, j)}{\sigma_j} \Rightarrow \delta_{ij}^0 = \delta_{ij} - R_{\text{mean}}(\delta_{ij}, j). \quad (3.37)$$

Finally, pyodine has implemented a reweight function from iSONG to weight down outlier chunk velocities within an observation:

$$\omega_{ij} = \frac{1}{1 + (|\delta_{ij}^0| / (a \cdot s))^\beta} , \quad (3.38)$$

where a , s and β are constants (in iSONG: $a = 1.8$, $s = 2$, $\beta = 8.0$).

Equation 3.34 is then modified to incorporate the ω_{ij} , and the weights thus become a 2D matrix:

$$w_{ij} = \left(\frac{\sigma_j}{\omega_{ij}} \right)^{-2} . \quad (3.39)$$

Accordingly, the equations for the weighted RV timeseries and the respective errors are updated to:

$$RV_i^w = \frac{\sum_j v'_{ij} w_{ij}}{\sum_j w_{ij}} , \quad \sigma_i^{RV} = \frac{1}{\sqrt{\sum_j w_{ij}}} . \quad (3.40)$$

Velocity combination algorithm based on the dop code

The velocity weighting algorithm from the dop package differs from the routine described above in that it does not use outlier-resistant means and standard deviations, but instead completely rejects bad chunks in the computation of the RV timeseries.

First, all observations are scanned and those with a median χ_{red}^2 from the chunk fits above a certain threshold are rejected, as well as those with too low or too high chunk-averaged median counts \bar{f} . The limits, $\chi_{\text{red,max}}^2$, \bar{f}_{min} and \bar{f}_{max} can be chosen by the user. For the remaining observations, a weights array w_{ij} is set up, and an error array is filled with the velocity errors from the fits Δv_{ij} . Similar to the previous section, a representation of the noise contribution of each chunk is found by subtracting an unweighted RV timeseries (here a simple median) from the barycenter-corrected chunk velocities $v_{i,j}$:

$$\eta_{ij} = v_{ij} - \text{median}(v_{ij}, j) \quad , \quad (3.41)$$

where $\text{median}(v_{ij}, j)$ denotes the median of the velocities along the chunks j . Using Chauvenet's criterion, outlier chunks from η_{ij} are identified for each observation i , and their respective weights w_{ij} are set to 0, and errors e_{ij} to 99 m s^{-1} .

In a next step, a sigma clipping of badly modelled chunks is performed, by setting the weights of chunks with $(\chi_{\text{red}}^2)_{ij}$ or e_{ij} above a certain user-defined percentile limit to 0. If the percentile limit delivers a χ_{red}^2 larger than $\chi_{\text{red,max}}^2$ from above, the latter is used as limit.

Now, again similar to the previous section, the noise of each chunk series is computed as the standard deviation $\text{std}()$ along the observations i :

$$\sigma_j = \text{std}(\eta_{ij}, i) \quad . \quad (3.42)$$

The chunk-for-chunk deviations in terms of σ_j are simply

$$\delta_{ij} = \text{median} \left(\frac{|\eta_{ij}|}{\sigma_j} \right) \quad , \quad (3.43)$$

and the non-zero elements of the weights array are filled with the inverse square of these deviations δ_{ij} multiplied with the average chunk series noise σ_j :

$$w_{ij} = (\sigma_j \cdot \delta_{ij})^{-2} \quad . \quad (3.44)$$

Finally, a second sigma clipping with a user-defined percentile limit is used to reject the chunk velocities with largest residuals from the median observation velocities. The remaining velocities are then combined to the final weighted RV timeseries and errors are calculated as specified in Equation 3.40.

3.4.8 Workflow of the observation modelling

I implemented the exact algorithms for modelling observations in main routines, which import the modules and functions described in the previous sections. Two

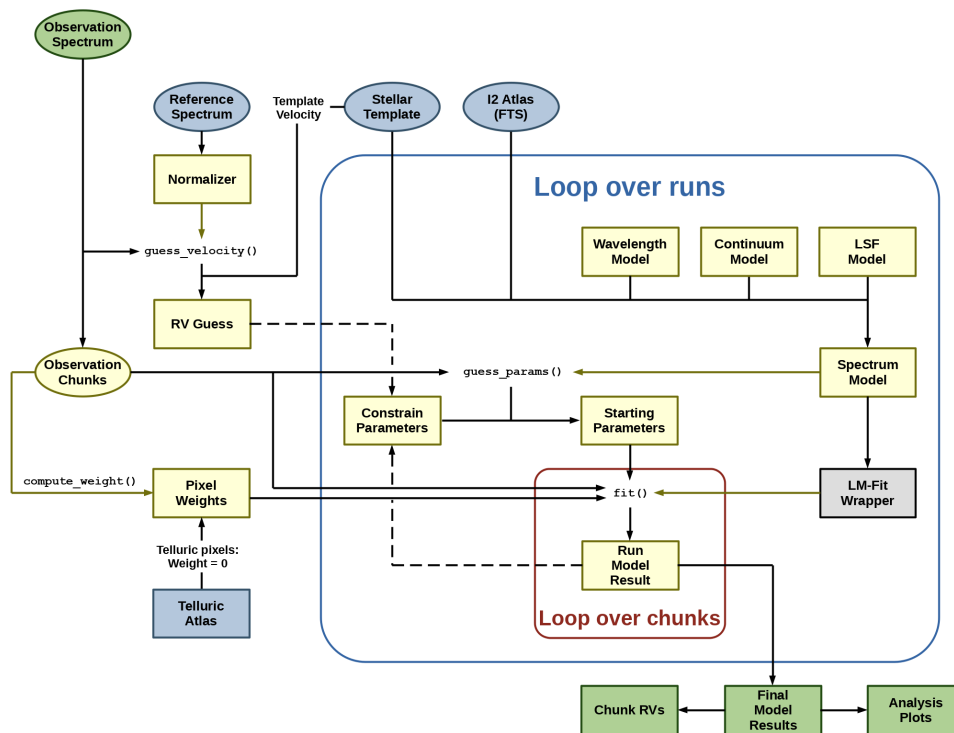


FIGURE 3.15: Flowchart of the extraction of radial velocities from the reduced spectra using the I2-cell method.

main routines have been created: `pyodine_multiprocess.py` is used to model observations using the stellar template implementation of the original `pyodine` code, `StellarTemplate`, which consists of continuous orders and where the observation chunks are created fixed in pixel space; `pyodine_multiprocess_chunked.py`, the second main routine, in contrast requires the `StellarTemplate_Chunked` object to be used as stellar template.

Both routines import and initialize a parameter input object, which contains all instrument- and observation-specific parameters, and definitions about the number of runs to be used in the model process. Furthermore, it possesses a method to constrain or alter starting parameters for each model run. I also included a multiprocessing feature in the main routines, so when analyzing multiple observations they can be modelled in parallel processes, thus effectively reducing the total time needed.

The basic workflow of the main routines is depicted in Figure 3.15, and can be summarized in the following points:

1. Load the observation spectrum to be analyzed, the suitable stellar template, and the I2 atlas (FTS).
2. Initialize a `Normalizer` object with a reference spectrum (either Arcturus or Sun, needs to be the same as used in the template creation), and use the method `guess_velocity()` to return an RV estimate of the observation spectrum relative to the reference spectrum; the difference of that estimate to the template velocity will serve as the initial RV guess in the modelling process.
3. Split the observation spectrum into chunks, either with the `simple()` or the `wave_defined()` algorithm.

4. Compute pixel weights to be used in the modelling process with the spectrum method `compute_weight()`. If no weights are required, the keyword "flat" will set all pixel weights to unity.
5. If required, load a telluric atlas. If a telluric atlas is initialized, it will be used to set weights of pixels affected by tellurics to zero in the chunk modelling process.
6. Model the observation; this can be performed in multiple runs, to determine the parameters better:
 - (a) Initialize the spectrum model, `SimpleModel`, with the stellar template, the I2 atlas, and the desired wavelength, continuum and LSF models for this run; initialize the fitter `LmfitWrapper` with the spectrum model.
 - (b) Create starting parameter estimates using the `guess_params()` method of the spectrum model for each chunk. Change, constrain or fix/vary the starting parameters if required (e.g. set the velocity parameter to the initial RV guess estimated in step 2, or constrain starting values around the respective model results from a previous run); this is done using the `constrain_parameters()` method of the parameter input object.
 - (c) Loop over all chunks and model the chunk spectra with the `fit()` method of the fitter, with the starting parameter estimates and pixel weights as additional input arguments. Store the resulting `LmfitResult` objects of all chunks in a results list.
 - (d) Optionally: save the chunk results to a file. Possibly create and save analysis plots (e.g. wavelength and velocity distributions, residual histograms, model and observation of individual chunks), and compute median parameters.
7. Save the final results from the last run to a file.

3.4.9 Workflow of the template creation

A second set of main routines has been written for the creation of stellar template spectra, again one for templates of the type `StellarTemplate`, and one for chunked `StellarTemplate_Chunked` template objects. As the main analysis routines for the observation modelling, these algorithms import a parameter input object which contains all relevant parameters, and whose method `constrain_parameters()` can be used to constrain or alter starting parameters of the model.

The workflow of the template creation process, depicted in Figure 3.16, follows these steps:

1. Load the stellar template observation, the I2 atlas (FTS), and the O-/B-star observations. Sum up all O-/B-star spectra to a `SummarizedObservation` object.
2. Initialize a `Normalizer` object with a reference spectrum (either Arcturus or Sun), and use the method `guess_velocity()` to return an RV estimate of the template observation spectrum relative to the reference spectrum; normalize the template observation spectrum.
3. Split the summed O-/B-star spectrum into chunks, either with the `simple()` or the `edge_to_edge()` algorithm.

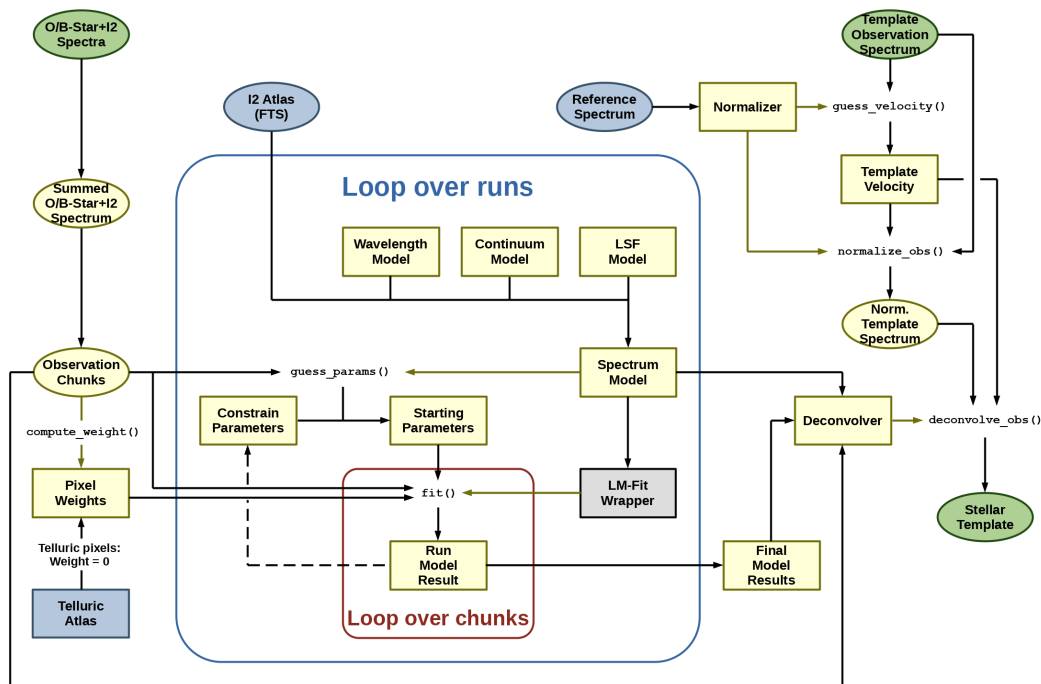


FIGURE 3.16: Flowchart of the creation of stellar templates, using the I2-cell method.

4. Compute pixel weights to be used in the modelling process with the spectrum method `compute_weight()`. If no weights are required, the keyword "flat" will set all pixel weights to unity.
5. If required, load a telluric atlas. If a telluric atlas is initialized, it will be used to set weights of pixels affected by tellurics to zero in the chunk modelling process.
6. Model the O-/B-star chunks; this can be performed in multiple runs, to determine the parameters better:
 - (a) Initialize the spectrum model, `SimpleModel`, with the I2 atlas, and the desired wavelength, continuum and LSF models for this run; initialize the fitter `LmfitWrapper` with the spectrum model.
 - (b) Create starting parameter estimates using the `guess_params()` method of the spectrum model for each chunk. Change, constrain or fix/vary the starting parameters if required (e.g. set the velocity parameter to the initial RV guess estimated in step 2, or constrain starting values around the respective model results from a previous run); this is done using the `constrain_parameters()` method of the parameter input object.
 - (c) Loop over all chunks and model the chunk spectra with the `fit()` method of the fitter, with the starting parameter estimates and pixel weights as additional input arguments. Store the resulting `LmfitResult` objects of all chunks in a results list.
 - (d) Optionally: save the chunk results to a file. Possibly create and save

analysis plots (e.g. wavelength and velocity distributions, residual histograms, model and observation of individual chunks), and compute median parameters.

7. Initialize a deconvolver object (`SimpleDeconvolver` or `ChunkedDeconvolver`, depending on the routine), and feed it with all necessary data and information to deconvolve the normalized template observation spectrum. Save the resulting stellar template object to file.

3.5 Testing of the RV analysis code

The original pyodine code was tested by R.T. Rasmussen both on simulated data and real SONG observations of the stars Arcturus and HD 185144. These tests are described in his master's thesis (Rasmussen, 2016) and generally showed good results: Using the simulated data, Rasmussen (2016) proved the ability of pyodine to extract the previously injected Doppler velocity correctly from a spectrum. For a short high-cadence timeseries of SONG observations of Arcturus, spanning a total time of ~ 45 minutes, pyodine showed a RV rms of roughly 1 m s^{-1} after subtraction of a linear trend, and individual measurement uncertainties in the same order of magnitude, which is comparable to the results from the dedicated iSONG reduction software. This proved the short-term precision of pyodine. Finally, tests on SONG observations of HD 185144, taken over 761 days, resulted in a RV curve very similar to the iSONG results. A modulation with a roughly 1 yr-period is seen in both RV timeseries and has been observed for some time in observations of several stars with SONG; probably this can be explained by inhomogeneities of the pixel spacings in the detector (Frank Grundahl, Sabine Reffert, private communication).

With these tests being only few in number and not published, and due to my significant modifications to the code, I again performed a number of tests of the updated version of the pyodine package which serves as the RV analysis tool in the Waltz DRS. Unfortunately, the Waltz project itself did not deliver suitable spectra in time for testing of the code, so I had to rely on other data. From our former Lick survey of G-/K-giant stars we have a large archive consisting of extracted spectra from all observing nights and corresponding RVs (and other analysis results) for all targets, the latter stemming from the original Doppler analysis code by Butler et al. (1996). Additionally, as the RV survey of evolved stars was partly resumed with observations at the SONG telescope, our group possesses a number of reduced SONG spectra for several of our targets, along with RV results from the dedicated iSONG pipeline for all of these observations.

The Lick and SONG data therefore allow two broad tests to be performed:

- performance of the RV analysis code in the Waltz DRS, by direct comparison of its results with the RVs (and possibly other output) from the dedicated reduction pipelines;
- feasibility of easy and fast adaptation of the package to a new instrument, through optimization of parameters.

In all the tests shown in this section, unless specified otherwise, I used the method of wavelength-defined observation chunks with similar spectral content, that is the `StellarTemplate_Chunked` class for stellar templates, and `wave_defined()` chunking algorithm for the observations. For the combination of chunk velocities to RV

TABLE 3.3: Properties of simulated observations

LSF model/ parameters	SingleGaussian $FWHM = 2.4$ pix	MultiGaussian $a_0 = 0.05, a_1 = 0.15, a_2 = 0.3,$ $a_3 = 0.45, a_4 = 0.62,$ $a_5 = 0.58, a_6 = 0.5,$ $a_7 = 0.29, a_8 = 0.12, a_9 = 0.05$	MultiGaussian_Lick
I2 spectrum/ parameters	SONG I2 FTS Spectrum depth: 0.98		
Stellar template/ parameters	Simulated template spectrum Spectrum depth: 1.0		
Continuum	Parabola over each order, simulating blaze-function		
Injected RVs	$[-2000, -1000, 0, 1000, 2000]$ m s ⁻¹		

timeseries, the SONG-based algorithm, which is part of the original pyodine package, has been used, as it is more flexible and in most cases showed better timeseries results during early tests. Sometimes I will report results from the dop-based algorithm for comparison.

When modelling spectra with an LSF model different from the `SingleGaussian` LSF, I always used multiple runs. Early tests showed that when immediately modelling an observation with one of the asymmetric LSFs, the free parameters would easily run away to extreme high or low numbers, resulting in highly distorted LSFs for some chunks. This can be explained by the fact that the wavelength scale of the observation, which was derived with the help of ThAr calibration spectra, usually does not correspond perfectly to the wavelengths of the I2 FTS. When the fitter then starts to model the observation, the differences between the wavelength scales not only influence the evolution of the wavelength parameters of the model but also of the LSF parameters. With the `SingleGaussian` LSF this does not happen due to its symmetric shape. Therefore, whenever I used either the `SuperGaussian`, `MultiGaussian` or `MultiGaussian_Lick` LSF, each observation was modelled in at least two runs: In a first run, using the `SingleGaussian` LSF, good estimates of wavelength and continuum parameters were determined; these then served as starting values for the second run with the asymmetric LSF model. This routine was also adopted for the simulated spectra.

First, Section 3.5.1 presents tests using simulated data as input, thus proving the general abilities of the package. Then, in Section 3.5.2, I describe first results from the analysis of Lick observations with the Waltz DRS, and finally I show an early test run of the software on SONG spectra of the star ϵ Cyg in Section 3.5.3.

3.5.1 Results from simulated spectra

The simulated template was created as follows: A real extracted spectrum of the star HIP 102488, which was obtained with the SONG telescope, was loaded to use its wavelength scale as a starting structure for the template. From the observation, 24 orders of 2048 pixels each were chosen around the wavelength range used in the I2 cell method. The templates used in the I2 analysis however are over-sampled, and the oversampling factor may vary between 4 and 10, depending on the reduction code and instrument. In the original pyodine code, an oversampling of 10 is used




FIGURE 3.17: One order of the simulated template (top), of the simulated observation using a `SingleGaussian` LSF and injected RV of 0 m s^{-1} (middle), and of the real observation used to generate the wavelength scale of the simulated spectra (bottom).

for the template creation, and I adopted this value. Thus, the wavelengths of the SONG observation were interpolated to an over-sampled pixel grid of 20,480 pixels per order. For the flux of the template orders, a continuous array of ones and of the same size as the wavelength array was created; then Gaussian profiles were subtracted from it to simulate stellar absorption lines. The number of Gaussians per order, and the pixel position, amplitude and standard deviation of each Gaussian were all drawn from normal or uniform distributions. Next, each order flux vector was chunked into 22 `TemplateChunk` objects, each covering 90 normally sampled pixels, and all were saved in a `StellarTemplate_Chunked` object. Figure 3.17 shows one order of the simulated template before chunking in the top panel and the corresponding order of the real observation used for the wavelength scale in the bottom panel.

Next, a number of simulated observation spectra were created to test the code with, using Equation 3.12 to produce each order of an observation: The wavelengths of the simulated, un-chunked template were shifted by a Doppler factor $1 + v_{\text{RV}}/c$, where v_{RV} is the injected RV. This shifted template spectrum was then multiplied with the corresponding part of the SONG I2 FTS, whose depth (i.e. strength of the I2 lines in the resulting spectrum) was before-hand set to 0.98 in all cases. The resulting spectrum was convolved with a LSF model, and finally multiplied by a parabola imitating the blaze function of the spectrograph and increasing the continuum. A total of 15 observations were created this way, using the three LSF models `SingleGaussian`, `MultiGaussian` and `MultiGaussian_Lick`, and five different injected RVs (-2000 , -1000 , 0 , 1000 and 2000 m s^{-1}) each. Table 3.3 gives an overview of the properties of the simulated observations. Most importantly, the RV analysis of these spectra should be able to return a good estimate of the originally injected RVs; however, the other parameters used in the creation of the simulated data, namely of the LSF models and the depths of the I2 and stellar template spectra, should also at least roughly




FIGURE 3.18: One chunk of the simulated modelling, with the SingleGaussian LSF and an injected RV of 2000 m s^{-1} . From top to bottom: The template (blue) and Doppler-shifted template (yellow), the I2 spectrum, the observation (blue), model (yellow) and fitted continuum (green), and the residuals between the observation and model.

be recovered.

All simulated observations of a given LSF model were then modelled with the main routine outlined in Section 3.4.8, using the same LSF as in their creation. A chunk size of 90 pixels was chosen, which roughly corresponds to 2 \AA for the SONG wavelength scale (which was used as basis for the simulated observations). Each order thus consisted of 22 chunks, and each observation of $23 \times 22 = 504$ chunks in total.

The Single-Gaussian LSF analysis

In a first test, the SingleGaussian LSF model was used to create the simulated observations, with the only variable parameter $FWHM = 2.4 \text{ pix}$, and the same LSF model was used for the analysis. Figure 3.18 shows the observation and model flux for an exemplary chunk (from the observation with injected velocity $v_{\text{inj}} = 2000 \text{ m s}^{-1}$). The best-fit velocity parameter for this chunk is 1995 m s^{-1} (see top part of the figure), which is very close to the injected velocity, and the rms of the residuals between the observation and model is 0.467% . This roughly corresponds to the average residual rms of all modelled chunks.

A histogram of best-fit chunk velocities for the same observation is shown in Figure 3.19, left; the velocities are roughly normally distributed, with a standard deviation of 27.7 m s^{-1} . The simple median of all chunk velocities is 1998.1 m s^{-1} , which

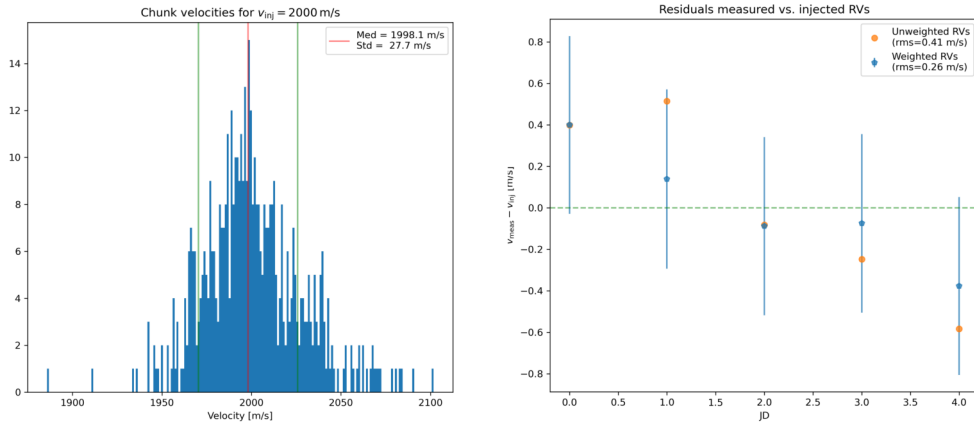


FIGURE 3.19: Results from simulated observations, with the `SingleGaussian` LSF model. *Left*: Distribution of chunk velocity results for the observation with injected velocity $v_{\text{inj}} = 2000 \text{ m s}^{-1}$. *Right*: Residuals between the measured RVs from the combination algorithm (blue: weighted RVs, orange: unweighted RVs) and the injected velocities of the simulated observations.

is already quite close to the injected velocity. Also for the other observations, the median velocities always fall within 2 m s^{-1} of the respective injected velocities, and the standard deviations of the observation chunk velocities are all between 26.8 and 27.8 m s^{-1} .

A weighted timeseries of the five simulated observations was created using the `SONG`-based combination algorithm. The right-hand side of Figure 3.19 shows the residuals between the weighted measured velocities and injected velocities in blue, and the unweighted measured velocities and injected velocities in orange. The rms values of the residuals for the two timeseries are 0.26 and 0.41 m s^{-1} , respectively; the weighted timeseries thus performs only slightly better than the unweighted one (although considerably better than the simple median), which is understandable given the comparably well-behaved distribution of velocities in these tests of simulated spectra. The derived uncertainties of the weighted velocities are of the order of 0.4 m s^{-1} , and in all observations they are large enough to incorporate the injected velocity. The overall downward trend of the measured velocity timeseries from left to right could be indicative of a systematic behaviour, but given the small overall variation of less than 0.8 m s^{-1} it could also just be a statistical artefact.

Finally, we can also compare the fitting results of the other free parameters in the model with the injected values: The average LSF *FWHM* recovered from all chunks and all observations is 2.39 pix , which is very close to the input value of 2.4 . The average model results for the depths of the I2 and stellar template spectra are 0.98 and 1.01 , respectively, and thus agree nearly perfectly with the respective injected values.

The Multi-Gaussian LSF analysis

To test the performance of the analysis code with spectra that are modulated by an asymmetric LSF, simulated observations were created using the `MultiGaussian` and `MultiGaussian_Lick` LSF models, with the ten free parameters (i.e. the amplitudes




FIGURE 3.20: Results from simulated observations, with the `MultiGaussian_Lick` and `MultiGaussian` LSF models on the left and right, respectively. Residuals between the measured weighted velocities from the combination algorithm and the injected velocities are plotted in blue, and the residuals for the unweighted velocities in orange.

of the satellite Gaussians) chosen as specified in Table 3.3. All other injected parameters are the same as in the previous section. For these tests, the recentering algorithms were included in the LSF models.

The simulated observations were then modelled with the same LSF models as used in the creation. As mentioned previously, the fitting of each observation was performed in two runs, first using the `SingleGaussian` LSF to approximate the wavelength, continuum and spectral depth parameters, then with the actual desired LSF model, using the non-LSF parameter results from the previous run as starting values. To get a good first guess for the LSF parameters of the `MultiGaussian` or `MultiGaussian_Lick` model, the desired model was fitted to the `SingleGaussian` LSF, median-averaged over all chunks, from the first run, using the `fit_lsfs()` method (see Section 3.4.3). Furthermore, by running the models for different parameter inputs and evaluating the resulting velocity results, it was found that by taking advantage of the capabilities of the `lmfit` parameter class and setting fitting limits to the LSF parameters the results can be improved. After a few iterations, best performance was achieved by limiting each LSF parameter (i.e. the amplitudes of the satellite Gaussians) to ± 0.3 around the starting value found by the fit to the `SingleGaussian` LSF described above. This approach reduces the number of outlier chunk results, where the modelling results in highly distorted LSFs leading to chunk velocities that greatly deviate from the observation median. After the fitting, the final model results of all observations were saved, and the velocity timeseries were evaluated using the original SONG-based weighting algorithm.

Figure 3.20 shows the residuals between the resulting measured timeseries and the injected velocities for the `MultiGaussian_Lick` and the `MultiGaussian` LSF models on the left and right, respectively, and both for the weighted and unweighted timeseries (blue and orange, respectively). The rms values of the residuals are 0.26 and 0.33 m s^{-1} , which gives a slight advantage to the `MultiGaussian_Lick` model. Overall the injected velocities however are recovered quite well, and the errors are large enough to incorporate the "true" values for all resulting weighted velocities except the last one in the `MultiGaussian_Lick` series.

3.5.2 Early results from Lick observations

The archive of observations from our survey of G-/K-giant stars, obtained with the Hamilton spectrograph at Lick Observatory, allows testing the creation of stellar template spectra, the modelling of the observations and the combination of resulting chunk velocities to RV timeseries. These results can then be compared to the RV results from the original Butler code, and the overall performance of the I2 analysis package of the Waltz DRS can thus be assessed. Already in the development of the code, Lick spectra were used to test modifications and upgrades to the original pyodine package and evaluate their results.

The Lick spectra in our archive are already extracted to 1D spectra for each order, but no wavelength solution is contained in the spectrum files. However, for each observation night we also possess extracted calibration spectra of a ThAr lamp and of a pure I2 spectrum, obtained by shooting light of a continuum source through the I2 cell. As the modified pyodine package in the Waltz DRS requires a relatively good approximation of the wavelength solution for each order when modelling spectra, I wavelength-calibrated all Lick spectra that I used in my tests with the ThAr spectra of the respective observation nights. The calibration pipeline is based on modules from the modified CERES package of the Waltz DRS and basically follows the process outlined in Section 3.3.2. It uses a ThAr line atlas that I created by applying ThAr atlases from other instruments in the original CERES package to the spectral format of the Hamilton spectrograph at Lick. Some of the Lick ThAr spectra are of bad quality, only showing few lines and being very noisy, which could lead to incorrect wavelength calibrations. In these cases the pipeline falls back on the wavelength solution from the most recent observation in time. The spectral format of the Hamilton spectrograph usually varied by less than a pixel in the dispersion direction over long times, and wavelength scale errors of that extent are small enough for the I2 reduction code to handle. Only few observations close to the end of our survey had to be dropped, as a correct wavelength determination for these spectra would have required a more detailed approach.

For the barycentric velocities and weighted mid-point times of the observations, the pipeline falls back to the file "bcvel.ascii" from our Lick archive, which contains information from the original Lick barycentric correction algorithm for all obtained spectra. This way, differences between the original RVs from the Butler code and the results from the Waltz DRS will only be due to differences in the I2 analysis.

For early tests of the Waltz DRS, the most useful observations were those of a star with little to no inherent systematic RV signatures, as the performance of the reduction code could then be measured simply by the rms of the resulting RV timeseries, with a lower rms indicating improvements of the analysis. In the Lick sample, the star HIP 96459 thus qualified as a good test object: Over a total of 38 observations its RV rms is 9.33 m s^{-1} as analyzed by the original Butler code, with the uncertainties of the individual RVs mostly lying between roughly 5 and 7 m s^{-1} . No clear periodic modulations are present in the RVs, so the RV scatter is mostly caused by astrophysical jitter, such as p-mode oscillations of the star. I therefore chose the observations of that star for tests during the development of the I2 analysis code, and for determining good parameter combinations.

Modelling procedures and parameters

In an iterative process, by modelling the spectra of the star HIP 96459 and evaluating the results, I determined a combination of model parameters and procedures to improve the performance of the code. However, some information about the required parameters exists in the literature, or can be found in the dop code. The setup of the `MultiGaussian_Lick` LSF model, specifically the positions and widths of the individual Gaussian functions of the model, are taken from the dop code: The satellite Gaussians have widths of 0.3 pixels, and their positions are 0.6, 1.1, 1.6, 2.1 and 2.4 pixels to the left and right of the central Gaussian, which has a width of 0.4 pixels and a fixed amplitude of 1. This description slightly differs from the one used by Butler et al. (1996), who positioned the satellite Gaussians at half-pixel steps from the center. However, in IDL session-files stemming from the RV analysis of spectra from our Lick survey, the same information as in the dop code is stored. Therefore, I used that LSF setup.

The Lick stellar templates, which are also stored as IDL-session files, consist of 704 spectrum chunks, each of which contains not only the template flux values and wavelength information, but also the order and pixel position of the chunk within the template observation. These templates cover the order indices 38 to 53 of the observation spectra, corresponding to wavelengths between roughly 5020 and 5870 Å. Each of the 16 orders is thus split into 44 chunks. From Butler et al. (1996), the dop code and the information in the Lick template files, we know that each chunk has a width of 40 pixels, and the chunk padding used in the template creation is 12 pixels at either side of the chunks. In the template creation with the Waltz DRS, these parameters were adopted to create the observation chunks with the `simple()` chunking algorithm. When modelling the observations, the `wave_defined()` algorithm was used to position the chunks such that they correspond to the template chunks regarding their spectral content.

As done in the Butler code, I used an oversampling factor of 4 when modelling the chunks. In the template creation process, a bad-pixel mask of the template observation was computed first to find flux artefacts, and corresponding pixel weights were set to zero in the modelling of the respective O-/B-star spectra. The flux values of these affected pixels in the template observation were corrected by linear interpolation as described in Section 3.4.5. For the modelling of I2-observations of the stars, I abstained from computing bad-pixel masks as early tests resulted in too many masked pixels for the later acquired spectra of each timeseries, probably because these were generally obtained with a different detector. In the future, this could be improved by using a varying threshold parameter in the computation of the bad-pixel mask, which depends on the detector used. For the tests presented in the following sections, modelling without a mask worked fine. To correct for tellurics, the CARMENES telluric mask was loaded, and the pixel weights in affected wavelength regions set to zero during the fitting (both in the template creation process and modelling of I2-observations). All other pixel weights were set to 1.

All spectra were modelled in at least two runs, first with the `SingleGaussian` LSF model, then with the `MultiGaussian_Lick` model (unless specified differently). Surprisingly, best results were achieved when using the LSF model without the recentering algorithm. While I could not find published information about whether the LSF used in the Butler code was recentered or not, in the iSONG analysis code of the SONG project the LSF model does employ a recentering algorithm (as described in Section 3.4.2, private communication with F. Grundahl); this is meant to prevent the

fitter from balancing out asymmetric LSFs with shifts of the wavelength intercepts, which would lead to changes of the modelled velocities. In my tests with Lick spectra of HIP 96459 however, RV timeseries created with a recentered LSF resulted in a considerably larger rms. This might be explained by the poorly constrained LSF, caused by the comparably large dispersion of the Lick spectra. However, a deeper analysis of this behaviour will definitely be required.

Prior to the first modelling run, after guessing the starting values of the model parameters, the wavelength-slope parameter guesses p_{slope} were smoothed over each order with a 3rd-degree polynomial fit. Furthermore, the FWHM parameter of the LSF was limited to values between 0.5 and 4.0, using the `constrain_parameters()` method of the input parameter object. After the first modelling run, the best-fit results of the wavelength parameters were modelled with 3rd-degree polynomials over each order. The evaluated polynomial values of the wavelength slope and intercept for each chunk served as starting guesses for the second modelling run; this way, the second modelling run is not affected greatly by outlier model results from the first run. For the continuum parameters, and the I2 FTS and stellar template depth parameters, the results from the first run were adopted as starting values in the second. The guess for the LSF parameters finally was found by modelling the first-run best-fit LSF, median-averaged over all chunks, with the LSF model for the second run. As done in the tests with simulated spectra, the LSF parameters were limited to ± 0.3 around the starting value, in order to prevent the fitter to produce too strong outliers.

After the second modelling run, the results were saved. When creating template observations, the same model was fitted again in a third run, but the wavelength parameters were limited to a narrow region around a 3rd-degree polynomial model of the results from the second run, to ensure that the final template would have a smooth wavelength scale over each order. The results from the third modelling run were used to deconvolve the stellar template observation, which was then stored in a `StellarTemplate_Chunked` object and stored to the disk. An overview of all important parameters, including the ones used in the deconvolution algorithm, can be found in Table A.2.

Stellar template creation for HIP 96459

In the creation of the deconvolved stellar template for HIP 96459, four observation spectra of the A-star δ Cyg (HR 7528), obtained with the I2 cell in the light path before and after the stellar template observation, were summed up to receive a high S/N spectrum, and the summed spectrum was modelled to determine the LSF of the instrument. Figure 3.23 displays some results from the second modelling run, where the uncentered `MultiGaussian_Lick` LSF model was used. On the left-hand side, the observation and model for chunk 300 are shown, with the I2 spectrum in the top, the observation and model spectrum (blue and orange, respectively) and the fitted continuum (green) in the middle, and the residuals between observation and model, relative to the observation flux, in the bottom. The rms of the residuals in this chunk is 0.35%, which shows that the spectral features are mostly very well fitted.

On the right-hand side of Figure 3.21, the resulting residuals of all chunks are shown, with a histogram in the top and a scatter plot over the chunk indices in the bottom; the green dotted lines in the scatter plot denote the order edges. The median value

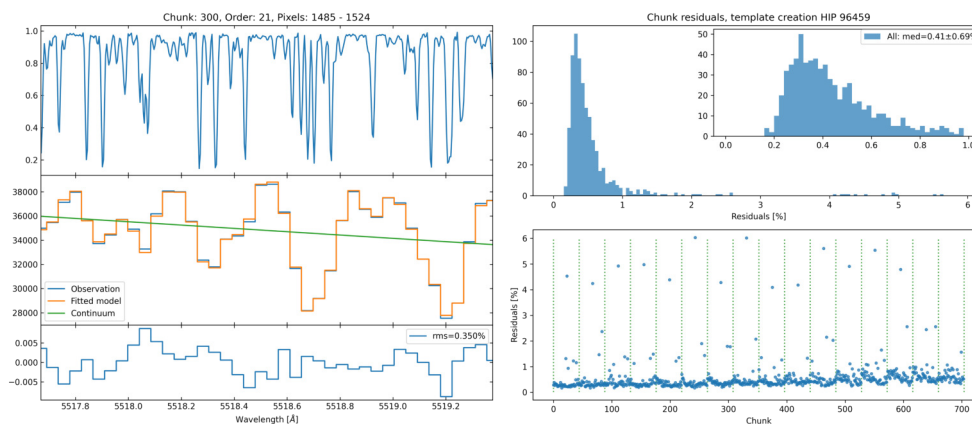


FIGURE 3.21: Modelling results of O-star spectra for the template creation of HIP 96459. *Left*: Model result of one chunk. *Right*: Histogram (with zoomed in subframe) and scatter plot of chunk residuals.

of the residuals is 0.41%, and most of the chunk results fall between roughly 0.2 and 0.6%. Within each order except the last one, a chunk in the middle between the order edges has quite high residuals of several percentage-points; these chunks are affected by a detector pixel column with very low counts, that runs through most of the extracted orders of the Lick spectra. While these pixels worsen the residuals of the respective chunks, they have been identified by the bad-pixel mask and their weights have been set to zero, so the model should not be affected. In the stellar template observation, the flux of these pixels will be interpolated over to produce a cleaner deconvolved template.

From the right-bottom panel in Figure 3.21, it is also obvious that the scatter and overall magnitude of the residuals increases towards higher chunk numbers, which corresponds to bluer wavelengths. This is at least partly due to the fact that the I2 spectral features decrease in strength at wavelengths a little over 5000 Å, and the relative noise of the spectrum thus grows stronger in these orders. Also, the bad-pixel mask contains a few more entries of affected pixels in the bluer orders as compared to the redder ones, which contributes to some chunks having larger residuals.

In Table 3.4, the median results for the third-run LSF parameters from all chunks, with the standard deviation as uncertainty, are summarized, with "lsf_left_5" being the amplitude of the satellite Gaussian at the most negative pixel position, and "lsf_right5" of the satellite Gaussian at most positive pixel position. Figure 3.22 displays the median and standard deviation of the resulting LSFs of all chunks. The median LSF is overall quite symmetric, with the right wing only being a little higher than the left one. To the extreme left and right sides, small bumps are visible before the LSF goes to zero, which are caused by the 4th parameter being negative on both sides. These bumps seem somehow unphysical, and they probably appear because the LSF in the Lick spectra is not well-sampled, as the Hamilton spectrograph has a dispersion of roughly 0.04 to 0.05 Å pix⁻¹; the dispersion of the Waltz spectrograph in contrast is roughly 0.034 Å pix⁻¹, and SONG even reaches ~ 0.02 Å pix⁻¹. Due to the badly sampled LSF and the large number of free parameters when using the MultiGaussian LSF, the fitter has problems to constrain the parameters well. The bumps could potentially be prevented by limiting all LSF parameters to always

TABLE 3.4: Template creation for HIP 96459, median LSF parameter results of all chunks

Parameter	Value
lsf_left_5	0.142 ± 0.114
lsf_left_4	-0.068 ± 0.162
lsf_left_3	0.271 ± 0.176
lsf_left_2	0.246 ± 0.138
lsf_left_1	0.549 ± 0.200
lsf_right1	0.578 ± 0.202
lsf_right2	0.268 ± 0.136
lsf_right3	0.293 ± 0.177
lsf_right4	-0.073 ± 0.162
lsf_right5	0.148 ± 0.117




FIGURE 3.22: Median evaluated LSF of all chunks in the template creation of HIP 96459.

remain positive. However, the LSF shown in Figure 3.22 is evaluated on the over-sampled pixel grid; when binned to the actual pixels of the observed spectrum, the bumps are smoothed out and act as slightly widened wings as compared to a normal Gaussian function.

The LSF parameter results from the third modelling run were used to deconvolve the stellar template observation of HIP 96459. Figure 3.23 shows four exemplary chunks of the deconvolved stellar template, with the results from the Waltz DRS plotted in blue; for comparison, the chunk flux values of the Lick stellar template which correspond best in wavelength are overplotted with gray dotted lines. Both the Waltz DRS and Lick chunks are shown over their full width including the padding of 12 pixels on the left and right. For the Waltz DRS template chunks, edge-effects caused by the deconvolution are visible, with the flux rising quickly to the continuum value over the last $\sim 0.2 \text{ \AA}$ at both sides; the Lick template chunks in contrast show good spectrum flux values over their complete width. Nevertheless, the edge-effects in the Waltz DRS chunks should not be a problem when modelling observations: With the `wave_defined()` chunking algorithm, the Doppler shifts caused by the relative barycentric velocities between template and observations will be compensated for, and the only remaining spectral shifts will be the ones caused by the RV changes of the star. With the edge-effects beginning no sooner than 0.3 \AA into the padding regions to both sides of the template chunks, corresponding to roughly 16 km s^{-1} in velocity space, all observations with RV changes smaller than this value will be unaffected.

Another interesting difference between the Waltz DRS and Lick template chunks is the slight mismatch of the wavelength scales, with the spectral features in the Lick chunks being roughly 0.2 \AA shifted towards longer wavelengths. In this context it is important to mention that I computed the plotted wavelengths of the Lick chunks on the basis of the wavelength intercept and slope parameters contained in the Lick template IDL-files. As wavelength intercept, denoted as "pixel 0", I assumed the central pixel of each chunk, as is done in the `pyodine` package in the Waltz DRS. It is possible however that the Butler and/or `dop` code use a different definition; unfortunately, I was unable to find the relevant passage in the `dop` code, and therefore I had to rely on the guess explained above. However, the Lick template files also




FIGURE 3.23: Four exemplary deconvolved stellar template chunks of HIP 96459, from the original Lick template (grey dotted) and from the result of the Waltz DRS (blue).

contain information about the absolute starting pixels of the chunks within their respective orders. These differ constantly by 5 pixels from the starting pixels of the Waltz DRS chunks, so all Lick chunks are simply shifted by that amount with respect to the Waltz DRS chunks. This number corresponds to roughly 0.2 to 0.3 Å in wavelength space, and could at least partly explain the wavelength shifts observed in Figure 3.23.

To better understand the mismatch, I compared the wavelength parameters (intercept and slope) derived from the Waltz DRS with the corresponding values from the Lick template file. Figure 3.24 displays the residuals between the Waltz DRS and Lick values for the wavelength intercept (top) and slope (bottom) for all chunks. The wavelength intercepts show clear systematic residuals, which vary between roughly 0.3 and 0.45 Å in the redder orders (lower chunk numbers), and 0.27 and 0.4 Å in the bluer orders (higher chunk numbers). These residuals are too large to be only explained by the 5-pixel shift between the Lick and Waltz DRS chunks, which points to the presumption of differently defined pixels for the wavelength intercepts of the chunks. The change of wavelength intercepts within the orders (and, less prominent, from order to order) simply reflect the changing linear dispersion of the spectrograph. In the worst case, the wavelength offsets however will only lead to an overall systematic offset of the RVs derived by the Waltz DRS when modelling observations.

The residuals between the Waltz DRS and Lick values for the wavelength slope in contrast scatter around roughly 0 with an rms of $\sim 10^{-4} \text{Å pix}^{-1}$, and also for the residual wavelength intercepts a small-scale scatter around the systematic trends described above can be observed. The scatter mostly arises due to a variation of the Waltz DRS modelled wavelength parameters around the 3rd-degree polynomials fitted prior to the second modelling run, as the wavelength parameter limits adopted in that run had been chosen wide enough to allow some variation. For the




FIGURE 3.24: Differences of wavelength intercept (top) and wavelength slope (bottom) parameters between the Lick and Waltz DRS templates, for all 704 template chunks.

wavelength intercepts, that scatter is of order 10^{-3}\AA , which corresponds to roughly 50 m s^{-1} in velocity space. The wavelength intercept scatter of the template therefore might lead to a variation of fitted chunk velocities of roughly that magnitude in the modelling of I2 observations. For comparison, I also modelled the wavelength parameters of the Lick template spectrum with polynomials of the same degree over each order, and barely found any scatter around the fits. It therefore seems that the wavelength parameters of the Lick template spectra might have been smoothed after the last modelling run. However, when smoothing the wavelength parameters of the Waltz DRS template and using it to model a series of I2 observations of HIP 96459, no improvements in the achieved rms of the RV timeseries could be found. Therefore, for the results presented in the following sections stellar templates without smoothed wavelength parameters have been used.

Modelling of I2 observations of HIP 96459

With the template which was created as described in the previous section I modelled 38 observations of the star HIP 96459, obtained with I2 cell in the light path, using the procedures and parameters described in Section 3.5.2. The individual chunk velocities were then combined to a RV timeseries with the SONG-based velocity weighting algorithm. For two observations, the combination algorithm produced large outliers falling more than 7000 m s^{-1} away from all other RVs. These observations are not included in the RV timeseries from the original Lick analysis, so they are possibly compromised, for instance by a wrong computation of barycentric velocities. I added both observations to the rejection mask, and computed the RVs again without them.

The final timeseries consists of 36 RVs, two less than the original Lick RVs, because the last two observations had been rejected in the wavelength calibration pipeline




FIGURE 3.25: *Top*: Timeseries of RVs for HIP 96459, from the Waltz DRS with the SONG-based velocity weighting algorithm (blue), along with the original Lick results (orange). *Bottom*: Residuals between the two timeseries; errors are propagated from individual uncertainties of both timeseries.

due to erroneous wavelength results. Figure 3.25, top, displays the resulting timeseries of the RVs in blue, and the original Lick RVs are shown in orange for comparison. Both have been centered around 0 by subtracting their overall median. The overall rms of the timeseries is 9.3 m s^{-1} for the original Lick results, and 10.4 m s^{-1} for the results from the Waltz DRS. The median scatter of the RVs thus is comparable. For many observations, the computed RVs from both reduction pipelines fall within each others' computed uncertainties, and for all but 5 observations the error-bars overlap.

In the bottom panel of Figure 3.25, residuals between the two timeseries are plotted, with errors propagated from the individual measurement uncertainties. The rms of the residuals between both timeseries is 8.2 m s^{-1} , and the errors of most residuals overlap with 0, thus indicating good agreement within the individual measurement uncertainties for these RVs. There are however some observations with residuals of 10 m s^{-1} and more, which might point to differences in the reduction pipelines that are not reflected by the computed uncertainties. Generally, the uncertainties computed with the SONG-based weighting algorithm are a little larger than the original Lick RV errors, with a median of 8.1 m s^{-1} as compared to 5.9 m s^{-1} . This could reflect a larger scatter of chunk velocities for each observation than achieved with the Butler code, but might also be caused by the different velocity weighting algorithm used.

Figure 3.26, left, shows the velocity scatter for each chunk timeseries around the observation medians (σ_j , Equation 3.33), which on average is $238.7 \pm 234.2 \text{ m s}^{-1}$. While most chunk series fall well below 300 m s^{-1} , some are very strong outliers. A deeper




FIGURE 3.26: Analysis results of RVs of HIP 96459, computed by the SONG-based weighting algorithm. *Left*: Velocity scatter of all chunk series around observation medians, as computed by Equation 3.33. *Right*: Robust standard deviation of chunk velocity scatters within the observations.

analysis into these bad chunk series showed that mostly their chunk velocities are not too far from the observation medians, but for a few observations the velocities deviate very strongly from all other chunk results, which inflates their σ . In the right panel of Figure 3.26, the observation-to-observation scatter of all chunk velocities is plotted (using the robust standard deviation to disregard outliers). The median scatter is roughly 300 m s^{-1} , which is quite large. I do not have information about what scatter was reached by the Butler code, but in the SONG project (using the iSONG data reduction pipeline) for most stars chunk velocity scatters around 100 m s^{-1} are achieved (F. Grundahl, private communication); this is however for spectra from the SONG spectrograph, which are of higher quality than the Lick spectra, and using the Waltz DRS on this data I achieve similar results as iSONG (see Section 3.5.3). It is therefore well possible that the large velocity scatter observed in the analysis of Lick spectra is at least to some part a consequence of the data, and not of the reduction code.

To better understand the influence of the velocity weighting algorithm on the resulting timeseries, I computed the RVs also with the dop-based algorithm. The absolute limits for chunk rejection, $\chi_{\text{red,max}}^2$, \bar{f}_{min} and \bar{f}_{max} , were set so wide that no chunks or observations were excluded, and the percentile limits for the sigma clipping were set to 0.997, corresponding to 3σ . This timeseries results in an overall RV scatter of 10.2 m s^{-1} , and the rms of the residuals to the original Lick RVs is 7.5 m s^{-1} ; both numbers are a little smaller than for the timeseries computed with the SONG-based algorithm, but the shape of the overall timeseries shows no substantial differences. The median of the RV uncertainties from the dop-based weighting algorithm however is 5.7 m s^{-1} , which is considerably smaller than for the SONG-based algorithm and even beats the median of the original Lick RV errors by 0.2 m s^{-1} . This could be a hint that the chunk velocity scatter of the observations achieved with the Waltz DRS is actually similar to the results from the Butler code, and the larger RV errors observed in the timeseries in Figure 3.25 are a consequence of using the SONG-based weighting algorithm instead.




FIGURE 3.27: RVs and Keplerian model results for the star Pollux (HIP 37826): The top two panels show the original Lick RVs and the RVs computed with the Waltz DRS in blue, respectively, with their respective best-fit Keplerian models as black lines; in the bottom two panels, residuals from the RVs to their respective Keplerian fits are displayed.

RV results for Pollux

Based on RV measurements collected as part of the G-/K-giant survey, Reffert et al. (2006) published strong evidence for a planetary companion to the K-giant star Pollux (HIP 37826), with an orbital period around 590 d, an eccentricity of 0.06 and a minimum mass of roughly $2.9 M_{\text{Jup}}$. Using the modified pyodine package of the Waltz DRS, I re-analyzed the Lick spectra of Pollux and computed an RV timeseries from the resulting chunk velocities, to test whether the precision reached by the code would allow identification of that planetary companion with similar parameters. For this run, I used the same settings and procedures as found suitable for HIP 96459 in the previous section.

The RV timeseries was again created with the SONG-based weighting algorithm, and resulted in a total of 94 RVs. The original Lick timeseries in contrast consists of 121 RVs; the discrepancy in the total number of RVs of the two timeseries is again due to problems with the automated wavelength calibration of the Lick spectra when preparing them for the modelling with the Waltz DRS, caused by problematic ThAr spectra and changes in the spectral format. The number of Lick RVs used today is also considerably larger than used in Reffert et al. (2006), due to additional observations having been obtained after publication.

Therefore, to allow a thorough comparison of the RVs computed with the Waltz DRS

TABLE 3.5: Keplerian parameters for the planetary companion to Pollux, from the Waltz DRS and Lick RVs

Parameter	Waltz DRS RVs	Lick RVs	
		Full dataset	Shortened dataset
P [days]	597.3	597.4	595.5
M_0 [deg] ^a	238.9	179.4	206.7
e	0.08	0.05	0.05
ω [deg]	44.9	77.9	106.2
K [m s ⁻¹]	46.0	46.4	46.0
$m \sin i$ [M _{jup}] ^b	2.88	2.88	2.91
a [AU]	1.71	1.71	1.70

^a The mean anomalies are calculated at the observational epoch $t_0 = 2451750.0$ JD.

^b The mass has been derived by solving the mass function for fixed inclinations i of 90° and with the primary mass $M_* = 1.86 M_\odot$.

and the original Lick RVs, I fitted a single-Keplerian model to each timeseries, using the companion parameters from Reffert et al. (2006) as starting values. Figure 3.27 displays the Lick and Waltz DRS timeseries in the top two panels, along with their respective best-fit Keplerian models, and the RV residuals to the models are shown in the bottom two panels for both timeseries. It is obvious that the periodic variations observed in the Lick RVs are generally well met by the results from the Waltz DRS. The rms of the Lick RV residuals around their best-fit Keplerian model is 11.0 m s^{-1} , which is a little smaller than the 12.1 m s^{-1} rms of the Waltz DRS RV residuals around their best fit — despite the Lick timeseries consisting of 27 more RV measurements. I performed the Keplerian fit on a shortened Lick timeseries again, this time only using RVs of observations that exist in both timeseries, which resulted in a smaller RV residual rms of 10.4 m s^{-1} for the Lick timeseries.

The χ_{red}^2 of the fit to the Waltz DRS RVs are smaller than for the Keplerian models to the Lick RVs; this is however due to the larger measurement uncertainties delivered by the Waltz DRS, which have a median value of 6.3 m s^{-1} , as compared to a median error of 4.0 m s^{-1} for the Lick RVs. The larger RV uncertainties returned by the SONG-based weighting algorithm were already discussed in the previous section; when I instead combined the chunk velocities for Pollux computed by the Waltz DRS with the dop-based weighting algorithm, it returned a much smaller median RV error of 4.3 m s^{-1} , which is closer to the median Lick RV error. The absolute values of the dop-weighted RVs change only slightly as compared to the SONG-weighted RVs, and a Keplerian fit to the dop-weighted timeseries resulted in a residual rms of 12.2 m s^{-1} , only slightly worse than the 12.1 m s^{-1} rms of the SONG-weighted residuals. The χ_{red}^2 of the fit to the dop-weighted RVs is 8.2, while the Keplerian model to the shortened Lick timeseries returned a χ_{red}^2 of 7.4, despite the still smaller RV errors.

Table 3.5 summarizes the Keplerian parameter results from the best-fit models to the SONG-weighted Waltz DRS RVs, the full and the shortened Lick RV timeseries. As can be seen, the results are mostly quite similar, and the largest differences appear for the mean anomalies M_0 and the longitudes of periastron ω ; due to the small eccentricity of the orbit however, the longitude of periastron is not well constrained, which also affects the mean anomaly. Both the orbital period P and the semi-amplitude K




FIGURE 3.28: Generalized Lomb-Scargle (GLS) periodograms of the RV timeseries of Pollux. The companion orbital period as derived by Reffert et al. (2006) is indicated by the green dotted line. *Top*: GLS of the original Lick RVs (orange dotted) and the results from the Waltz DRS (blue). *Bottom*: GLS of the RV residuals from Lick and the Waltz DRS to their respective best-fit Keplerian models.

however show very good agreement particularly between the models to the Waltz DRS RVs and the Lick RVs. Interestingly, the model to the shortened Lick timeseries, consisting of the same observations as the Waltz DRS timeseries, shows the strongest differences, but even these values do not fall far from the other results. Generally, the comparison between the fitted Keplerian models show that using the Waltz DRS RVs instead of the Lick RVs only results in minor changes of the model, and the general physical interpretation of the data remains untouched.

To test whether any periodic systematics caused by the code itself appear in the RV timeseries from the Waltz DRS, I computed Generalized Lomb-Scargle (GLS) periodograms (Zechmeister and Kürster, 2009) of both the Lick and Waltz DRS RVs of Pollux, and compared the found periodicities. Figure 3.28 displays the power distribution over a wide range of periods, for the SONG-weighted RV timeseries from the Waltz DRS (blue) and Lick (orange dotted) in the top panel; the bottom panel shows the periodogram for the residuals of the RVs to their respective best-fit Keplerian models. Generally, the periodograms show good agreement between the Waltz DRS and Lick results. The highest power is found for periods of 596 and 593 d for the Lick and Waltz DRS RVs, respectively, which is both quite close to the 590 d orbital period of the planetary companion published by Reffert et al. (2006). In the periodograms of the residuals, some peaks with powers below 0.25 remain, most of which are already visible in the periodograms of the RVs and not greatly enhanced after subtraction of the Keplerian models. In some cases, the heights of the peaks show some variation between the Lick and Waltz DRS residuals, but their periods agree quite well. Overall, it seems that the Waltz DRS does not add any systematics to the RVs.

It can be concluded that the Waltz DRS reaches a similar precision as the original Lick

data reduction, and its RVs allow a clear identification of the previously published planetary companion to Pollux. The Keplerian parameters as returned by the fit to the SONG-weighted RV timeseries from the Waltz DRS only differ slightly from the best-fit parameters for the Lick timeseries. Furthermore, the Waltz DRS RVs seem to be free of any systematic artefacts added by the code itself, at least when compared to the Lick RVs.

RV results for HIP 36616

Another published K-giant system from the Lick survey is HIP 36616 (HD 59686), which consists of an eccentric stellar binary system and an S-type planet of $\sim 7 M_{\text{Jup}}$, orbiting the main component with a roughly 300 d period (Ortiz et al., 2016). The orbit of the stellar companion ($M_* \approx 0.5 M_{\odot}$) has a period of 32 yr and eccentricity of 0.73, which leads to a periastron distance of only 3.6 AU, and thus greatly varying gravitational forces are exerted on the planetary companion. In a dynamical analysis, using the RVs from the Lick survey, Trifonov et al. (2018) identified islands of stable orbital configurations of the system in the parameter space.

Due to the high mass and eccentricity of the stellar companion HIP 36616 B, the measured RVs from the primary component undergo quick and high-amplitude variations on the order of $\pm 4 \text{ km s}^{-1}$ around periastron passage. I therefore chose the Lick observations of that star to test the reliability and precision of the Waltz DRS when analyzing heavily RV-modulated spectra. The comparably small-scale signal of roughly 140 m s^{-1} induced by the planetary companion can serve as a testbed, by checking whether it is recovered accurately.

For modelling the observations, I again used the same settings as for the previously discussed stars, and the final RV timeseries was again computed with the SONG-based weighting algorithm. The modelling resulted in a timeseries with 75 RVs, while the full Lick timeseries consists of 88 RVs; again, for a number of observations acquired later in the survey the wavelength calibration pipeline failed to deliver good results — on the other hand, three earlier observations with RVs contained in the Waltz DRS timeseries are missing in the Lick timeseries. The shortened timeseries therefore consist of 72 RV measurements each for the Waltz DRS and Lick.

As in the previous section, I modelled both the Lick and Waltz DRS timeseries with a Keplerian model; however, this time I included two companions, and chose the results for the Keplerian parameters from Ortiz et al. (2016) as starting values. To allow better comparison of the models, I did the same fits for the shortened timeseries of the Waltz DRS and Lick, and the results from these fits are shown in Figure 3.29, with the Lick and Waltz DRS RVs and corresponding Keplerian models in the top two panels, and the respective residuals in the bottom two panels. Generally, the curves look very similar, and the high-amplitude variation caused by the stellar companion is well recovered by the Waltz DRS RVs. Also the modulation induced by the planetary companion is clearly visible in the Waltz DRS timeseries. The Keplerian parameters of the different best-fit models to the timeseries only show slight differences on the order of 1% and less (see Table 3.6), which indicates a very good agreement of the models.

However, the residuals of the Waltz DRS RVs show a smaller rms of 14.1 m s^{-1} as compared to a value of 19.0 m s^{-1} for the Lick RVs, using the shortened timeseries. The difference here can be mostly attributed to a few data points around JD 2 454 500, that fall far from the best-fit curve in the Lick timeseries and nearly seem to follow




FIGURE 3.29: RVs and Keplerian model results for the star HIP 36616: The top two panels show the original Lick RVs and the RVs computed with the Waltz DRS in blue, respectively, with their respective best-fit Keplerian models as black lines; in the bottom two panels, residuals from the RVs to their respective Keplerian fits are displayed. Only RVs of observations that exist in both timeseries have been used here.

a systematic modulation that is not accounted for by the model, decreasing from a residual value of nearly 80 m s^{-1} to -37 m s^{-1} over seven observations. In fact, in the best fit of the full Lick timeseries this pattern is even slightly enhanced, and it can also be found in the published Keplerian and coplanar dynamical models in Ortiz et al. (2016) and Trifonov et al. (2018), respectively (Figure 1 in both publications). The residuals of the same observations in the Waltz DRS timeseries in contrast seem to scatter more or less randomly around 0, and do not show a significantly increased rms (both in the shortened and full timeseries).

To examine whether the mismatch between the Lick and Waltz DRS RVs for these observations are due to inherent errors or systematics of the latter, I correlated the absolute RVs and residuals of the Waltz DRS with several output metrics of the modelling process, such as results of the other fitting parameters, median chunk χ_{red}^2 and chunk velocity scatters of the respective observations; additionally, I checked correlations with properties of the input spectra, for instance the S/N , and instrument settings used when obtaining the spectra. However, no clear relations could be found that would explain the observed mismatch between the Lick and Waltz DRS RVs, and it therefore seems that the RVs computed by the Waltz DRS are not compromised by any internal artefacts.

In the analysis of HIP 36616 performed by Trifonov et al. (2018), the Lick RVs that

TABLE 3.6: Keplerian parameters for the HIP 36616 system, from the Waltz DRS and Lick RVs

Parameter	Waltz DRS RVs			
	Full dataset		Shortened dataset	
	Planet. comp.	Stellar comp.	Planet. comp.	Stellar comp.
P [days]	299.9	11427.9	299.7	11560.0
M_0 [deg] ^a	195.6	254.2	201.8	255.5
e	0.04	0.72	0.05	0.73
ω [deg]	129.7	149.5	122.6	149.5
K [m s ⁻¹]	135.7	4012.9	134.3	4022.9
$m \sin i$ [M_{jup}] ^b	6.87	556.19	6.79	557.13
a [AU]	1.09	13.36	1.09	13.47

Parameter	Lick RVs			
	Full dataset		Shortened dataset	
	Planet. comp.	Stellar comp.	Planet. comp.	Stellar comp.
P [days]	299.5	11618.6	299.3	12103.9
M_0 [deg] ^a	196.5	256.0	203.4	260.2
e	0.06	0.73	0.06	0.74
ω [deg]	126.6	149.4	119.4	149.4
K [m s ⁻¹]	136.5	4015.2	135.4	4040.2
$m \sin i$ [M_{jup}] ^b	6.90	555.04	6.84	558.03
a [AU]	1.09	13.51	1.09	13.89

^a The mean anomalies are calculated at the observational epoch $t_0 = 2\,451\,400.0$ JD.

^b The mass has been derived by solving the mass function for fixed inclinations i of 90° and with the primary mass $M_* = 1.9 M_\odot$.

produce outlier residuals when fitted with a double-Keplerian model actually help to constrain a dynamical model with mutually inclined orbits, which results in a significant model improvement as compared to coplanar edge-on dynamical models and greatly reduces the rms of these RVs. The authors conclude however that due to the sparse sampling of observations around that time, and throughout the periastron passage in general, it is well possible that the model just fitted noise-induced outliers instead of a real signal — particularly because the model of mutually inclined orbits mainly depends on these few data points.

As the RV results from the Waltz DRS greatly reduce the rms of these observations with respect to a double-Keplerian model, it seems probable that the Lick RV outliers could be an artefact, or at least that both pipelines do not deliver an exact representation of the true stellar RVs at these observation times. It is however unclear to me what might cause such a mismatch between the results of the data reduction codes; with the observations being obtained close to periastron passage in the left wing of the RV curve, there could actually be an astrophysical reason to it, such as tidally induced oscillations in the primary star which lead to a deformation of the stellar absorption lines. Another possibility would be that the stellar companion is actually not completely "invisible" in the spectra, so that its absorption lines, while being too weak to be noticed directly, lead to a slight deformation of the primary's absorption lines once the stars' RVs become comparable during periastron passage. If such a deformation is absorbed by the LSF model in one code, while it instead leads to a slight shift of the wavelength scale in the other software, then these effects could

lead to different RV results.

To better examine the outlier RVs and understand their cause and the implications on the physical nature of the HIP 36616 system, a more thorough analysis is required, including new dynamical models with varying inclinations fitted to the Waltz DRS RVs. This is beyond the scope of this work, which focuses on the development of the Waltz DRS and an early evaluation of its results. In any case, it can be concluded that the RV results for HIP 36616 computed by the Waltz DRS again suffice to clearly identify the stellar and planetary companions, and the Keplerian model delivers very similar results as when using the original Lick RVs. The Waltz DRS thus performs on a comparable level of precision as the original Lick data reduction.

3.5.3 An early test on SONG spectra

The 1 m Hertzprung SONG telescope on Tenerife has been in operation since 2014 and is equipped with an I2 cell and an Echelle spectrograph, connected to the Coudé focus of the telescope, with a resolving power of up to roughly 110 000 (depending on the used slit width); RVs are extracted with the software iSONG, and the project reaches an RV precision down to 1 m s^{-1} for some stars (e.g. Antoci et al., 2013; Grundahl et al., 2017; Fredslund Andersen et al., 2019). As our group has resumed RV measurements for a number of giant stars with promising planet candidates at the SONG telescope, we possess an archive of SONG spectra and respective RVs as computed by iSONG. Due to the much better precision achieved by SONG as compared to Lick, this data allows a better determination of the actual performance of the Waltz DRS. Furthermore, analyzing SONG spectra with the Waltz DRS allows the flexibility of the software to be tested, and shows whether the goal of making it easily accessible to different instruments has been attained.

Due to time constraints, I have thus far only been able to analyze one complete time-series of SONG observations with the Waltz DRS, namely the spectra obtained of ϵ Cyg (HIP 102488), which is a highly eccentric stellar binary whose RV curve shows short-period variations in addition to the long-period, high-amplitude modulation caused by the stellar companion. The short-period signal at first sight looks like it might be caused by a planetary companion on an S-type orbit around the primary component. The extracted RVs by iSONG, up until JD 2 458 480 (late December 2018), have been used in my work on the system in Chapter 4 (and published in Heeren et al., 2021), which revealed that a planetary cause of the short-period RV variations is highly unlikely (for a more thorough interpretation of the data I point to my analysis in the mentioned chapter). Since then, additional spectra have been obtained with SONG, which I also use in my test of the Waltz DRS; observations from an asteroseismic campaign on the star in contrast have been excluded. The dataset thus consists of 227 spectra of ϵ Cyg with I2 cell, 11 template observations of the star without I2 cell, and 5 A-star observations of HR 6410 for the determination of the LSF in the template creation. The A-star and template observations were shared with me by F. Grundahl from the SONG consortium.

The chunk setup is the same as for the simulated spectra in Section 3.5.1, so 24 spectral orders of 22 chunks each, totalling 528 chunks per spectrum; however, the chunk width was chosen to be the same as used in iSONG, so 91 pixels, and a padding of 23 pixels was used. The MultiGaussian LSF model was also constructed as in iSONG: The satellite Gaussians have widths of 0.7 pixels, and their positions are 0.7, 1.2, 1.6, 2.0 and 2.6 pixels to the left and right of the central Gaussian, which has a width




FIGURE 3.30: RVs (from SONG spectra) and Keplerian model results for the star ϵ Cyg (HIP 102488): The top two panels show the original iSONG RVs and the RVs computed with the Waltz DRS in blue, respectively, with their respective best-fit Keplerian models as black lines; in the bottom two panels, residuals from the RVs to their respective Keplerian fits are displayed. Only RVs of observations that exist in both timeseries have been used here.

of 0.4 pixels and a fixed amplitude of 1. Furthermore, the recentering algorithm was included in the LSF model. Pixel weights were all set to 1, and neither a bad-pixel-mask nor a telluric mask were used in this first test. Finally, the modelling was again performed in two runs for all observations, first using the `SingleGaussian` LSF model and second the `MultiGaussian`. In the second run, the chunk parameter results from the first run were used as starting values for the chunk velocities, the I2 and stellar template depths, and continuum parameters, and evaluated 3rd-degree polynomials were again used to determine starting values for the wavelength parameters; the LSF starting values were determined by fitting the median 1st-run LSF, and they were limited to ± 0.4 times their starting value. All these settings were used in the modelling of I2-observations as well as for the determination of the instrument LSF in the template creation process. For the stellar template, all 11 template observations were summed up to reach the highest possible S/N , and the same deconvolution parameter settings as for the Lick template creation were used. An overview of the most important general parameters is given in Table A.3.

Figure 3.30 displays the resulting RV timeseries in the second panel from top, as computed by the SONG-based weighting algorithm. The RV results from the iSONG reduction pipeline for the corresponding observations are shown in the top panel. Both timeseries were fitted with a double-Keplerian model, where I used the best-fit

TABLE 3.7: Keplerian parameters for the HIP 102488 system, from the Waltz DRS and iSONG RVs extracted from SONG spectra

Parameter	Waltz DRS		iSONG	
	Planet. comp.	Stellar comp.	Planet. comp.	Stellar comp.
P [days]	298.9	18937.7	299.4	18745.7
M_0 [deg] ^a	116.6	238.4	131.8	237.1
e	0.21	0.93	0.24	0.93
ω [deg]	243.9	275.3	239.1	275.3
K [m s ⁻¹]	48.1	4616.4	48.0	4612.4
$m \sin i$ [M_{jup}] ^b	1.65	277.59	1.64	277.33
a [AU]	0.90	15.44	0.91	15.34

^a The mean anomalies are calculated at the observational epoch $t_0 = 2451400.0$ JD.

^b The mass has been derived by solving the mass function for fixed inclinations i of 90° and with the primary mass $M_* = 1.103 M_\odot$.

results from the double-Keplerian model in Heeren et al. (2021)/Section 4 as starting values. The bottom two panels in Figure 3.30 show the residuals of the SONG and Waltz DRS RVs to the respective best-fit models of the data. Both timeseries clearly reflect the long-period, high-amplitude modulation caused by the stellar companion, and the residuals of the RVs to their respective models look very similar. The rms of the residuals is 16.3 m s^{-1} for the original SONG RVs, and 16.0 m s^{-1} for the Waltz DRS RVs, so the performance as measured by the scatter around the respective models is very comparable. The best-fit results for the Keplerian parameters of the two models are displayed in Table 3.7, and differences between the model results to the Waltz DRS and iSONG timeseries are mostly on the 1%-level or below. The largest relative mismatch can be noted for the angular parameters, that is the longitude of periastron ω and the mean anomaly M_0 , of the planetary orbit, which might hint at some differences between the computed RVs of the two pipelines which lead to a slight phase-shift of the RV curve.

Figure 3.31 shows GLS periodograms of the RVs (top) and RV residuals (bottom), both for the Waltz DRS results (blue) and the iSONG data (orange dotted), respectively. In the periodograms of the RVs, only small differences are again visible, notably at periods slightly larger than the fitted period of the planetary companion between 300 and 400 d, where the iSONG RVs apparently possess additional power. Similarly, the periodograms of the RV residuals mostly produce peaks at the same periods for both timeseries, but the heights of the peaks vary; particularly the two peaks around the best-fit planetary period of roughly 300 d are nearly twice as pronounced in the periodogram of the Waltz DRS data as compared to the one of the iSONG RV residuals. It should be noted at this point, however, that from the analysis in Section 4 we know that the double-Keplerian model certainly does not deliver an impeccable description of the data, which is reflected for instance in the appearance of the two peaks in the periodograms of the residual RVs; I have used the model here simply to allow a first rough comparison of the Waltz DRS results with the iSONG RVs, as it fits the data well enough to identify highly significant and systematic differences. The differing peak heights around 300 d clearly show that some differences between the RV timeseries do exist, as is expected; nevertheless, they do not allow a conclusion to be drawn about which pipeline performs better, because we know of the difficulties with the adapted model and therefore also expect deviations from




FIGURE 3.31: GLS periodograms of the RV timeseries of ϵ Cyg (HIP 102488). *Top*: GLS of the original iSONG RVs (orange dotted) and the results from the Waltz DRS (blue). *Bottom*: GLS of the RV residuals from iSONG and the Waltz DRS to their respective best-fit Keplerian models.

that model.

As a more robust measurement for comparison, the computed RV uncertainties can be used, which reflect the scatter of individual chunk velocity results. For the iSONG RV timeseries, the median measurement error is 2.32 m s^{-1} , while the Waltz DRS RVs have a median error of 2.73 m s^{-1} . The iSONG RVs thus seem to perform a little better; however, they have been originally computed on a much longer timeseries, which included a total of 5 063 observations acquired as part of the before-mentioned asteroseismic campaign. These observations were acquired with longer exposure times (120 s as compared to the 90 s generally employed for the other ϵ Cyg observations), thus achieved a higher S/N and smaller chunk velocity scatter. Due to the chunk velocity weighting taking into account all chunks from all observations, calculating RVs with that full timeseries will bring down the chunk errors σ_j as computed in Equation 3.33, and consequently all RV uncertainties will be smaller. It is therefore likely that the difference between the iSONG and Waltz DRS median errors can be largely attributed to the longer timeseries used by iSONG.

To better understand the velocity scatter achieved by the Waltz DRS, I computed the outlier-resistant standard deviations of the chunk velocities for each observation; the median of these chunk velocity scatters is 118 m s^{-1} , and the best observation achieved a scatter of 92 m s^{-1} . With iSONG, generally velocity scatters between 90 and 150 m s^{-1} are achieved for stars like Arcturus and ϵ Cyg (F. Grundahl, private communication). This corresponds well to the results of the Waltz DRS, and proves that the two reduction pipelines perform on a very similar level of precision.

3.6 Computational performance of the Waltz DRS

An important aspect of the Waltz DRS is its computational performance: In the future, when observations are regularly conducted with the Waltz, the code should

be able to reduce all data obtained within a night quickly the following day, ideally right after the end of the observations. This is particularly true for the extraction of spectra, so that the data will be available quickly for further analysis. At Lick Observatory, when obtaining spectra for the G-/K-giant survey, a typical observation night consisted of up to 50 observations of stellar targets, and a number of calibration spectra obtained at the beginning and end of the night (usually 10 to 40 flatfields, a few ThAr spectra, 10 to 30 bias frames and two I2 spectra). With the Waltz project, we aim at a similar number of observations in good nights, so the number of spectra stated above should serve as a benchmark for the capacity of the Waltz DRS.

In my tests of the code, I therefore monitored its computation time, both for the reduction of spectra using the modified CERES package, as well as for the I2 analysis. The performance obviously greatly depends on the computer that is used, and I tested the code on various machines. The following results were achieved with an Asus laptop, running an 8th generation 1.6 GHz Intel Core i5 processor with four cores. Furthermore, in the optimal extraction process of spectra of the data reduction, and in the I2 analysis of extracted spectra, parallel processing can be employed in the Waltz DRS to speed up the overall computation. I chose four processor cores in all my tests.

Reduction of spectra

As expected, the most costly step in the reduction of spectra is the optimal extraction: Using the specs described above, the computation of extraction weights took roughly 2 to 3 min, with slight variation depending on the quality of the masterflat spectrum used for the computation. The optimal extraction itself then depends on the input spectrum: Calibration frames generally required roughly 40 s per frame; for stellar spectra however, the computation time varied widely between 40 to 200 s, apparently depending on the (spectral) type of star and the quality of the spectrum. This is most probably due to the iterative procedure employed to ignore cosmics and other artefacts in the extraction, which takes longer for spectra with many absorption lines and low S/N . Simple extraction (which I performed for stellar spectra) in contrast took less than 10 s.

All other steps are performed much quicker than the optimal extraction: The pre-processing of the input spectra took on the order of 1 s, the order tracing was performed within 10 to 15 s (again with some dependence on the quality of the masterflat frame), and the computation of a global wavelength solution required not more than 5 s per extracted ThAr spectrum. The determination of the instrumental drift throughout the night finally took on the order of 1 to 2 s.

A test run, performed on six flatfield frames, five ThAr frames and five stellar spectra from an observation night, thus required an overall computation time of 15.5 min to produce the reduced science spectra both in "SONG" and "CERES" output formats. From these measurements, the total reduction time for the benchmark observation night of 50 stellar spectra can be estimated: Assuming an average time of 120 s required for the optimal extraction of each stellar frame, the reduction of all spectra should not take longer than 2 h in total. When adopting the longest measured extraction time of 200 s for each frame, the total reduction time still stays below 3 h.

Extraction of RVs

The extraction of RVs from the reduced spectra can be split up into three broad steps: creation of stellar template spectra, modelling of observation spectra, and combination and weighting of chunk velocities to the final RV timeseries. For the first two, the achieved speed again greatly depends various metrics: the quality of the spectra; the number of chunks, and pixels per chunk, to be modelled; the parameters, modules and processes employed, such as oversampling, number of modelling runs, or LSF model; and the limits put on model parameters. The computation times that I report here are for the tests performed on Lick and SONG spectra as described in Section 3.5.

The creation of stellar templates generally required on the order of 4 to 6 min both for SONG and Lick spectra, and a large part of the time was spent on the deconvolution of the template observations. Modelling of the observations showed greater variance: Whereas SONG spectra only required 2 to 5 min per observation, the modelling of each Lick spectrum took between 8 and 12 min. The discrepancy between SONG and Lick observations is probably mostly caused by the higher total number of chunks employed for Lick spectra, and their mostly lower S/N as compared to SONG spectra. In both cases, as the modelling of observations was always performed in four parallel processes, the total computation time for a batch of observations was roughly the sum of the individual computation times divided by four. The 38 observations of the Lick target HIP 95459 thus required 1.6 h, and the 227 modelled SONG spectra of HIP 102488 took 4 h. The computation of the RV timeseries finally only takes a few seconds.

Due to a lack of test spectra from the Waltz project, I do not know how computationally expensive the modelling of Waltz spectra will be. If we assume however that the spectra will resemble the Lick spectra concerning their S/N and spectral format, and we adopt the longest modelling time required for Lick spectra, the total computation time for the benchmark of 50 spectra obtained in the best observation nights will be roughly 2.5 h (using four processor cores on the same machine as utilized for the tests).

3.7 Conclusion and outlook

In this chapter, I have presented the current state of the Waltz DRS, a data reduction pipeline intended for the reduction and calibration of spectra from the Waltz telescope and extraction of precise RVs on the m s^{-1} -level. The main focus in the development of the software was placed on usability and functionality, to allow users a quick understanding of the code when starting to work with it, and make it easy to change parameters or completely adapt the DRS to a new instrument. The philosophy concerning the structure of the software thus was to strictly separate fundamental and generic algorithms from user- and instrument-specific modules and parameter definitions, and bundle the latter in a clear and easy-to-access manner.

The software builds upon three existing code packages, namely CERES for the extraction of spectra, barycorrpy for the computation of the barycentric velocity correction, and the unpublished pyodine package for the determination of RVs using the I2 cell method. All three packages were incorporated into the common Waltz DRS repository, and modifications were made where necessary to comply with the guidelines explained in the previous paragraph. The CERES package thus required

a complete translation from Python 2 to Python 3, and a restructuring to trace all important reduction parameters back to a common input parameter file. Furthermore, it was modified to allow the usage of Python objects for the storing of basic data, such as spectra and reduction results, which serves a clearer structuring of the top-level code and eases access to the data. For barycentric correction, a wrapper module was created to enable communication between the main reduction routine developed for the Waltz project and the underlying functions of the barycorrpy package.

Similar to the CERES package, the pyodine package was modified to bundle all important parameters in a single parameter input file. Furthermore, it was upgraded with various modules to offer all functionalities that could possibly be required in the determination of RVs, such as the usage of non-uniform pixel weights in the modelling process, and employment of bad-pixel masks and telluric masks. The biggest update concerned the usage of chunks that are constant in barycenter-corrected wavelength space, thus always covering roughly the same spectral content; this method is employed by other I2-cell Doppler codes, and was not implemented in the pyodine package before. Finally, main routines were written for the creation of deconvolved stellar template spectra, the modelling of observations and the computation of the RV timeseries.

Due to a lack of suitable data from the Waltz telescope at time of writing this dissertation, the Waltz DRS, thus far, could not be completely tested on Waltz spectra. Only the spectrum reduction, using the modified CERES package, was performed on early Waltz spectra. The main reduction routine and the employed parameters were tailored to achieve the best results, and some analysis was presented in Section 3.3.4. Most importantly, the optimal extraction of the spectra seems to work as intended, and the wavelength calibration delivers a sufficient precision for the later usage in the I2 analysis code.

The I2 Doppler analysis code, built upon the modified pyodine package, was mainly tested on simulated data and archival spectra from the Lick survey of G-/K-giant stars. On simulated data, the software proved its general ability to model observation spectra and determine precise RVs, accurate on the m s^{-1} -level. Lick observations of the star HIP 96459 served as a test bed in the development of the software, to find a good combination of parameters and procedures to achieve the best possible results. The resulting RV timeseries for that star results in a rms of 10.4 m s^{-1} , as compared to 9.3 m s^{-1} for the original Lick RVs, which suggests a similar performance of both pipelines. Tests on several other stars that show RV variations induced by planetary and stellar companions confirm this impression, and the resulting RV timeseries from the Waltz DRS allow a clear identification of all previously published companions with very similar Kepler parameters. Furthermore, a first test of the Waltz DRS on SONG spectra of the star ϵ Cyg successfully proved the flexibility of the software: The code could be quickly adapted to the new instrument, and the resulting RV timeseries from the first full run on SONG observations suggest a similar level of precision reached by the Waltz DRS as the dedicated iSONG pipeline.

With development of the main structure of the Waltz DRS finalized, and first results showing a very promising performance of the software, future work on the code will focus on a better determination of the actual precision reached. This is best done on observations of a star which shows very little intrinsic RV variation. The Lick spectra of the star HIP 96459, which I used for early optimization and determination of

the Waltz DRS precision, still carry an RV scatter of roughly 10 m s^{-1} due to intrinsic stellar jitter, an effect which generally prevents us from reaching a much better understanding of the precision when testing on stars from the G-/K-giant sample. Therefore, I plan to evaluate the Waltz DRS performance more thoroughly on SONG spectra: The SONG consortium has collected a large archive of spectra from RV standard stars, and generally stars with very low RV scatter, some of which have resulted in RV rms well below 1 m s^{-1} over several consecutive nights when analyzed with the iSONG pipeline. F. Grundahl from the SONG consortium has agreed to share the required data with me to make these tests possible. Using this data, I will further optimize the parameters and procedures employed in the Waltz DRS, and better determine the minimum precision achieved by the code.

In a next step, the Waltz DRS will be adapted to work with Waltz spectra as soon as the first suitable data from the Waltz telescope is available. It should be straight forward to implement the proper general parameter settings: For instance, the Waltz spectrograph has an angular dispersion of roughly $7.4 \cdot 10^{-4} \text{ rad}/\text{\AA}$ (at a wavelength of 5500 \AA), which results in a linear dispersion of $\sim 29 \text{ pix}/\text{\AA}$ (using the camera focal length and detector pixel size as defined in Section 2.2.4). The chunk size used in the I2 method generally is 2 \AA in wavelength space, which then corresponds to roughly 60 pixels of a Waltz spectrum. Similarly, most of the other parameters can be determined. A more complicated task will then be to find a good LSF description, which will require testing different LSF models, for instance on O-star spectra, and evaluating the achieved residuals between model and observations. Then, stellar templates can be created with well-performing LSF models, and the overall performance of the Waltz DRS on Waltz spectra can again be evaluated through the resulting RV scatters from modelled timeseries. We are therefore planning to schedule regular observations of RV standard stars, which can be used to optimize the Waltz DRS implementation for the Waltz spectrograph by varying the parameter settings, similarly as done for Lick.

Furthermore, I want to include and test some additional methods in the Waltz DRS: For instance, Díaz et al. (2019) developed a technique which "cleans" spectra obtained with an I2 cell by employing the forward model used in the determination of RVs to reconstruct and subtract the I2 spectrum from the observations. The resulting I2-free stellar spectra reached an rms scatter on the 1.5%-level, as compared to the original stellar template, and could potentially be used to compute activity indicators of spectral lines that fall within the I2 region (e.g. Sodium doublet, Fe I, Mg I), as well as performing line-diagnostics such as bisector spans (depending on the LSF and stability of the used spectrograph). This method might be a beneficial upgrade to the Waltz DRS. On a similar note, in the data reduction software SERVAL (Zechmeister et al., 2018) stellar template spectra are created through a combination of all observations of a given star, shifted in wavelength by their respective RVs, which delivers extremely high S/N . This technique could be used in the I2 analysis code as well, but would have to include the additional steps of deconvolving each observation and subtracting the I2 spectral features from it, which will certainly introduce additional errors. Still, the resulting spectrum might deliver a better stellar template, and it would be interesting to test the performance of the Waltz DRS with that method.

Finally, I am planning to publish the Waltz DRS and make it accessible to the scientific community under an open-access license. Thus far, there are no I2 analysis codes openly available, which makes it difficult to reanalyze data, and compare RV

results of I2 spectra processed with different data reduction pipelines. Furthermore, having access to a well-tested and comparably easy-to-adapt reduction software will greatly benefit small telescope projects such as the Waltz project, which often cannot afford to pay experts in the field to develop a tailored reduction software. As I built the Waltz DRS on the existing CERES and pyodine packages, but have greatly extended them (and translated the former to Python 3), I want to cooperate with the authors of these packages. Working together, the CCF-based RV computation and spectral classification features of the original CERES package could be included and the resulting code would then offer an "all-inclusive" package, allowing all steps to be performed from spectrum reduction, over basic spectral analysis, to RV determination through both the CCF and I2 methods.

IS THERE A PLANET IN THE HIGHLY ECCENTRIC STELLAR BINARY ϵ CYGNI?

based on

Precise radial velocities of giant stars XV. Mysterious nearly periodic radial velocity variations in the eccentric binary ϵ Cygni

Paul Heeren¹, Sabine Reffert¹, Trifon Trifonov², Ka Ho Wong⁴, Man Hoi Lee^{3,4},
Jorge Lillo-Box⁶, Andreas Quirrenbach¹, Torben Arentoft⁵, Simon Albrecht⁵, Frank
Grundahl⁵, Mads Fredslund Andersen⁵, Victoria Antoci^{7,5}, and Pere L. Pallé^{8,9}

published in A&A, 29.03.2021

¹ Landessternwarte, Zentrum für Astronomie der Universität Heidelberg, Königstuhl 12, 69117 Heidelberg, Germany

² Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany

³ Department of Earth Sciences, The University of Hong Kong, Pokfulam Road, 40000 Hong Kong, PR China

⁴ Department of Physics, The University of Hong Kong, Pokfulam Road, Hong Kong, PR China

⁵ Stellar Astrophysics Centre, Department of Physics and Astronomy, Aarhus University, Ny Munkegade 120, DK-8000 Aarhus C, Denmark

⁶ Centro de Astrobiología (CAB, CSIC-INTA), Depto. de Astrofísica, ESAC campus 28692 Villanueva de la Cañada (Madrid), Spain

⁷ DTU Space, National Space Institute, Technical University of Denmark, Elektrovej 328, DK-2800 Kgs. Lyngby, Denmark

⁸ Instituto de Astrofísica de Canarias, E-38200 La Laguna, Tenerife, Spain

⁹ Universidad de La Laguna (ULL), Departamento de Astrofísica, E-38206 La Laguna, Tenerife, Spain

The contents of this chapter have been published in an almost identical version in *Astronomy&Astrophysics* (Heeren et al., 2021), and include some contributions from co-authors. My own work comprises the following parts: Section 4.1, "Introduction"; Section 4.1, "Stellar properties", with exception of the asteroseismic analysis (last paragraph, contributed by T. Arentoft); Section 4.3, "Observations", also except of the discussion of the asteroseismic measurements (third paragraph, contributed by T. Arentoft); Section 4.4, "Analysis of the RV data", with exception of the co-orbital analysis (Section 4.4.6, performed by J. Lillo-Box), and parts of the stability analysis in Section 4.5.2 (first and second paragraph, contributed by K.H Wong and M.H. Lee); finally, all analysis in Section 4.6, "Possible alternative explanations for the short-period RV variations", and the discussion in Section 4.7, "Summary & conclusions", are my own work. RV data used in this analysis can be found in Appendix B.

4.1 Introduction

To date, more than 4000 extrasolar planets have been confirmed, and more than 800 of these were discovered with Doppler spectroscopy, also known as the radial velocity (RV) method.¹ While most of the discovered planets orbit main-sequence (MS) stars, the number of detections around evolved stars has also risen quickly: Since the first discovery of a planet around a giant star in 2001 (Frink et al., 2002), 112 exoplanets orbiting evolved intermediate-mass stars have been published.²

Giant stars allow us to extend planet surveys to higher stellar masses, while also still being sensitive in the lower mass regime: Whereas intermediate-mass MS stars show only few and rotationally broadened absorption lines, their evolved counterparts, such as K- and G-giant stars, have many sharp spectral lines and are therefore perfect targets for RV measurements. Furthermore, the discoveries of planets around these stars help to improve our understanding of the evolution of planetary systems once the stars evolve into giants (Villaver and Livio, 2009; Reffert et al., 2015).

Despite the overall large number of extrasolar planets discovered to date, only a comparably small fraction of $3 \sim 4\%$ has been detected in stellar binaries³. In the case of giant stars, this number is especially small, with only five known cases: 11 Com (whose companion falls into the brown dwarf regime, Liu et al., 2008), γ Leo (Han et al., 2010), 91 Aqr (Mitchell et al., 2013), 8 UMi (Lee et al., 2015), and HD 59686 (Ortiz et al., 2016), and only one of these (HD 59686) is a spectroscopic, that is, rather close binary. To a large extent, this small fraction of discovered planets in binary star systems can be explained by the fact that most exoplanet surveys focus on single stars, which is unfortunate since binary systems harboring planets serve as good special cases to constrain the theory of planet formation and evolution.

In this work we study the K giant star ϵ Cyg, which has been observed spectroscopically for more than 100 years and is known to undergo large RV changes, hinting at the existence of a close stellar companion to the primary component that is not directly visible. We use our own RV measurements of the star both from the Lick and Stellar Observations Network Group (SONG) telescopes to derive a precise orbit of the spectroscopic stellar companion and to investigate whether additional short-period RV variations are caused by an S-type (i.e., circumstellar) planet in the system.

In Section 4.2 we present the known properties of the ϵ Cyg system and its primary component. Section 4.3 describes the RV data sets used in this analysis. Next, in Section 4.4, we perform the RV modeling. In Section 4.5 we then analyze the stability of the system using an N -body code. In Section 4.6 we investigate alternative explanations for the short-period signal. Section 4.7 summarizes and concludes our analysis.

4.2 Stellar properties

ϵ Cyg A is a bright ($m_V = 2.48$ mag) K0 III giant star. Its HIPPARCOS parallax is 44.86 ± 0.12 mas (Leeuwen, 2007), which puts it at a distance of 22.29 ± 0.06 pc. It has been observed spectroscopically since the beginning of the 20th century (see

¹<https://exoplanets.nasa.gov>

²<https://www.lsw.uni-heidelberg.de/users/sreffert/giantplanets/giantplanets.php>

³<https://www.univie.ac.at/adg/schwarz/multiple.html>

TABLE 4.1: Stellar properties of ϵ Cyg A

Parameter	Value
Apparent magnitude m_V [mag] ^a	2.48 ± 0.01
Luminosity L_\star [L_\odot] ^b	$57.1^{+0.5}_{-0.4}$
Color index $B - V$ [mag] ^a	1.04 ± 0.01
Effective temperature T_{eff} [K] ^b	4805^{+16}_{-14}
Surface gravity $\log g$ [cm s^{-2}] ^b	$2.45^{+0.16}_{-0.05}$
Metallicity [Fe/H] [dex] ^c	-0.11 ± 0.03
Stellar mass M_\star [M_\odot] ^d	1.103 ± 0.042
Stellar radius R_\star [R_\odot] ^b	$10.94^{+0.08}_{-0.13}$
Parallax [mas] ^e	44.86 ± 0.12
Distance [pc] ^e	22.29 ± 0.06
Age [Gyr] ^b	9.62 ± 0.12
Spectral type ^f	K0 III
Frequency of max. power [μHz] ^d	32.16 ± 0.81

^a Oja, 1993

^b Stock, Reffert, and Quirrenbach, 2018

^c Montes et al., 2018

^d Arentoft et al. (to be published)

^e HIPPARCOS, the new reduction (Leeuwen, 2007)

^f Keenan and McNeil, 1989

e.g., Campbell and Moore, 1906; Kustner, 1908; McMillan et al., 1992), and has long been known to host a spectroscopic binary companion, for which Gray (2015) derived an orbital period around 55.1 yrs without putting an error on that value. The spectroscopic companion has never been observed visually though, and the spectra do not show any evidence of a second set of spectral lines nor any temporal changes of line asymmetry (Gray, 1982). McMillan et al. (1992) also detected short-period RV variations on a time scale on the order of a few 100 d with much smaller amplitude, which they did not analyze in detail, but simply compared them qualitatively to other examples of K-giants with fast variations. Further RV measurements of the ϵ Cyg system were performed from 1999-2010 by Gray (2015), who present a spectroscopic analysis of the binary system. As the formal uncertainties of these measurements are too large to clearly identify the short-period signal, Gray (2015) only reference the data taken by McMillan et al. (1992) and propose that the variations might be caused by stellar activity, modulated by the rotation of the star.

In addition to the spectroscopic companion, ϵ Cyg A has an optical companion, called ϵ Cyg B, at a separation of $71''$ (Montes et al., 2018), whose Gaia DR2 parallax and distance are 4.35 ± 0.02 mas and 229.9 ± 1.3 pc (Gaia Collaboration et al., 2018), respectively, and which is therefore not gravitationally associated. A second optical companion C has a separation of $78''$ (Montes et al., 2018), and its Gaia DR2 parallax is 45.51 ± 0.03 mas, putting it at a distance of 22.97 ± 0.01 pc and therefore very close to ϵ Cyg A (Gaia Collaboration et al., 2018). According to Montes et al. (2018), it is an M4V+ star, and as its proper motion also resembles that of the primary star (ϵ Cyg C: $+354.62 \pm 0.04$ mas/yr in r.a., $+329.18 \pm 0.05$ mas/yr in declination (Gaia Collaboration et al., 2018); ϵ Cyg A: $+355.66 \pm 0.08$ mas/yr in r.a., $+330.60 \pm 0.09$ mas/yr in declination (Leeuwen, 2007)), it can be expected to be a very wide physical companion. We do not pay any further attention to it within this

work though, as its separation from the primary at the distance of the system is at least $22.29 \text{ pc} \cdot 78 \text{ mas} = 1739 \text{ AU}$, and any influence therefore can be neglected.

We have two independent sources for the stellar properties of ϵ Cyg A: The first one comes from Stock, Reffert, and Quirrenbach (2018), who used a Bayesian interpolation scheme for evolutionary tracks. This method provides a probability for the star to either fall onto the red giant branch (RGB) or onto the horizontal branch (HB). It also delivers nonsymmetrical probability density functions (PDFs), which allows us to compute asymmetric 1σ confidence intervals for each of the stellar parameters. The mode values of the PDFs are used as the most probable values. According to this method, ϵ Cyg A is most probably an HB star ($P = 99.5\%$), with a stellar radius of $R_\star = 10.94_{-0.13}^{+0.08} R_\odot$ and a mass of $M_\star = 1.21_{-0.09}^{+0.46} M_\odot$. This mass estimate is smaller than the numbers that are usually adopted in the literature, which are closer to $2 M_\odot$ (see e.g., Gray, 2015). It is possible that previous authors assumed ϵ Cyg A to be an RGB star and thus overestimated its stellar mass. This is plausible considering the (very unlikely) RGB solution by Stock, Reffert, and Quirrenbach (2018), which yields a mass of $M_\star = 1.46_{-0.13}^{+0.43} M_\odot$, somewhat closer to the past estimates.⁴

In order to get a second independent measurement for the stellar properties, we also performed an asteroseismic campaign on ϵ Cyg A with SONG, and derived its mass from the frequency of maximum power ν_{max} of the star, using the relation as in Stello et al. (2017). We used the same approach as in Arentoft et al. (2019), who analyzed a similar SONG-dataset for the red giant ϵ Tau, to derive ν_{max} and its uncertainty. In short, following the methods described in Mosser and Appourchaux (2009) and Stello et al. (2017), we applied a Gaussian fit combined with a linear (background) trend to the oscillation signal in the power spectrum to determine ν_{max} . The uncertainty was found by performing the same fit to slightly modified versions of the power spectrum, in each of which a single oscillation mode had been subtracted, and using the variations in the determined ν_{max} -values to estimate the uncertainty; see Arentoft et al. (2019) for details. The analysis resulted in a $\nu_{\text{max}} = 32.16 \pm 0.81 \mu\text{Hz}$, which gives a mass estimate of $M_\star = 1.103 \pm 0.042 M_\odot$; the 1σ -error bars of this result overlap with the ones of the HB mass estimate from Stock, Reffert, and Quirrenbach (2018), and we can rule out the possibility of ϵ Cyg A being an RGB star with high certainty. Given the robustness and high fidelity of the asteroseismic method, we adopt that mass estimate for our analysis. Table 4.1 lists all the stellar parameters.

4.3 Observations

From 1999 to 2012, we monitored a sample of 373 G- and K-giant stars and measured their RV variations, using the Hamilton Échelle Spectrometer at the Lick observatory in California, USA. The spectrograph was fed by the 0.6 m Coudé Auxiliary Telescope (CAT) and covers a wavelength range of 3755 - 9590 Å at a resolving power of $R \sim 60\,000$. High RV precision is achieved using the iodine cell method as described by Butler et al. (1996). Our observations led to several planet detections (e.g., Frink et al., 2002; Reffert et al., 2006; Trifonov et al., 2014), among them a massive circumprimary planet in a close and eccentric binary system (Ortiz et al., 2016). The measurements also proved valuable in gaining a statistical understanding of the properties of extrasolar planets around giant stars (Reffert et al., 2015).

⁴The RGB solution is not published; we thank the authors for providing us with the result.




FIGURE 4.1: Uncorrected SONG RV measurements from the last week of our asteroseismic campaign on ϵ Cyg. Oscillations with periods shorter than a day are clearly visible.

The K-giant star ϵ Cyg (HIP 102488) is a member of our sample, and we collected a total of 109 RV measurements at Lick observatory. They cover a time span of more than eleven years (from June 2000 until November 2011), and their median measurement precision is 4.8 m s^{-1} .

In addition to the Lick data set, we obtained a total of 5 272 RV measurements of ϵ Cyg with the 1 m robotic SONG telescope on the island of Tenerife, Spain (Andersen et al., 2014; Fredslund Andersen et al., 2019). A set of 5 063 of these individual measurements were part of an asteroseismic campaign to better constrain the stellar properties of the K giant (see Section 4.2); these observations were therefore performed at very high cadence, with about 100 to 150 measurements per night, for a little more than a month from June 20 until July 27, 2018. From these measurements we can get an estimate of the short-term RV jitter, that is, the expected stochastic astrophysical stellar RV variation. Data from a few nights of this campaign are shown in Figure 4.1; the stellar oscillations are clearly visible, with peak-to-peak variations of $\sim 40 - 50 \text{ m s}^{-1}$ on timescales of several hours, in agreement with the value for ν_{max} found above. The observed oscillations are a combination of a number of individual oscillation modes, which together give rise to a short-term variation in the time series with a standard deviation of $\sim 11 \text{ m s}^{-1}$. This scatter agrees well with other stars with similar maximum-power frequencies ν_{max} as ϵ Cyg (compare Yu et al., 2018). Following Kjeldsen et al. (2008), we isolated the p-mode signal and determined the amplitude per radial mode to be $1.45 \pm 0.06 \text{ m s}^{-1}$. That value is approximately 9 times larger than the amplitude of the oscillation signal seen in the Sun, and accordingly the standard deviation and peak-to-peak variations of our asteroseismic measurements are roughly one order of magnitude larger than the solar values (see Kjeldsen et al., 2008; Fredslund Andersen et al., 2019).

For our orbital analysis of the system this large number of data points is unsuitable, as it puts too much weight on the SONG measurements and requires too much computational power when performing orbital fits. For this reason, we computed the median RV for each night with more than one observation. Our updated SONG data set then consists of 228 RV measurements, which fall between April 2015 and December 2018. They cover exactly the last periastron passage of the eccentric stellar

companion to ϵ Cyg, which occurred around February 2017, and thus enable us to determine the binary orbit with very high precision. Just as the Lick measurements, the high-resolution spectra taken by SONG are calibrated by the means of an iodine cell (Grundahl et al., 2017), and the measurements reach a similar precision, with a median of 2.3 m s^{-1} .

In addition to the long-period signal induced by the close stellar companion, both the Lick and SONG data sets show RV variations with periods just below 300 d (see 4.4.2), which could be caused by a planetary companion orbiting the primary star in an S-type configuration. To gain a better coverage of the orbits of the binary and the possible planetary companion, we complement our own two data sets with the older measurements from McMillan et al. (1992). The McMillan data set consists of 213 individual RV measurements of ϵ Cyg, collected between May 1987 and May 1992 with an interferometer at the 0.9 m-telescope of the Steward Observatory at Kitt Peak, USA. Calibrations were done with emission lines of a Fe-Ar hollow cathode lamp, yielding a precision of 12 m s^{-1} for all individual measurements.

Finally, we use six RV measurements from Kustner (1908) to cross-check our results for the orbital solution of the long-period stellar companion to ϵ Cyg. The observations by Kustner were carried out in Bonn, Germany, with a 30 cm refractor and a spectrograph constructed out of three prisms, and RVs were computed by comparing the positions of the stellar absorption lines to a reference iron spectrum. The observations of ϵ Cyg cover the time between July 1904 and September 1906. Their errors are large in comparison to modern values, and vary between 1000 and 2100 m s^{-1} , so we do not include these measurements in our fits. The merit of this data set however lies in the fact that it seems to record the periastron passage of the stellar companion two full periods back, and we can overplot and check our results.

We decided to refrain from using the RVs collected by Gray (2015) from 2001 until 2010 as they cover mostly the same time as our own Lick data and have much larger measurement errors, which is why they do not help to constrain our models. All in all, our combined data set for the fitting procedures consists of 550 RV measurements from McMillan, Lick and SONG.

4.4 Analysis of the RV data

4.4.1 Determining the orbit of the close stellar companion

We use a Keplerian model in combination with a Levenberg-Marquardt (LM) minimization scheme to fit the combined RV measurements of ϵ Cyg and determine orbital parameters for the spectroscopic companion, which we hereafter also refer to as the close stellar companion; the main host star is denoted as in the literature, ϵ Cyg A. The Keplerian fit incorporates eight free parameters: RV semi-amplitude K , orbital period P , eccentricity e , argument of periastron ω , and mean anomaly M_0 of the orbit, as well as zero-point offsets for the McMillan, Lick and SONG data sets, respectively. In order to account for the RV jitter of the star, we quadratically added a fixed value of 20 m s^{-1} to all individual measurement errors (see Section 4.4.3). Error estimation was performed with a Markov-Chain Monte Carlo (MCMC) approach, using the `emcee` python package (Foreman-Mackey et al., 2013): We constructed 32 walkers (four times the number of free parameters) with initial values drawn randomly from a small Gaussian ball around the best-fit solution from above, and let the sampler run for altogether 6000 steps (where we discarded the first 1000 steps).




FIGURE 4.2: *Top plots:* RV measurements of ϵ Cyg by Kustner (1908) (left), and McMillan et al. (1992), Lick observatory and the SONG telescope (right), along with two orbital models of the binary companion: **Fit 1** (black line) is the best-fit solution for the three modern data sets of McMillan, Lick and SONG, with an orbital period of 19618.6 d; **Fit 2** (gray line) was derived from fitting all four data sets and yields a period of 20175 d. *Middle plots:* Residuals of the RV measurements after removing the orbital solution of Fit 1. The last measurement of Kustner then is a clear outlier. *Bottom:* Residuals of the RV measurements after removing the orbital solution of Fit 2. Now all the Kustner RVs agree with the model within their 1σ -uncertainties.

This method produced 36 649 unique samples of the parameters; the 1σ intervals of these parameter sets around the mean values of the posterior distribution serve as our uncertainties.

The LM fitting scheme delivers a best-fit model with an orbital period of 19619 d at $\chi_{\text{red}}^2 = 1.93$, which lies within the errors from the mean period of the MCMC posterior distribution of $19611.5^{+117.4}_{-119.9}$ d $\approx 53.7 \pm 0.3$ yrs (see column “Single-Keplerian” of Table 4.2). Our results therefore are roughly one-and-a-half years less than the orbital period of about 55.1 yrs derived by Gray (2015), who used 53 RV measurements taken between 2001 and 2011 and complemented them with the McMillan data set as well as data by Griffin (1994).

When plotting our best-fit solution with a period of 19619 d over the measurements taken by Kustner (1908), one clearly notices that the last RV from that data set falls far away from the curve, by a value of about 7400 m s^{-1} , which corresponds to more than three times the formal error of 2100 m s^{-1} (compare Figure 4.2). If we fit all four data sets combined, the best Keplerian model delivers an orbital period of 20175 d and a χ_{red}^2 of 1.95. In this model, the former outlier now also agrees with the fit curve within its 1σ -uncertainty and falls exactly at the sharp peak of the curve; however, the fit leads to systematic (albeit small) deviations from the RVs of the three modern data sets by McMillan, Lick and SONG. As our further analysis focuses on these measurements, we decide to adopt the model with the shorter orbital period of 19619 d for now. We do not exclude the possibility that the result for the orbital period of the close stellar companion might still change on the level of a few percent with future observations coming in, but given our current data it seems most probable that the last Kustner measurement is actually an outlier.

Using the primary mass of $M_{\star} = 1.103 \pm 0.042 M_{\odot}$ and solving the posterior results of the mass equation $f(m)$ numerically for fixed inclinations $i = 90^{\circ}$, we derive a minimum mass of $278 \pm 7 M_{\text{jup}} = 0.265 \pm 0.007 M_{\odot}$ for the close stellar companion.

Our estimate places it therefore well above the brown dwarf regime. As the spectroscopic companion has never been imaged directly to our knowledge, and its absorption lines are not visible in the spectra, its luminosity must be much lower than that of the primary star ϵ Cyg A. If we additionally assume that both stars have been formed around the same time, as is expected for binaries, this opens up two possibilities for the nature of the companion: The first is a white dwarf, which would mean that it was originally the primary, more massive component in the system and therefore evolved more quickly.

The second possibility is that the spectroscopic companion to ϵ Cyg A is a main-sequence star with a mass lower than that of the primary, therefore evolving more slowly and being much less luminous. This places it anywhere between an early M-dwarf to an early G-type star, and constrains the possible orbital inclination to values larger than $\sim 14^\circ$, to keep its mass smaller than the mass of the primary, $\sim 1.103 M_\odot$. Therefore the apparent brightness in the V band should be between approximately 6 and 12 mag.

Griffin (1994) speculated about the possibility of imaging the companion directly (either at visual wavelengths or the infrared) and predicted that its angular separation on the sky could get nearly as large as $2''$ during apastron, but he had no certain solution for the orbit yet and used a semi-major axis of approximately 20 AU and eccentricity of 0.9 for his calculations. Our orbit solution puts the apastron at 30.5 AU from the primary. At the distance of the system, 22.1 pc, this would translate to a maximum projected separation of $1.38''$ in the case of an optimal position of the system on the sky. However, we were also able to derive a solid estimate of the argument of periastron of the orbit, which is at 275° and therefore places the periastron of the orbit nearly exactly in-between the observer and the primary component, and the apastron behind the primary ϵ Cyg A as seen from Earth. Hence the angular separation at apastron will only become large for a low inclination of the orbit. For an inclination higher than 80° , the maximum separation of the two components would be reached around the co-vertices of the orbital ellipse (semi-minor axis: $b \approx 5.8$ AU, projected separation: $0.26''$).

Modern direct imaging instruments offer the capability to resolve the two stars: As an example we used the exposure time calculator of the ESO instrument *SPHERE-IRDIS*⁵ to calculate the achievable contrast for ϵ Cyg within an exposure time of 1800 s with median weather conditions, which resulted in a magnitude difference of ~ 8.4 mag at a separation of $0.1''$, and ~ 14 mag at the co-vertex separation of $0.25''$. With the magnitude difference of the two stars probably not being higher than 12 mag in the visual, it might be an interesting option to try and image the spectroscopic companion directly. According to our calculations, it will reach the co-vertex around JD 2459797, that is, in August 2022, which is also the best time within the year to observe ϵ Cyg directly from Cerro Paranal Observatory.

4.4.2 GLS periodogram of the RV residuals

We aim to find relevant short-period variations in the RV measurements of ϵ Cyg by calculating generalized Lomb-Scargle (GLS) periodograms as described in Zechmeister and Kürster (2009). As the RV data are largely dominated by the signal of the stellar companion with a semi-amplitude of $K \approx 4.6 \text{ km s}^{-1}$, we first subtracted the best-fit orbital solution of the binary from the measurements and then calculated

⁵<http://www.eso.org/observing/etc/bin/gen/form?INS.NAME=SPHERE+INS.MODE=IRDIS>

TABLE 4.2: Keplerian parameters of the ε Cyg binary system, from the single- and double-Keplerian models

Parameter	Single-Keplerian Stellar comp.		Double-Keplerian	
	MCMC	Best-fit	Stellar comp. MCMC	Best-fit
P [days]	$19611.5^{+117.4}_{-119.9}$	19618.6		
M_0 [deg] ^a	160.8 ± 1.2	160.9		
e	0.9295 ± 0.0003	0.9295		
ω [deg]	275.30 ± 0.06	275.29		
K [m s ⁻¹]	4600.7 ± 1.7	4600.9		
$f(m)$ [M _{jup}] ^b	10.411 ± 0.016	10.407		
$m \sin i$ [M _{jup}] ^c	278 ± 7	278		
a [AU]	15.8 ± 0.2	15.8		
Parameter	Stellar comp.		Planetary comp.	
	MCMC	Best-fit	MCMC	Best-fit
P [days]	$19502.9^{+90.5}_{-88.6}$	19575.7	291.1 ± 0.1	291.2
M_0 [deg] ^a	159.7 ± 0.9	160.4	$138.8^{+17.1}_{-17.7}$	129.2
e	0.9295 ± 0.0002	0.9297	$0.150^{+0.056}_{-0.058}$	0.173
ω [deg]	275.36 ± 0.06	275.38	$267.04^{+17.88}_{-16.75}$	276.24
K [m s ⁻¹]	4607.3 ± 1.8	4607.5	$29.7^{+1.5}_{-1.6}$	30.3
$f(m)$ [M _{jup}] ^b	10.390 ± 0.016	10.386	$(8.009^{+1.178}_{-1.230}) \cdot 10^{-7}$	$8.366 \cdot 10^{-7}$
$m \sin i$ [M _{jup}] ^c	277^{+7}_{-6}	277	1.02 ± 0.06	1.04
a [AU]	15.7 ± 0.2	15.8	0.89 ± 0.01	0.89

^a The mean anomalies are calculated at the first observational epoch in the data set of McMillan, $t_0 = 2446945.9465$ JD.

^b The expression $f(m)$ denotes the mass function: $\frac{(m \sin i)^3}{(M_* + m)^2} = \frac{P}{2\pi G} K^3 \sqrt{(1 - e^2)^3}$.

^c Masses and errors have been derived by solving the mass function numerically for fixed inclinations i of 90° , taking the uncertainty of the primary mass M_* into account.

the periodograms of the residuals. Figure 4.3 shows the significance of variations with periods between 100 and 500 d, for the three individual data sets by McMillan, Lick, and SONG (lower three panels), as well as for the combined measurements (uppermost panel). False-alarm probabilities (FAPs) of 5, 1, and 0.1% are depicted by the gray dash-dotted, dashed, and dotted lines, respectively.

All three data sets show a highly significant peak at periods slightly shorter than 300 d that greatly exceeds the FAP of 0.1%, but the exact positions of those peaks vary: In the McMillan data the highest peak lies at 277.9 ± 2.8 d, the Lick data set shows its strongest signal at 287.3 ± 1.7 d, and the SONG data set at 291.2 ± 3.0 d. The uncertainty estimation used by Zechmeister and Kürster (2009) is based on the curvature of the peaks in the periodograms (as described in Ivezić et al., 2014), and in simulations, using fixed periods with varying white noise amplitudes and the sampling of the original data sets, we found a good agreement with these errors.

The period of the most dominant signal changes by more than 13 d over the course of 30 years from the first until the last measurement. All the peak periods lie at least 1σ away from each other; in the case of the McMillan and Lick periods the separation is

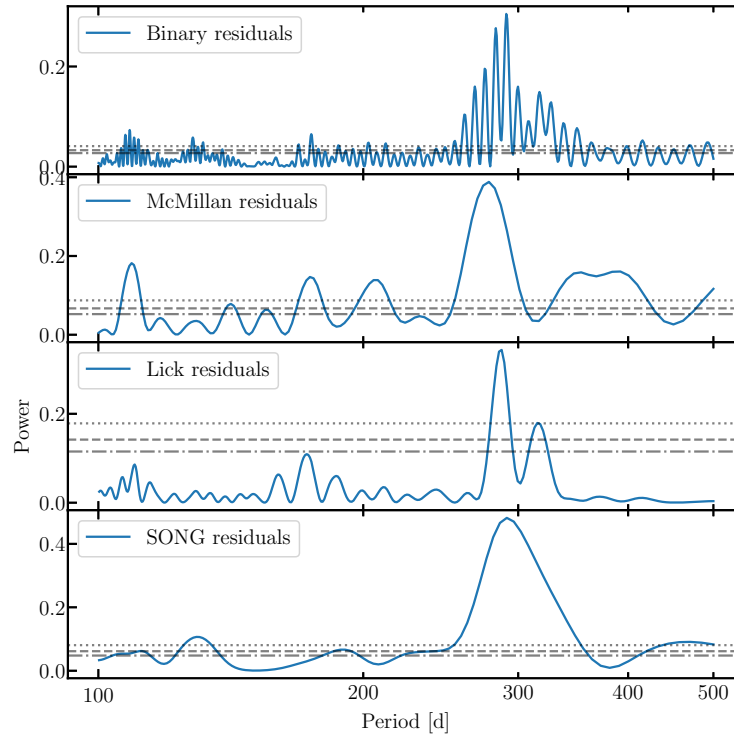


FIGURE 4.3: GLS periodograms of the residuals of the RV measurements after removal of the binary signal. The periodogram of the combined data sets (top panel) shows two peaks with similar heights at ~ 282.5 d and ~ 290.8 d. The gray dash-dotted, dashed and dotted lines represent false-alarm probability (FAP) levels of 5, 1 and 0.1%, respectively.

even larger than 3σ . The change of the period over time therefore seems to be quite significant, and we examine that behavior more closely in Section 4.4.4.

The GLS periodogram of the combined data sets shows a forest of peaks around periods of 300 d, with the strongest one at 291 d. The side peaks can at least partly be attributed to the sampling of the data: Between the individual data sets there are gaps of several years without any data, which can produce such patterns (see e.g., VanderPlas, 2018).

4.4.3 Fitting a double-Keplerian model to the data

We first investigate the hypothesis that this short-period signal is caused by a possible additional companion of planetary nature in the system orbiting the main component. Therefore we use a double-Keplerian model to constrain the orbits of the known stellar component and the possible second companion simultaneously. The fit then incorporates 13 free parameters: RV semi-amplitude K , orbital period P , eccentricity e , argument of periastron ω , and mean anomaly M_0 for each of the companions, as well as zero-point offsets for the McMillan, Lick and SONG data sets, respectively. Again we estimated the errors through the MCMC method, similar to the single Keplerian model, but now with 52 walkers. The 1σ -intervals of the 62 535 samples around the mean values of the posterior distribution serve as our uncertainties.

The best-fit results from the model and the mean orbital parameters from the MCMC posterior distribution are shown in Table 4.2 (columns “Double-Keplerian”), and




FIGURE 4.4: *Top panel:* RVs of the three modern data sets McMillan, Lick and SONG, phase-folded by the best-fit period of the putative planet $P = 291.2$ d, are shown in symbols along with the RV curve of the best fit as a black line. *Bottom panel:* Residuals of the individual data points.

they agree within the errors. The parameters of the stellar companion only change slightly as compared to the single-Keplerian model. The best-fit orbital period of the putative planet is 291.2 d, and it has a low eccentricity of ~ 0.173 . Using the primary mass of $M_{\star} = 1.103 M_{\odot}$ puts the minimum mass of the planet at about $1 M_{\text{Jup}}$, and its semi-major axis at 0.89 AU.

Without taking any jitter into account, the χ_{red}^2 of the best fit is 37.49. When we fit our Lick data set alone, letting all parameters vary, a jitter value of 17.4 m s^{-1} brings the χ_{red}^2 of that fit down to unity. This value is larger than the jitter estimate from the asteroseismic measurements (see Section 4.3), which is no surprise as the Lick data set covers a much longer baseline. Stellar jitter stems from a number of stellar phenomena, and oscillations are just one of those – others are granulation, stellar rotation, and magnetic cycles (Dumusque, 2016; Dumusque et al., 2017). These jitter sources act on considerably different time scales (few hours to several years), which might well explain the discrepancy between the jitter estimates for the whole Lick data set (covering 11 years) and the asteroseismic data (~ 1 month). For consistency, similarly as for the Lick data we also fitted only the SONG measurements, but left out the asteroseismic data, which resulted in a jitter estimate of 15.0 m s^{-1} to achieve $\chi_{\text{red}}^2 = 1$. This value falls in-between the jitter estimates from the longer-baseline Lick data and the shorter-baseline asteroseismic measurements.

We can also compare these results to the color-dependent distribution of jitter estimates found for the stars in our own K-giant sample (see Frink et al., 2001; Trifonov et al., 2014): For the color of ϵ Cyg, $(B - V) = (1.04 \pm 0.01)$ mag, the measured values scatter from about 10 to 30 m s^{-1} , so all of the jitter results from above lie within the distribution. In the end we decided to use a fixed jitter value of 20 m s^{-1} for all our

Keplerian and dynamical fits in this analysis. This corresponds roughly to the mean of the above-mentioned jitter distribution for stars of similar color ($B - V$), and allows to account for additional jitter sources that only come into play over the 32-years long baseline of the three datasets McMillan, Lick and SONG combined. The double-Keplerian fit to all these measurements then delivers a χ_{red}^2 of 1.21, which is a clear improvement to the single-Keplerian model with $\chi_{\text{red}}^2 = 1.93$ from Section 4.4.1. The combined rms of all data sets is 23.85 m s^{-1} , which corresponds roughly to our adopted jitter value.

Figure 4.4 shows the RVs of the three data sets along with the best Keplerian fit to the data phased to the best-fit period $P = 291.2 \text{ d}$ (top) and the residuals of the RVs from the fit (bottom). While many of the data points lie close to the best-fit solution, there are a number of McMillan and Lick measurements that fall far away from the curve. They form an additional “branch”; they are systematically shifted in phase and offset in vertical direction with respect to the fit. The amplitude of the best-fit curve is a little too small to perfectly match the distribution of the SONG RVs, but overall the SONG data set is much better represented by the model than the other two. The fact that the SONG data dominate the fit is not surprising as this is the largest data set with the smallest measurement uncertainties. However, the outliers especially in the McMillan and Lick data sets show that the double-Keplerian model does not provide a satisfactory fit to the combined data. A possible explanation for that would be a temporally changing RV signal; as the Keplerian model assumes undisturbed orbits, changes of the Keplerian elements are not incorporated.

4.4.4 Investigating the temporal evolution of the Keplerian signal

To further investigate the problems of the Keplerian model, we examine the three data sets individually by fitting a double-Keplerian model to each one of them separately and comparing the resulting parameters. By this we are hoping to find out whether a changing RV signal is the cause for the bad fit of the combined model. In a second step, we split each data set into shorter sections and fit each of these separately, in order to follow changes of the Keplerian elements over even shorter times. It is important to assure a large number of data points that is more or less evenly sampled as well as a time coverage that is longer than the period of the signal for each of the sections. The McMillan data set covers 5 years with 213 measurements, but the measurement uncertainties are comparatively large (12 m s^{-1}), and the observations were mostly taken in chunks within a few days, leaving gaps of several days or weeks in between. This leads to a potentially poor phase coverage of the period, which is why we split the McMillan data set into only two sections of 106 and 107 data points, respectively. The Lick observations cover the longest time span of the three data sets used here, but also contain the smallest number of measurements. Therefore we also split them into two sections of 55 and 54 data points, respectively. The SONG data set in contrast consists of a large number of measurements with small uncertainties, which allowed us to split it into three sections: The first consists of 46 data points and covers the time right before the very quick RV change during periastron passage of the close stellar companion occurs; the second is made up of the 128 data points taken during the close approach of the two stars; and the third consists of 54 data points and focuses on the time after periastron passage.

Each single fit to a complete data set incorporates 11 parameters, two times five Keplerian elements for the stellar and the planetary companion, and one RV zero

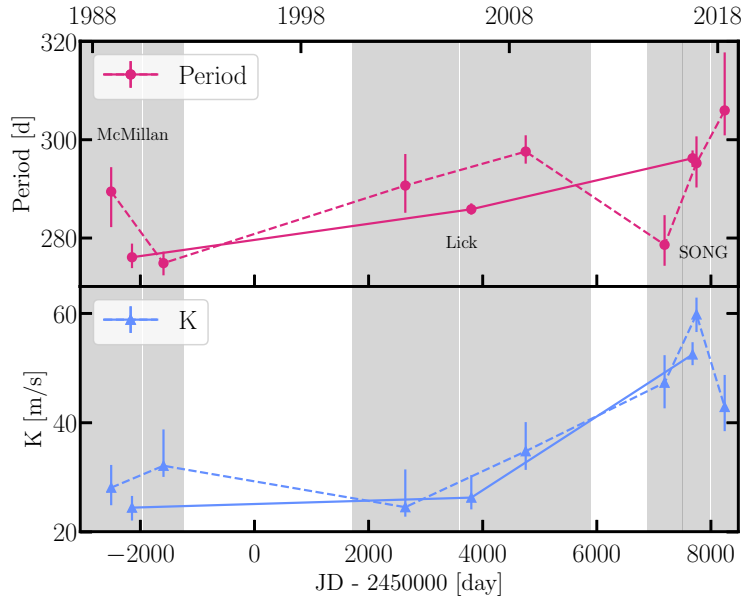


FIGURE 4.5: Orbital period P (top) and semi-amplitude K (bottom) of the presumed planet, derived from double-Keplerian fits to the individual data sets (solid lines) and to shorter sections of the data (dashed lines). The gray-shaded areas denote the time spans from the first until the last measurement within each section.

point. When fitting the shorter sections of the data sets however, the long-period orbit of the spectroscopic companion is not constrained sufficiently, due to the smaller number of data points and shorter time spans; therefore, in the models for the individual sections, we left the Keplerian parameters of the stellar companion fixed at the best-fit results from the whole respective data set. This leaves 6 free parameters for each fit to a section: 5 Keplerian elements for the planetary companion and one RV zero point. Uncertainties on the free parameters were again computed from an MCMC analysis for each data set and each section, in analogy to the single and double-Keplerian models to the combined data set.

Each model to a subset of the full data, either individual data set or even shorter section, results in a much better fit than the complete double-Keplerian model, with only very few random outliers. The χ_{red}^2 for the McMillan, Lick and SONG data sets are 0.96, 0.77, and 0.47, respectively. Figure 4.5 shows the evolution of the orbital period P and the semi-amplitude K of the short-period signal over time: The solid lines depict best-fit results of the individual data sets, while the dashed lines denote the results from fitting the sections. In both cases distinct changes of the parameters over time become apparent: The period clearly follows the pattern already observed in the GLS periodograms of the three data sets, changing from 278 d in McMillan over 286 d in Lick to 296 d in the SONG data. The fits to the shorter sections follow a similar trend, but show more scatter. Some of that variance might be explained by the fewer measurements within each section and the shorter time spans covered by the data. Especially in the case of the SONG data set, the first and the last section only contain around 50 measurements, and none of the three sections covers two full periods of the short-period signal.

The semi-amplitude K of the signal in contrast stays mostly constant for McMillan and Lick around a value of $26 \sim 27 \text{ m s}^{-1}$, before it drastically increases to 53 m s^{-1} in the SONG data, effectively doubling. Again, the fits to the sections show some

more scatter but generally follow the same trend. Interestingly, of the three SONG sections the second one, which covers the time around periastron, results in a much larger K than the other two. This might hint at an actual physical process that increases the semi-amplitude exactly during periastron passage, but could also be simply explained by the fitting problems described above. According to our calculations, for the edge-on configuration the gravitational redshift z_{grav} only changes by 0.6 to 0.7 m s^{-1} over the eccentric orbit of the stellar companion, so it fails as an explanation for the varying short-period semi-amplitude.

4.4.5 Fitting a dynamical model to the RV data

One possibility to explain the changes of the Keplerian elements over time would be the gravitational interaction of the planet in question with the close stellar companion: The double-Keplerian model assumes undisturbed orbits around the center of mass and neglects any other forces that might be present. However, for two bodies in such close proximity as the planetary and stellar companion in our model, and with such high masses, this assumption certainly does not hold true anymore, and the orbits of the two bodies would be subject to considerable changes. This is even more true because our total observational time span covers more than 30 years by now, over which any evolution of the orbits can be traced. Furthermore, the SONG measurements were taken around the periastron passage of the close stellar companion, where all three bodies come closest to each other and the biggest changes can be expected.

To take this effect into account, we attempt to better fit the data using a fully dynamical model, as done before for HD 59686 in Trifonov et al. (2018). It consists of a modified version of the Bulirsch-Stoer N -body integrator in the *SWIFT* package, that outputs the RV data of the primary component in the system, and a Levenberg-Marquardt minimization algorithm. As for the double-Keplerian model, the free parameters comprise the RV semi-amplitude K , orbital period P , eccentricity e , argument of periastron ω , and mean anomaly M_0 for each of the companions, as well as zero-point offsets for the McMillan, Lick and SONG data sets, respectively.

Additionally, in a dynamical model the data can potentially constrain the inclinations of the orbits relative to the sky plane $i_{1,2}$ and the longitudes of the ascending nodes $\Omega_{1,2}$ of the close stellar companion (subscript 1) and the putative planet (subscript 2). However, it is not feasible to keep both parameters free in our case as there are many local minima in the $\log(\text{likelihood})$ plane, and we do not have any independent knowledge of the true orientation of the orbits that we could use to find suitable starting parameters. Therefore, we applied our model only to strictly coplanar prograde orbits, where $i_1 = i_2 = 90^\circ$ and $\Omega_1 = \Omega_2 = 0^\circ$, and strictly coplanar retrograde orbits, that is, $i_1 = i_2 = 90^\circ$ and $\Omega_1 = 0^\circ, \Omega_2 = 180^\circ$.

Our best dynamical models however show no improvements to the double-Keplerian fits, with a χ_{red}^2 of 1.42 and 1.23 for the prograde and retrograde configuration, respectively. Both models are still left with many systematic outliers from the best-fit curve, so it seems as if the changes of the Keplerian parameters over time cannot be fully accounted for by the expansion to a dynamical approach. When integrating the two models further over time, with a modified version of the Wisdom-Holman N -body integrator in the *SWIFT* package, we find stability time of only 85 and 454 yr for the prograde and retrograde configuration, respectively, after which the planetary eccentricity becomes larger than 1 and the planet is kicked out of the system. The

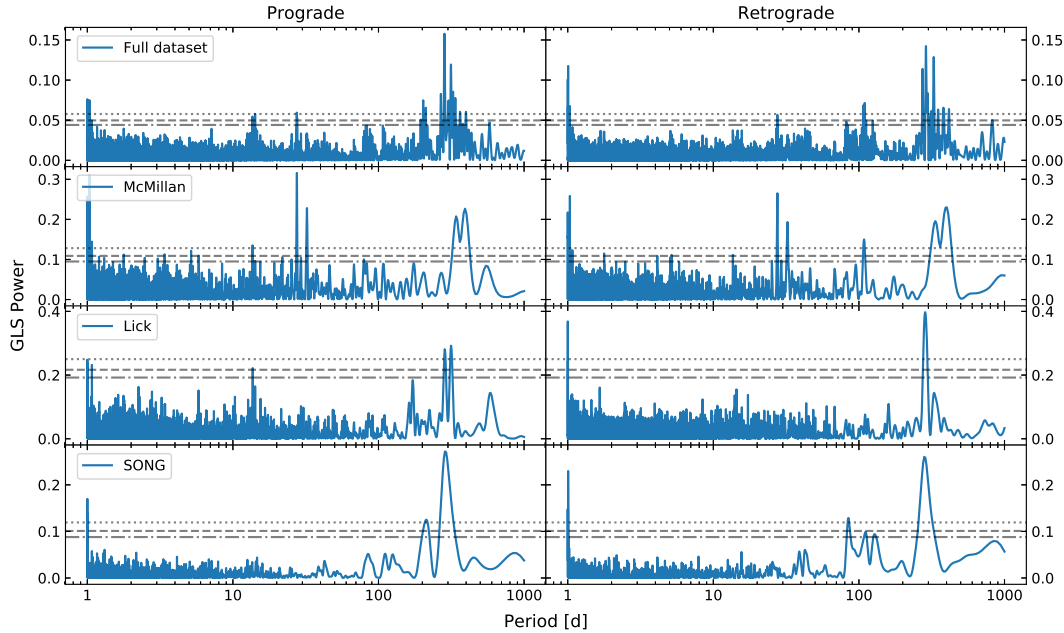


FIGURE 4.6: GLS periodograms of the residuals of the best dynamical models for the prograde (left) and retrograde configuration (right). The top-most plot shows the periodograms of the residuals of the combined data, the bottom three of the individual data sets. FAPs of 5, 1 and 0.1% are indicated by the gray dash-dotted, dashed and dotted lines, respectively.

long-term survival of the putative planet therefore seems highly questionable, and we present a more in-depth analysis of the dynamical stability in Section 4.5.

Figure 4.6 shows the GLS periodograms of the residuals of the dynamical models and the RVs, for both the prograde and retrograde configuration and for all individual data sets as well as the combined data. In all periodograms there are still peaks left that exceed the FAP level of 0.1% and mostly lie at periods close to the fitted orbital periods just below 300 d, which illustrate the above-mentioned systematics in the residuals. It is evident that the dynamical models do not fully describe the signals present in the RV data.

4.4.6 The co-orbital scenario

In this section we investigate the source of the significant peak of the RV residuals after subtracting the possible planet signal in the context of the potential presence of a co-orbital body. Co-orbital configurations have not yet been detected outside of the Solar System although several candidates have recently been published (of particular interest is the case of TOI-178, Leleu et al. 2019; but see also Hippke and Angerhausen 2015, Janson 2013 or the TROY project⁶ by Lillo-Box et al. 2018b; Lillo-Box et al. 2018a). Interestingly, co-orbital planet pairs are stable under a very relaxed condition developed in Laughlin and Chambers (2002), stating that such configurations would remain long-term stable as long as the total mass of the planet and its co-orbital companion is smaller than 3.8% of the mass of the star.

⁶www.troy-project.com




FIGURE 4.7: Lomb-Scargle periodogram of the radial velocity residuals (blue) and the corresponding demodulated radial velocities (red). The red vertical dotted lines correspond to the range of allowed libration periods for this star-planet system based on a minimum ($m_t = 0$) and a maximum ($m_t = m_p$) mass for the co-orbital body.

The motivation of this analysis for our system comes from the theoretical study of Leleu, Robutel, and Correia (2015), who derived the implications on the RV periodogram of a co-orbital planet. In such case, the long-term libration of the co-orbital motion would introduce two signals in the periodogram at $n \pm \nu$, where n is the frequency of the orbital period and ν corresponds to the libration frequency. However, the sampling and precision of the data can hide these two signals in the periodogram. Leleu, Robutel, and Correia (2015) propose a technique to enhance their detection through the so-called demodulation technique: The RV residuals are convolved with a sinusoidal function of same frequency as the carrier (the Keplerian frequency of the planet), reproducing the peak at the libration period. The libration period depends on the Trojan, planet and star masses so a minimum ($m_t = 0$) and maximum ($m_t = m_p$) libration period can be estimated. We applied this technique and obtained the periodogram of the demodulated RVs, which is shown in Figure 4.7. The result displays a peak in the expected period range of possible libration frequencies, which opens up the possibility of a librating co-orbital body to be the cause of these RV residual peaks. A more in-depth analysis is needed to properly investigate this possibility, which is beyond the scope of this paper. It is also necessary to point out that the issue of instability of a single circumprimary planet, which we examine in depth in Section 4.5, potentially also poses a problem for the co-orbital scenario; further investigations into this direction therefore require a thorough dynamical study.

4.5 Dynamical stability analysis

4.5.1 Theoretical considerations about the orbital configuration

An S-type planet around ϵ Cyg A would be subjected to substantial gravitational forces from the close stellar companion, which could potentially alter its orbit considerably over time, as is explained above in Section 4.4.5. Even though the semi-major axis of the orbit of the stellar companion is $a_1 = 15.8$ AU, due to its high eccentricity the periastron distance is only $q_1 = 1.11$ AU, which is just 0.22 AU larger than the proposed semi-major axis of the planet, $a_2 = 0.89$ AU. According to Hamilton and Burns (1992), the Hill radius of the stellar companion can be approximated as

$$r_H \approx a_1(1 - e_1) \sqrt[3]{\frac{m_1}{3m_0}} \approx 0.267 \text{ AU}, \quad (4.1)$$

with $m_1 = 0.265 M_\odot$ being the minimum mass of the stellar companion, $m_0 = 1.103 M_\odot$ the mass of the primary, and a_1 and e_1 the semi-major axis and the eccentricity of the orbit of the stellar companion. For the best coplanar double-Keplerian model the orbital path of the putative planet thus would pass through the Hill sphere of the close stellar companion during its periastron passage, making its long-term survival very questionable. For low inclinations and thus higher masses of the stellar companion, the problem becomes even worse as the Hill sphere becomes larger. However, despite the small distance of the orbits during periastron, the two companions could in principle always be much farther apart than the Hill radius of the stellar companion, as they do not necessarily have to have the same true anomalies at the time of closest approach.

We know of a number of extreme multi-companion systems today that are dynamically stable because their orbits are locked in very specific configurations, for instance, secular alignments of their periastra and/or mean-motion resonances (MMR), such as HD 59686 (Trifonov et al., 2018) or HD 82943 (Tan et al., 2013). In order to gain a better understanding of the possible orbital configurations of the ε Cyg system, we ran a comprehensive dynamical analysis: Using a modified version of the Wisdom-Holman N -body integrator in the *SWIFT* package, we tested the temporal evolution of many different orbital configurations by varying the starting parameters in a broad range around our best-fit results. First we tested both prograde and retrograde coplanar orbits, with a special focus on configurations that are locked in secular apsidal alignment or mean-motion resonance. In a second step we extended our analysis to mutually inclined orbits, for which we concentrated on regions of the parameter space near the fixed point of the Kozai-Lidov mechanism as these would offer the greatest chances for long-term stability (for more details see e.g., Kozai, 1962; Lidov, 1962; Lithwick and Naoz, 2011).

4.5.2 Stability analysis of coplanar orbits

We set up a system of two bodies, using the derived masses of the primary and secondary stellar components as well as the semi-major axis of the secondary from Section 4.2 and 4.4.3, and inserted a variety of massless test particles around the best-fit orbital solution of the planet. As starting parameters, we chose the periastron of both the test particles and the stellar companion to be at 0° (so $\Delta\omega = 0^\circ$), and the mean anomalies of the test particles to be 0° , while that of the stellar companion was set to 180° . The semi-major axes of the test particles were varied within a range of 0.06 AU around the best-fit solution of the planet, and their eccentricities varied between 0 and 0.8. In order to understand the influence of the high binary eccentricity on the inner orbits, these simulations were repeated for eccentricities of the stellar companion ranging between 0.8 and the best-fit value of 0.93. The whole analysis was done both for prograde and retrograde orbital configurations, and a system was considered stable if the test particle stayed within a critical distance interval $\{a_{c,1}, a_{c,2}\}$ from the primary component for the whole integration time of 10^4 orbits of the stellar companion ($\sim 5.4 \cdot 10^5$ yrs). The inner boundary $a_{c,1} = 0.22$ AU was chosen to be the distance at which one orbit of a test particle is only resolved by 10 integration time steps (the time step of the simulation being 3.4 d), the outer boundary $a_{c,2}$ corresponds to the semi-major axis of the stellar companion.

As was expected from our considerations in Section 4.5.1, the system generally shows a very high degree of chaotic behavior. The most stable coplanar configurations were achieved for retrograde MMR or secularly aligned orbits with test particle eccentricities around 0.4. A 1: n MMR is established for systems where at least one of the MMR angles

$$\lambda_2 - n\lambda_1 + (m - 1)\omega_2 - (m - n)\omega_1, \quad (4.2)$$

librates around a constant, with $\lambda_{1,2}$ being the mean longitudes of the stellar companion and the test particle, respectively. For prograde orbit, n is positive and $m = 1, \dots, n$. For retrograde orbit, n is negative and $m = 1, \dots, |n| + 2$. For secular alignment the angle between the pericenters of the two orbits,

$$\Delta\omega = \omega_1 - \omega_2, \quad (4.3)$$

librates around zero. For all our simulated systems we checked whether these requirements were fulfilled; we were thus able to identify the MMR and secularly aligned configurations. Nevertheless, even for those configurations general long-term stability only occurred for eccentricities of the stellar companion being smaller than 0.84 (for prograde orbits) or smaller than 0.9 (for retrograde ones). At the best-fit binary eccentricity of 0.93 the instability timescale is on the order of only 10 binary periods ($\cong 500$ yrs) in the retrograde case and on the order of just 1 binary period in the prograde case. Altering the semi-major axes of the test particles within the range defined above has barely any effect on the stability.

We redid the simulations for 1000 test particles with initial semi-major axes evenly distributed between 0.65 and 1.1 AU and fixed eccentricities at 0.1, while varying the binary eccentricity between 0.8 and 0.93 in steps of 0.01, to illustrate our general findings for the pro- and retrograde case in Figure 4.8. Here it becomes clear that orbiting closer-in to the primary component improves the stability of the test particles, but the position of the best fit (indicated by the white diamond) is far away from the stable region even in the retrograde configuration. With the eccentricity of the orbit of the stellar companion being our best-constrained parameter ($e_{AB} = 0.9295 \pm 0.0003$ in the single-Keplerian model), there is no doubt that the putative planet falls into a highly unstable region of the parameter space.

4.5.3 Stability analysis of mutually inclined orbits

From our considerations in Section 4.4.1 we know that the inclination of the binary orbit must be larger than 14° , but apart from that we do not have any constraints on the orbital inclinations. Therefore, in order to test the stability of mutually inclined orbits, we decided to focus on the region of possible Kozai librations within the parameter space, which we deem most promising to guarantee long-term stability. We ran simulations for retrograde orbits only (these being more stable than prograde orbits), varying the mutual inclination Δi between 90° and 180° in steps of 5° , with $\omega_2 = 90^\circ$ (this corresponds to the stable Kozai libration regime) and $\Delta\Omega = 0^\circ$ (to avoid close encounters). A smaller time step of 0.35 d was chosen, which puts the inner boundary (for at least 10 time steps per orbit) at $a_{c,1} = 0.055$ AU and therefore allows the possibility of highly eccentric orbits of the test particles, as are expected to occur in Kozai oscillations. The outer boundary $a_{c,2}$ again is equal to the semi-major




FIGURE 4.8: Survival times for test particles placed into the ϵ Cyg stellar binary, in prograde (top plot) and retrograde coplanar configuration (bottom). The eccentricity of the orbit of the stellar companion and the initial semi-major axes of the test particles were varied in order to find the stability boundary. The position of the best-fit solution from Section 4.4.3 is plotted by the white diamond; its size is at least by a factor of 10 larger than the uncertainties on the Keplerian parameters.

axis of the stellar companion. This time we did not test different binary eccentricities, but only used the best-fit value of 0.93 in our simulations.

The most stable solutions were found for mutual inclinations $135^\circ < \Delta i < 140^\circ$, but even for those configurations the instability timescale is on the order of 10 binary periods ($\cong 500$ yrs), after which the test particles are ejected from the system; only a small percentage of systems survive longer than 100 binary periods, but fully stable regions do not appear in the parameter space. We also did not find any signs of Kozai oscillating orbits; if they exist, they must be very short-lived and are unlikely to provide long-term stable solutions.

4.6 Possible alternative explanations for the short-period RV variations

4.6.1 Hierarchical triple

In a series of publications (Morais and Correia, 2008; Morais and Correia, 2011; Morais and Correia, 2012) it was shown that RV variations of a star can be caused by two bodies orbiting each other while orbiting that star, as is the case in hierarchical triples. In Morais and Correia (2012) the authors develop a secular theory for such systems and derive an expression for the precession rate of the orbit of the companion pair around the main component. Furthermore they suggest this effect as an alternative explanation to the planet hypothesis for the ν Octantis system (Ramm et al., 2009; Ramm et al., 2016), which is known to be a close single-line spectroscopic binary where the two stellar components orbit each other with a period of 1050 d. In addition to the long-period RV signal caused by the stellar companion, Ramm et al. (2009) observed short-period RV variations with a period of 417 d, which they interpreted as a possible planet with a minimum mass of $2.5 M_{\text{jup}}$. Stable orbital configurations have been published by Ramm et al. (2016), but previously the stability of the proposed planet seemed questionable. Morais and Correia (2012) argue that the short-period RV variations could also be caused by a very close pair orbiting the main component, meaning that the stellar companion ν Octantis B would itself be made up of two bodies. They also show that the precession rate of the long-period orbit as measured from the RVs is in agreement with their hypothesis.

Here we aim to test the hierarchical triple configuration for the ϵ Cyg system. Unfortunately, our RV data do not even cover one full phase of the long-period orbit, which makes any measurement of the precession rate very unreliable. Therefore we chose a different approach: Using a modified version of the Wisdom-Holman N -body integrator in the *SWIFT* package we modeled hierarchical systems and tried to reproduce the RV variations observed in our data. As the integrator was originally created to model the Solar System, it only allows to set up orbits of bodies around the central star of the modeled system, not of two outer bodies orbiting each other whilst orbiting together around the central star. By changing from barycentric to Jacobi coordinates however, we are able to choose one of the two components of the companion pair (denoted as m_1) as central object. In this transformed setup, the other component of the companion pair (m_2) as well as the main component (m_0) are then orbiting m_1 .

We used a brute-force technique to try to find systems that produce RVs similar to what we observe. The long-period orbit of the companion binary is effectively given by our solution in Section 4.4.1; as described above, these (Jacobian) orbital elements were now ascribed to the main component m_0 . We then replaced the single companion by a pair of two bodies m_1 and m_2 ; m_1 was chosen as the central star in the system and we systematically varied the Keplerian elements of the orbit of m_2 about m_1 as well as the masses of the two bodies. After setting up the system, we let it run for 200 years with a time step of 0.5 days and recorded the RVs of the main component.

In these simulated RVs the long-period orbit of the m_1 - m_2 pair around the main component m_0 was in general clearly visible and closely resembled our data (see top plot of Figure 4.9). We then fitted a Keplerian model to that RV variation and subtracted it from the RVs. By computing a GLS periodogram of the residuals, we




FIGURE 4.9: *Top*: Simulated RVs of the main component of a hierarchical triple, with the orbit of the companion pair around the primary chosen according to our solution for the close stellar companion to ϵ Cyg A. *Middle*: Residuals of the simulated RVs after subtracting the high-amplitude, long-period RV signal modeled by a Keplerian orbit. *Bottom*: Zoom in on the time around periastron passage of the eccentric long-period orbit.

searched for short-period variations. While in most simulated systems the binarity of the companion was too small to produce any considerable signature, or one of the components was lost within a very short time, some simulated configurations did show peaks at the periods in question. These were all systems with a comparably large separation a_1 between the two bodies m_1 and m_2 , and all of them showed signs of long-term instability. However, the short-period variations were always limited to the time around periastron passage of the m_1 - m_2 binary around the main component.

Figure 4.9 shows a characteristic example, for a system with $a_1 = 0.3$ AU, $m_1 = 0.16 M_\odot$ and $m_2 = 0.12 M_\odot$: After subtracting the long-period orbit, the residuals show variations with an amplitude of 30 m s^{-1} near periastron passage (middle plot), with a quasi-periodic nature (lower plot). During that time, all three bodies are very close to each other (the periastron distance of the m_1 - m_2 system to m_0 is approximately 1.2 AU), so the varying quadrupole (and higher) moments of the companion pair induce a large motion on the main component. But this signal quickly decays to amplitudes below 1 m s^{-1} further away from the time of periastron passage. In contrast, we clearly observe RV variations with amplitudes around 25 m s^{-1} in the McMillan and Lick data of ϵ Cyg far from periastron passage.

We conclude that while hierarchical triples may for some systems be a valid alternative explanation for the observed RV signals, this is not a viable solution for the observed RVs of ϵ Cyg. In contrast to the case discussed in Morais and Correia (2012), the long-period orbit of ϵ Cyg is far too eccentric to be compatible with large-amplitude short-period RV variations over the whole orbit induced by a companion

pair.

4.6.2 Stellar spots

Many stars exhibit stellar activity, that is, stellar spots, flares etc., which can mimic a planet in RV observations by blocking or enhancing light from parts of the surface of the stellar photosphere. This effect can become especially problematic when searching for planets around M dwarfs, which are often quite active and therefore can show activity-induced RV signals of several m s^{-1} (see e.g., Reiners et al., 2010; Barnes, Jeffers, and Jones, 2011; Tal-Or et al., 2018). By checking time-series of activity indicators, such as photometry, bisector spans of the absorption lines, or the depths of the $\text{H}\alpha$ and Ca II lines of these stars, and comparing them to the RV time series one can distinguish between activity-induced and planetary signals. In addition, the so-called chromaticity, that is, wavelength-dependence, of RV signatures caused by stellar activity can serve as a tracer (Reiners et al., 2010; Zechmeister et al., 2018).

Evolved stars in contrast are known to be less affected by large-scale activity. Our sample stars were chosen specifically because they were comparably quiet stars in the HIPPARCOS photometry measurements. The top panel of Figure 4.10 displays the HIPPARCOS data of ϵ Cyg, which have been averaged whenever there was more than one data point within a day. The time series largely overlaps with the second half of the McMillan RVs, which are displayed in the plot underneath. Apart from some outliers, most of the HIPPARCOS measurements lie very close around the mean of 2.643 mag with an rms scatter of 3.7 mmag. This translates to a relative change in flux of about 0.34%.

Nevertheless, when taking a periodogram of the HIPPARCOS photometry, there is a peak around a period of 300 d (see bottom panel of Figure 4.10), which matches the observed period of the RV signal (indicated by the red line) quite well. This might hint at a signal induced by spots in the photometry data and would mean that the rotation period of the star would be approximately 300 d or an integer multiple of that. To our knowledge, there are two modern measurements of the projected rotational speed of ϵ Cyg: Massarotti et al. (2008) determined it to be $v_{\text{rot}} \sin i_{\star} = 1.2 \text{ km s}^{-1}$, while the analysis of Gray (2015) resulted in $v_{\text{rot}} \sin i_{\star} = 1.0 \pm 0.2 \text{ km s}^{-1}$ (where i_{\star} denotes the inclination of the stellar spin axis). Assuming $i_{\star} = 90^{\circ}$ these values correspond to rotation periods of 461 and 554 ± 111 d, so in-between 1 and 2 times the period that is visible in our data. Taking the uncertainty in the measurement of Gray (2015) into account as well as the fact that for smaller values of the unknown inclination i_{\star} the rotation periods become shorter, it seems possible that the peak in the periodogram of the photometry actually traces the rotation of the star.

To further investigate this we searched the literature for additional photometric time series, but only found two more data sets: The BRITE nano-satellites BRITE-Toronto (BTr) and UniBRITE (UBr) (Weiss et al., 2014) observed ϵ Cyg in the years 2014 and 2015, which coincides with the beginning of our SONG measurements (see Figure 4.11, top). Unfortunately, the two data sets only cover 50 and 146 d, respectively, which is much shorter than one full phase of the short-period signal and therefore does not help to shed more light on the question whether that signal has a direct counterpart in photometry (see Figure 4.11, top). The rms scatter of the data sets BTr and UBr are 2.0 and 1.1 mmag, respectively, which is even lower than in the HIPPARCOS data. Interestingly however the periodograms of the measurements reveal

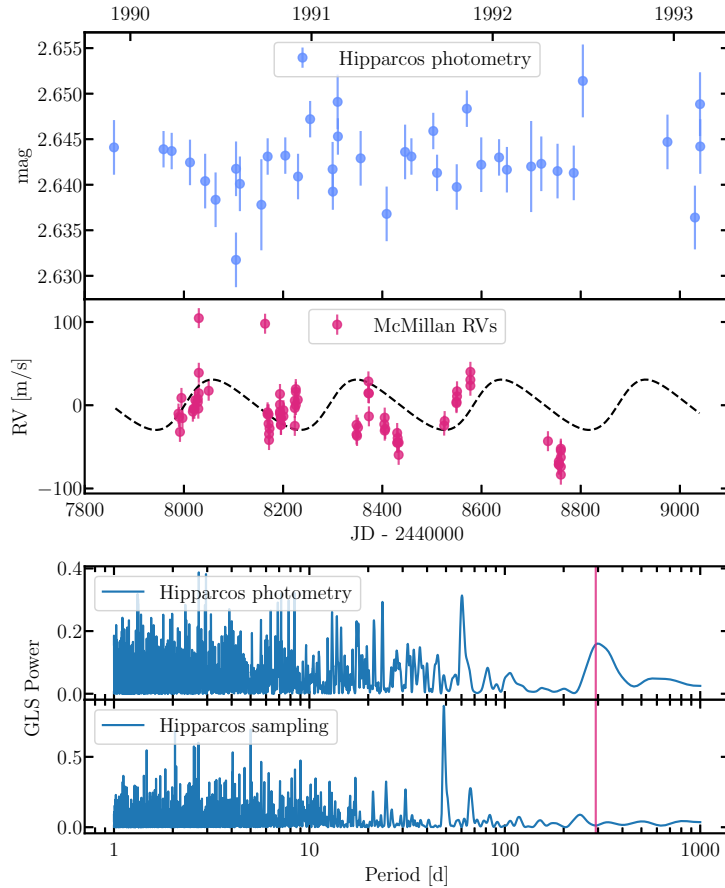


FIGURE 4.10: *Top*: HIPPARCOS photometry for ϵ Cyg, taken around the same time as the last RVs of the McMillan data set. The displayed photometry data are median values for all cases of multiple measurements within one day; the RVs are the residuals of the McMillan measurements after subtracting the long-period signal induced by the binary companion (see Section 4.4.1), and the dashed line denotes the best-fit model. *Bottom*: GLS periodogram of the HIPPARCOS photometry data and of the window function of the observations. The best-fit Keplerian orbital period of the signal of 291 d is marked with a red line.

some variability on even shorter periods, with a strong peak around 40 d in both data sets (see Figure 4.11, bottom). The HIPPARCOS periodogram seems to have a rough analog with a peak around 60 d, although in that case it might also be caused by the sampling frequency (see Figure 4.10, bottom).

Generally, it would be surprising if the observed short-period RV variations of ϵ Cyg were caused by spots, as it is an evolved star with a radius of about $11 R_{\odot}$. In order to cause RV variations on the order of the observed semi-amplitude of $20 \sim 30 \text{ m s}^{-1}$, stellar spots would have to be extremely large, which in turn should become clearly visible in photometry as well. To study the effect qualitatively, we used the online tool of the star spot simulator *SOAP 2.0* (Dumusque, Boisse, and Santos, 2014) to find a combination of spots that could possibly mimic the observed RV signal. Setting the rotation period of the simulated ϵ Cyg A to the observed modulation of 291 d, we found that a single spot would need to have a radius of $0.2 R_{\star}$ to induce a RV semi-amplitude of about 25 m s^{-1} , and the resultant flux change would be around 2.7%, much more than the scatter of 0.34% in the HIPPARCOS photometry. Using more than one single spot to create the same RV amplitude generally allows to decrease

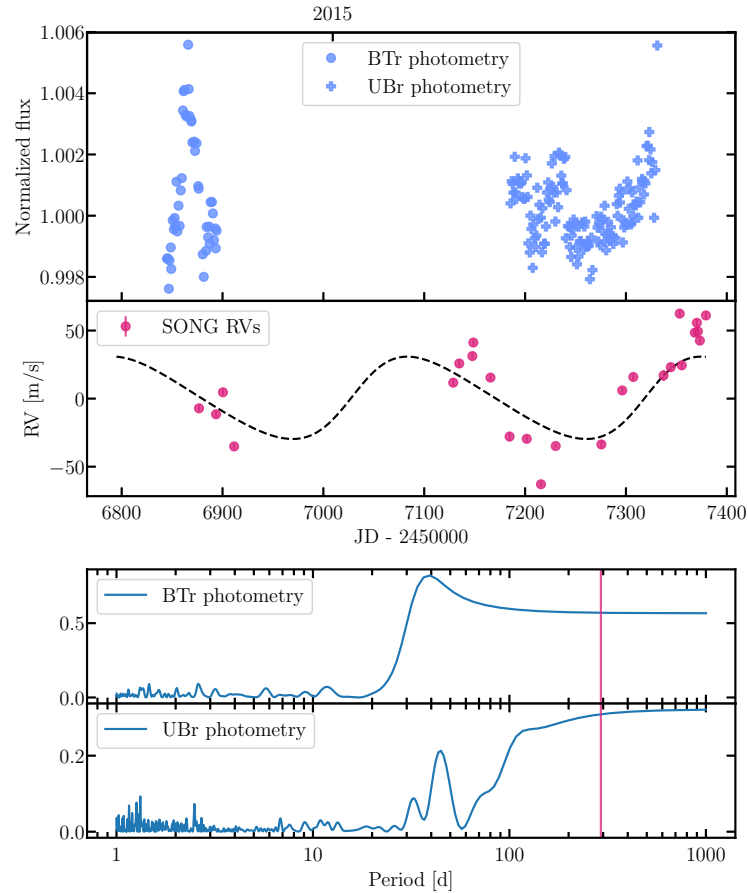


FIGURE 4.11: *Top*: BRITE-Toronto (BTr) and UniBRITE (UBr) photometry for ϵ Cyg, taken contemporaneously with the first RVs of the SONG data set. The displayed photometry data are median values for all cases of multiple measurements within one day, and both data sets have been normalized by their respective median values; the RVs are the residuals of the SONG measurements after subtracting the long-period signal induced by the binary companion (see Section 4.4.1), and the dashed line denotes the best-fit model. *Bottom*: GLS periodogram of the Brite photometry data sets. The best-fit Keplerian orbital period of the signal of 291 d is marked with a red line.

the individual spot sizes a little, but the photometric variability remains the same or even increases. Four spots with radii of $0.12 R_*$ each at very close longitudes and latitudes, for example, keep the RV semi-amplitude roughly constant but lead to a flux variability of 3.5%. Similarly, placing spots on opposite sides of the star, or assuming an inclination $i < 90^\circ$, only increases the predicted photometric variability.

We conclude that even though there are hints of a signature of the rotation of the star in the HIPPARCOS periodogram, the modeling suggests that spots are an unlikely cause of the observed RV variations.

4.6.3 Oscillatory convective modes

In many evolved stars with very high luminosities, photometric oscillations are known with typical periods between about 400 to 1500 d, but the causes of these oscillations are not fully understood (Wood et al., 1999; Hinkle et al., 2002; Wood, Olivier, and Kawaler, 2004). Those stars are classified as long secondary period (LSP) variables because their primary oscillation periods are much shorter, with the

period ratios between the secondary and primary oscillations ranging from about 5 to 13 (Wood et al., 1999). While the oscillations in most LSP stars have only been detected photometrically, Hinkle et al. (2002) and Wood, Olivier, and Kawaler (2004) also obtained RVs for some of these stars and found the long variations present in most cases, with amplitudes of a few km s^{-1} and periods consistent with the photometric variations. Saio et al. (2015) examine the observed periodic variations of the LSP stars and suggest so-called oscillatory convective modes, that is, nonadiabatic g^- modes in the deep convective envelopes of the stars, as a possible explanation.

However, all of the known LSP stars have luminosities of $L_\star > 300 L_\odot$ (Saio et al., 2015), whereas ϵ Cyg has a much lower luminosity of $L_\star = 57 L_\odot$. It is therefore not clear whether the mechanism described by Saio et al. (2015) might be responsible for the observed RV variations in this case, but it is certainly possible that LSP oscillations also exist in stars with lower luminosities, and that the distribution of known LSP stars is constrained by observational biases: Most of them have been found and characterized in the *OGLE* survey (see e.g., Udalski, Kubiak, and Szymanski, 1997; Soszyński et al., 2009), which is aimed at discovering microlensing events and focuses on very distant stars in fields toward the Galactic Center and the Large Magellanic Cloud; stars similar to ϵ Cyg therefore are probably just not bright enough to be observed with sufficient S/N in order to detect the photometric variations.

Two other examples of evolved stars whose RV variations might well be caused by this mechanism are γ Draconis (Hatzes et al., 2018) and Aldebaran (Reichert et al., 2019); both show signals with periods around 700 and 600 d, respectively, which could easily be confused with Keplerian signals if one just examined parts of the data. Similar to ϵ Cyg, in both cases only the very long time span of the collected RV measurements allowed the authors to discover the varying nature of the signal and therefore reject the planet hypothesis. As the luminosities of γ Draconis and Aldebaran are $L_{\text{fIDra}} = 510 \pm 51 L_\odot$ (Hatzes et al., 2018) and $L_{\text{Ald}} = 402^{+11}_{-10} L_\odot$ (Reichert et al., 2019), they fall comfortably inside the distribution of known LSP stars and therefore oscillatory convective modes seem to be a plausible explanation for the observed RV signals. With only these two examples of similar (non-)detection histories, it could be that ϵ Cyg is another example of that category, even though its stellar parameters might suggest otherwise.

4.6.4 Potential stellar oscillations through the heartbeat phenomenon

A different mechanism that is known to produce oscillations in some eccentric binary systems is the so-called heartbeat phenomenon, where a stellar companion passes close by the primary component and excites tides within the latter. With the *Kepler* mission alone, more than 150 of these systems have been observed (Kirk et al., 2016). They are characterized by a sudden increase in brightness of the primary star during periastron passage of the companion, due to the deformation of the star caused by the increased gravitational force, which is called the equilibrium tide. Furthermore, sometimes one or even several periodic luminosity-signals can be observed over the whole orbit, the so-called dynamical tides. These oscillate in characteristic stellar oscillation modes, which are excited through the regular close passage of the stellar companion (for a full theoretical description see e.g., Fuller, 2017).

All of the known binary systems with heartbeat variations are however much more compact than ϵ Cyg, with orbital periods typically shorter than 1 yr (Kirk et al., 2016).




FIGURE 4.12: Primary stellar radii over binary orbital periods for a number of heartbeat systems, with the orbital eccentricity color-coded. Triangles mark the main-sequence stars from Shporer et al. (2016), circles the RGB stars from Beck et al. (2014). For comparison ϵ Cyg has been added, denoted by a cross. The plot has been created in analogy to Beck et al. (2014).

Still, in an attempt to investigate this possibility further, we revisit two publications on heartbeat systems that have been monitored in RVs to better constrain their orbits: Shporer et al. (2016) analyze a sample of 19 main-sequence stars in heartbeat binaries, with orbital periods between 8 – 90 d, while Beck et al. (2014) examine 18 RGB heartbeat stars whose stellar companions orbit them with periods between 20 – 438 d. In analogy to Figure 18 from Beck et al. (2014), Figure 4.12 displays important orbital and stellar parameters of all systems from both samples. The plot shows a clear positive correlation between primary stellar radius and binary orbital period for the RGB stars, while the main-sequence stars seem to be scattered more or less uniformly at the small-radius, short-period end of the distribution. Generally, for a given primary stellar radius heartbeats may occur in systems with longer orbital periods P if the eccentricity e is also larger, as the periastron distance scales as $q \propto P^{2/3}(1 - e)$. ϵ Cyg however falls far away from all other systems at a much longer orbital period; even though its eccentricity is also larger than any other in the combined sample and its stellar radius rather large, it is questionable from this figure whether ϵ Cyg could be a heartbeat system.

Shporer et al. (2016) further investigate the ratio of the tidal forcing due to the stellar companion to the surface gravity of the primary, using the relation

$$\frac{F_{\text{tide}}}{F_{\text{gravity}}} = \left(\frac{GR_1M_2}{q^3} \right) \left(\frac{GM_1}{R_1^2} \right)^{-1} = \left(\frac{R_1}{q} \right)^3 \frac{M_2}{M_1}, \quad (4.4)$$

where $R_{1,2}$ and $M_{1,2}$ are the radius and mass of the primary and secondary component, respectively, G is the gravitational constant and q is the periastron distance of the two stars. To compute the masses of the secondary components, the authors used the mass functions from their RV analysis and assumed $\sin^3 i = 0.6495$; this value corresponds to the median of the $\sin^3 i$ distribution for the unknown inclinations and it was applied instead of the mean due to the very asymmetric shape of the distribution (see Shporer et al., 2016, for more information). We redid these calculations for both samples of Shporer et al. (2016) and Beck et al. (2014), using the stellar and orbital parameters from their works, as well as for the ϵ Cyg system (scaling its




FIGURE 4.13: Primary stellar radii (top) and tidal forcing ratio (bottom) over periastron distances in units of stellar radii (q/R_*) for heartbeat systems, consisting of the main-sequence star sample from Shporer et al. (2016) (green) and the RGB stars from Beck et al. (2014) (blue). ϵ Cyg has been added for comparison (orange). Symbol sizes are proportional to the stellar radii. The tidal forcing ratio is calculated as in Eqn. 4.4, and the large errors of the main-sequence stars stem from the high uncertainties on the radii R_1 and the masses M_1 of the primary components. The plot has been created in analogy to Shporer et al. (2016).

companion mass M_2 by the same inclination as used for the other systems). In Figure 4.13 we plot the primary stellar radius (top) and tidal forcing ratio as calculated in Eqn. 4.4 (bottom) over the periastron distance in units of the primary stellar radius (q/R_*) for each system; the stellar radii are also indicated by the sizes of the symbols in the plots.

Plotted against q/R_* , the stars of the two samples now fall much closer to each other, with most of them having periastron distances between 4 and $12 R_*$. The RGB heartbeat stars populate the area at small values of q/R_* , whereas the main-sequence stars extend the distribution to larger values of q/R_* . One main-sequence heartbeat system even has a $q \approx 24 R_*$ (this one being KIC 10334122), which is comparable to the periastron distance of ϵ Cyg, $22 R_*$. Due to its larger radius however, ϵ Cyg falls above the distribution of the samples of RGB and main-sequence stars.

Finally, as is to be expected from Eqn. 4.4, the ratio of the tidal forcing to the surface gravity of the primary shows a clear correlation with q/R_* . Most of the main-sequence and RGB stars fall between values of $4 \cdot 10^{-3}$ (when q/R_* is small) and $4 \cdot 10^{-4}$ (when q/R_* is larger), and the rms scatter of the two samples combined is $1.55 \cdot 10^{-3}$; the main-sequence system at $q \approx 24 R_*$ has a tidal forcing ratio around $7 \cdot 10^{-5}$. Calculating the same quantity for ϵ Cyg gives a value of $2.3 \cdot 10^{-5}$, which is not far from the main-sequence outlier. Given that ϵ Cyg lies so close to a known heartbeat system in these quantities makes it appear plausible that its primary star

ϵ Cyg A might also undergo tidally induced stellar oscillations. Additionally, it is important to keep in mind that the tidal forcing ratio scales with the companion mass M_2 , and that ϵ Cyg might therefore be even closer to the other stars in the case of a low inclination of the system. For example, for the lower limit on the inclination of $i = 14^\circ$, the tidal forcing would be $1 \cdot 10^{-4}$, which is greater than that of the main-sequence outlier and nearly within the rms scatter from the mean of the RGB and main-sequence distribution.

We also emphasize the fact that the variability of ϵ Cyg has been observed in RVs, whereas the systems from Beck et al. (2014) and Shporer et al. (2016) have all been discovered through their photometric variations in Kepler observations. Despite the fact that both these studies also monitored their systems in RVs to constrain their orbits, they do not mention any detection of RV variations corresponding to the heartbeat oscillations. This could have several reasons: They might not have picked up the signals in RVs due to sparse sampling or large errors of their measurements. Also some of their heartbeat systems barely seem to show any noticeable dynamical tides and mostly just equilibrium tides during periastron passage of the stellar companion, so there might not even be an RV signal to detect during most parts of the orbit. For ϵ Cyg in contrast we have a large number of RV measurements over most parts of the binary orbit, which clearly record a short-period signal that changes in period and semi-amplitude especially around periastron passage of the stellar companion. Given the differing detection method as compared to the other heartbeat stars, it is possible that ϵ Cyg experiences a similar phenomenon even though it does not follow all correlations exactly.

4.7 Summary & conclusions

Short-period RV variability of the spectroscopic binary ϵ Cyg had already been noticed in past publications (McMillan et al., 1992; Gray, 2015) and had been attributed to intrinsic stellar variations rather than an extrinsic source. We observe a similar RV signal in our measurements of the star at Lick and with SONG and further investigated the possibility of a planetary origin.

In a first step we determined the long-period binary orbit with very high precision by fitting a Keplerian model to the three RV data sets of McMillan et al. (1992), Lick and SONG. The results indicate the possibility of directly imaging the secondary component in the years 2021/2022. A direct observation would constrain the orbital inclination and the mass of the spectroscopic companion, possibly revealing whether it is a white dwarf or a main-sequence star, which could give insight into the past evolution of the binary system.

Next we computed GLS periodograms of the residuals after subtracting the Keplerian model from the RV data; they show clear signs of additional periodic signals at periods just below 300 d for all three data sets. We fitted the combined RV data with a double-Keplerian model to explore the possibility of an S-type planetary companion as cause for the short-period RV variations. The best fit leaves a considerable number of systematic outliers that look like a phase shift of the signal at some point in time or a change of orbital period. By modeling each data set, and smaller sections of the data sets individually, we revealed that the planetary period gradually increases over time, whereas the RV semi-amplitude mostly stays constant until it increases by a factor of 2 at periastron passage of the stellar companion. This indicates that a single circumprimary planetary companion is an unlikely explanation for the

observed RV signal, but the residuals of the double-Keplerian model could point toward a scenario with a second low-mass, co-orbital companion in the system. Using a fully dynamical model for the planetary companion, treating it as a test particle in the stellar binary system, does not help to solve the issue, as it cannot reproduce the systematic changes of the Keplerian elements, neither for the prograde nor for the retrograde case. Additionally, our stability analysis shows highly unstable behavior in large regions around the best-fit orbital configurations, thus strengthening the arguments against a planet.

Following Morais and Correia (2008), Morais and Correia (2011), and Morais and Correia (2012), we explored whether the RV variations could be due to an additional component not orbiting the primary, but the secondary star, making the system a hierarchical triple. Our models clearly show that in our case such a setup cannot reproduce the data. We therefore investigated possible intrinsic stellar origins of the signal, such as stellar spots. Data from HIPPARCOS suggest that ϵ Cyg A is a photometrically relatively quiet star, even though a periodogram of the measurements reveals a possible weak signal around 300 d, which is close to the period observed in RVs and could correspond to the rotation period of the star. If the RV variations were due to stellar spots, we would however expect the photometric variation to be much larger than the observed scatter.

Two other examples of ambiguous RV variations of giant stars, γ Draconis (Hatzes et al., 2018) and Aldebaran (Reichert et al., 2019), might belong to the class of LSP variables described by, for example, Hinkle et al. (2002), Wood et al. (1999) and Wood, Olivier, and Kawaler (2004), and Saio et al. (2015). As ϵ Cyg A is much smaller and less luminous than the known LSP stars, this mechanism can offer a satisfactory explanation in this case only if the distribution of LSP stars known today is incomplete due to observational biases.

A more likely, even though also not completely convincing solution is the possibility of ϵ Cyg being a heartbeat system, with the close stellar companion inducing tides in the primary component during each periastron passage that excite long-standing oscillations with the observed period. However, there are some distinct differences between the known heartbeat systems and ϵ Cyg, with the periastron distance of the latter being much larger. Also the question remains whether there are oscillation modes at periods around 300 d in HB stars that could easily be excited through this mechanism.

Generally, we showed that none of the phenomena discussed above offer a fully satisfactory explanation for the RV variability by themselves at the moment. Nevertheless, it is possible that a combination of some of the mechanisms could explain the observed RVs while also eliminating some of the problems encountered in the analysis. For example, the oscillatory convective modes might be present in ϵ Cyg A and, although not strong enough to cause recognizable RV variations by themselves in this type of star, could be excited regularly by the tidal interaction with the stellar companion. On the one hand, this would explain why ϵ Cyg A is so far away from the distribution of known LSP stars and still shows oscillations; on the other hand, the fact that these oscillation modes are present and probably easily excited offers a solution to the problem of the comparably large separation of the stellar binary when compared to other heartbeat systems. Similarly, Saio et al. (2018) show the presence of tidally excited Rossby waves (so-called r modes) in some heartbeat systems, which typically appear at periods slightly longer than the rotation period. If the observed RV signal of ϵ Cyg was caused by this mechanism, the star's rotation

period therefore would have to be shorter than ~ 280 d, and the inclination of its spin axis i_* smaller than $\sim 30^\circ$ (see Sec. 4.6.2, given that the v_{rot} measurement by Gray (2015) is correct). As the work by Saio et al. (2018) is focused on upper MS stars, it obviously remains questionable whether r modes can also be present in HB stars like ϵ Cyg A.

Another possibility is that there actually is a planet present around ϵ Cyg A, and its orbital period corresponds to the rotational period of the star, meaning the system is tidally locked. The observed RV variations could then be a sum of the Keplerian signal of the star and the rotational modulation, therefore solving the mismatch between the low photometric variations and the large RV amplitude for a rotational signal alone. Additionally, a tidally coupled system might improve the stability of a planetary companion in an orbit around the best-fit parameters derived in Section 4.4.3. In order to investigate this, it is necessary to pair an N -body code with a state-of-the-art algorithm for the computation of tidal forcing, which is beyond the scope of this work.

For the sake of completeness, we also mention the recent work by Maciejewski et al. (2020), who show that the slightly eccentric RV signature of WASP-12, which has been explained by a hot Jupiter on an unusual eccentric orbit, could also be explained by a circular orbit of the planet and the signature of tides in the host star. Arras et al. (2012) delivered the theoretical foundation, proving that tides in stars, induced by massive planets on tight orbits, can manifest themselves in RV signals with periods of half the orbital period and an amplitude of a few m s^{-1} ; due to a phase shift between the planetary and tidal RV signals the combined signature can then mimic an orbit with nonzero eccentricity and longitude of periastron equal to 270° . This raises the question whether a similar effect could be at work in ϵ Cyg, where tides raised in the primary ϵ Cyg A during the close periastron passage of the spectroscopic companion might contribute to the recorded RV signature; with the eccentricity of the long-period orbit being constrained mostly by the data around periastron passage, the true orbital eccentricity might then be lower than our models suggest. This in turn could solve the stability issues of the putative planet discussed in Section 4.5 if the eccentricity actually would drop below a value of 0.9. It is questionable though whether the RV contribution of tides would be large enough, and a more thorough analysis is needed to assess this possibility.

Generally however, we must conclude that the cause of the short-period RV variations of ϵ Cyg remains a mystery for now. This analysis confirms what has been shown in other recent publications about evolved stars: With more years of continued RV measurements we are now detecting more stars with signals that mimic planets but are more likely caused by other phenomena. Only thanks to the long time spans of the observations and through careful analysis are we able to identify these false positives. This should be a warning, as a considerable number of the published planets around giant stars, especially the early discoveries, had been declared on the basis of very sparse data sets, and some of them might have to be revoked once more measurements are available.

CONCLUSIONS AND FUTURE PROSPECTS

My research and work done during my PhD focussed on different elements that are all concerned with the detection of extrasolar planets via the radial velocity (RV) method. More specifically, I contributed to the RV survey of G- and K-giant stars, performed by the exoplanet group at the Landessternwarte (LSW) Heidelberg. In the stellar sample, in addition to previously published planetary systems, several stars with candidate planets exist; however, it can be complicated to identify whether a planetary companion is the true cause of an observed RV signal, as a variety of intrinsic stellar effects can produce misleading RV modulations. To clearly discern between these false positives and actual planetary systems, a profound analysis of the data is required, and in many cases additional RV measurements obtained over longer baselines are needed.

The main goals of my work were to reach a better understanding of the validity of planet candidates in the sample of G-/K-giant stars, to improve our knowledge about false-positive signals induced by other astrophysical phenomena, and to help in building observing capabilities and develop analysis tools which can be used to further pursue this research on an extended data basis in the future. My research on a target of the K-giant sample, ϵ Cygni, which was presented in Chapter 4 and resulted in a publication (Heeren et al., 2021), serves as a prime example for the importance of long-baseline RV measurements and provides a thorough analysis of that data. By fitting Keplerian models to RV measurements from our Lick survey, from the SONG telescope and archival data, spanning roughly 30 yr in total, I constrained the highly eccentric orbit of the unseen stellar companion to ϵ Cyg A to higher precision than achieved previously. In addition to Keplerian models, I applied fully dynamical models to explore short-period RV variations observed in the data, which seem indicative of a Jupiter-mass planet on a circumprimary orbit, and I executed a broad dynamical analysis of the system to determine its long-term stability. Thanks to these methods I could show that a planetary companion is a highly unlikely explanation for the short-period RV signal, as the amplitude and period of the modulation change over time, and the best-fit orbits lie far from dynamically stable regions in the parameter space.

Similar cases in the literature are the evolved stars γ Draconis (Hatzes et al., 2018) and Aldebaran (Reichert et al., 2019), whose RVs have been monitored over decades, and both of which show sudden changes in the observed RV signals only after many years of gathered measurements. The cause of the observed modulation has not been concretely identified in any of these stars. It is probable that both γ Dra and Aldebaran experience nonadiabatic g^- oscillations as described by Saio et al. (2015);

as my analysis showed, ϵ Cyg in contrast might be a very extreme example of a heartbeat system, where oscillations in the primary star could be induced through tidal interaction with the stellar companion. Despite the uncertainty about the actual astrophysical nature at play though, all three examples showcase not only the value of thorough analysis techniques, but also the necessity of long-baseline observations to find deviations from Keplerian or dynamical models and consequently reject the planet hypothesis for these kinds of stars. This is of high importance, as falsely identified and published planetary companions will distort estimations of planet occurrence rates, even more so due to the relatively small overall number of planets discovered to date around evolved stars.

For our research on the RV survey of G-/K-giant stars, the Waltz telescope project is therefore key to trace remaining planet candidates in the sample and ultimately improve our understanding of planet occurrences and intrinsic stellar sources of RV variations. I presented the current status of the Waltz project and described my contributions to its instruments in Chapter 2. Notably, I assembled and mounted the fiber-feeding unit (FFU), which was the last missing component to reach first light on sky and obtain stellar spectra. While RV results from the Waltz telescope are still missing, first observations indicate that all projected target characteristics of the project are met, as the spectrograph delivers good spectral properties and the overall light efficiency reached under good conditions is at least on the same level as for the original Lick observations despite a lack of guiding at this point.

My main contribution to the project was the development of the Waltz data reduction software (DRS), which I have presented in Chapter 3. First tests of the software on Lick and SONG spectra show a performance similar to the precision reached by the dedicated instrument pipelines, which is on the low m s^{-1} -level. Furthermore, the successful extraction of high-precision RVs from spectra obtained with these two different instruments proves the flexibility and functionality of the Waltz DRS, which was an important design goal in the development.

In the future, I plan to adapt the software to work with Waltz spectra, further test its performance on SONG spectra, and include and test additional methods, such as an advanced template creation and cleaning of the observations from the I2 features; ultimately, I want to work together with the authors of the original CERES and pyodine packages to publish the code of the Waltz DRS under an open-access-license (see Section 3.7).

Also, it will be an interesting project to re-analyze all Lick spectra in our archive to compute new RV timeseries for the stars in our G-/K-giant sample with the Waltz DRS, and compare these results to the original Lick RVs. This will not only help to better assess the performance of the Waltz DRS, but moreover, it might lead to new astrophysical insights wherever systematic differences occur between the Lick and Waltz DRS RVs. An example of such a potentially different interpretation has been presented for the star HIP 36616 in Section 3.5.2, whose RV timeseries as computed by the Waltz DRS differs significantly from the Lick RVs for a short observational epoch, indicating that the Lick results of the affected observations may be influenced by errors; redoing the dynamical analysis presented in Trifonov et al. (2018), but using the Waltz DRS RVs as foundation, could likely lead to different conclusions than before.

WALTZ DRS PARAMETERS

TABLE A.1: Parameters used for reduction of Waltz spectra, as contained in the parameter input file

Parameter name	Parameter description	Value
General information:		
input_directory	Path to raw input spectra	"/path/"
avoid_plot	No interactive mode with step-by-step analysis plots (True, False)	True
npools	Number of processor cores to use in extraction	4
output_directory	Path to reduction output	"/path/"
order_dir	Path to ThAr wavelength atlas	"CERES/wavcals_waltz"
n_useful	Maximum nr. of useful orders to extract	46
ro0	Physical order number of lowest extracted order	79
#rotate	Whether raw spectra need to be rotated (True, False)	False
#flip	Whether raw spectra need to be flipped in cross-disp. direction (True, False)	True
object2do	What Objects to do ("all" for all in directory, "new" for not yet reduced ones)	"all"
force_pre_process	Force pre-processing of spectra (True, False)	False
force_flat_extract	Force extraction of masterflat frame (True, False)	False
force_thar_extract	Force extraction of ThAr frames (True, False)	False
force_thar_wavcal	Force wavelength calibration using ThAr frames (True, False)	False
force_sci_extract	Force extraction of science frames (True, False)	True
force_spectral_file_build	Force building of science output frames (True, False)	True
bary_wave_scale	Transform wavelength scale to barycentric frame (True, False)	False
wavelength_type	Convert wavelengths from air to vacuum, or vice versa, or not at all ("air_to_vac", "vac_to_air", or None)	None
output_science	Output format of reduced spectra ("CERES", "SONG", "both")	"both"
Scattered light removal:		
scatter_type	Averaging of inter-order regions ("median", "min")	"median"
scatter_neg	Allow negative values in scattered light frame (True, False)	False
scatter_option	Force scatter estimate close to CCD edge (1) or not (0)	1
Order tracing and extraction:		

TABLE A.1: continued.

Parameter name	Parameter description	Value
order_multiplicity	Number of extracted fibers/slits for each order	1
order_width	Height of the Echelle orders in pixels (in cross-disp. direction)	8
order_separation	Rough separation of the Echelle orders in pixels	30
nSIGMAS	Min. nr. of standard dev. of orders above background	4.
min_extract_col	Min. pixel column for order extraction (in disp. direction)	0
max_extract_col	Max. pixel column for order extraction (in disp. direction, -1 for last column)	-1
trace_degree	Degree of polynomial for fitting of order traces	5
Marsh_alg	Use Marsh (0) or Horne (1) optimal extraction algorithm	0
NSigma_Marsh	Nr. of standard dev. to ignore cosmics in weight calculation	5
NCosmic_Marsh	Nr. of standard dev. to ignore cosmics in optimal extraction	5
S_Marsh	Spacing between polynomials in weight calculation, in pixels	0.4
N_Marsh	Degree of polynomials in weight calculation	3
Order mapping (code order numbers to physical order numbers):		
code_order_corr	Code order used for correlation	34
d_order_corr	Range of orders around code_order_corr used for correlation	3
order_conv_fac	Guess of first physical order/conversion factor between code orders and physical orders	79
max_pix_shift_disp	Max. pixel shift of spectrum in dispersion direction	100
shift_increment	Cross-correlation step in pixels	0.1
binning	Pixel binning of the ThAr spectrum	1
Wavelength calibration:		
Inverse_m	Spectrum is flipped in cross-disp. direction (True, False)	False
use_cheby	Use Chebyshev polynomials (True, False)	True
Dump_Argon	Reject all Argon lines in the analysis (True, False)	False
do_xc	Do a cross-correlation for each order to better determine pixel shift (True, False)	True
del_width	If do_xc is True: pixel limit for outliers in cross-correlation	5.0
line_width	Half-width of zones around each ThAr lines used for Gaussian fits	5
position_fact	ThAr atlas pixel positions stretch factor	1
rms_max	Max. rms in m/s of initial wavelength solution	150
minlines_init	Min. nr. of lines per order for initial wavelength solution	10
line_sigma	Approx. sigma of ThAr lines in pixels	2.0
MRMS	Max. rms in m/s of global wavelength solution	150
minlines_global	Min. total nr. of lines for global wavelength solution	500
ncoef_x	Degree of polynomials in pixel (disp.) direction	3

TABLE A.1: continued.

Parameter name	Parameter description	Value
ncoef_m	Degree of polynomials in order (cross-disp.) direction	8
npar_wsol	Nr. of free parameters for global wavelength solution fit	31
Continuum normalization:		
cont_deg	Degree of polynomials for continuum fit	3
lower_lim	Nr. of standard dev. for lower flux limit	1.
upper_lim	Nr. of standard dev. for upper flux limit	5.
min_frac_bins	Min. fraction of fitted bins for the final solution	0.1

TABLE A.2: Parameters used for the I2 analysis of Lick spectra with the Waltz DRS

Parameter name	Parameter description	Value
General information:		
osample_obs	Oversampling factor for the model	4
ref_spectrum	Reference spectrum in normalizer and for velocity guess	"arcturus"
velgues_order_range	Order range used for velocity guess	(2,15)
telluric_mask	Telluric mask used	"carmenes"
order_range	Code order range of the input spectra to analyze	(38,53)
chunk_width	Width of the chunks in pixels	40
chunk_padding	Left and right padding of the chunks in pixels	12
chunks_per_order	Number of chunks per order	44
lsf_conv_width	Half-width in pixels to evaluate LSF over	6
Pixel weights:		
weight_type	Pixel weight type ("flat" or "lick")	"flat"
bad_pixel_mask	Whether to compute a bad-pixel mask	True (for template)
bad_pixel_cutoff	Cutoff parameter for the bad-pixel mask	0.22
Deconvolution parameters:		
osample_temp	Oversampling of deconvolved stellar template	10
jansson_niter	Max. number of iterations in deconvolution (k_{\max})	1200
jansson_zerolevel	Spectrum zero-level in deconvolution (a)	0.0
jansson_contlevel	Spectrum continuum-level in deconvolution (b)	1.02
jansson_conver	Convergence parameter in deconvolution (r_0)	0.1
jansson_chi_change	Threshold iteration change at which to stop deconvolution	10^{-6}

TABLE A.3: Parameters used for the I2 analysis of SONG spectra with the Waltz DRS

Parameter name	Parameter description	Value
General information:		
osample_obs	Oversampling factor for the model	6
ref_spectrum	Reference spectrum in normalizer and for velocity guess	"arcturus"
velgues_order_range	Order range used for velocity guess	(4,17)
telluric_mask	Telluric mask used	None

TABLE A.3: continued.

Parameter name	Parameter description	Value
order_range	Code order range of the input spectra to analyze	(18,41)
chunk_width	Width of the chunks in pixels	91
chunk_padding	Left and right padding of the chunks in pixels	23
chunks_per_order	Number of chunks per order	22
lsf_conv_width	Half-width in pixels to evaluate LSF over	6
Pixel weights:		
weight_type	Pixel weight type ("flat" or "lick")	"flat"
bad_pixel_mask	Whether to compute a bad-pixel mask	False
Deconvolution parameters:		
osample_temp	Oversampling of deconvolved stellar template	10
jansson_niter	Max. number of iterations in deconvolution (k_{\max})	1200
jansson_zerolevel	Spectrum zero-level in deconvolution (a)	0.0
jansson_contlevel	Spectrum continuum-level in deconvolution (b)	1.02
jansson_conver	Convergence parameter in deconvolution (r_0)	0.1
jansson_chi_change	Threshold iteration change at which to stop deconvolution	10^{-6}

MEASURED RADIAL VELOCITIES OF ϵ CYGNI

TABLE B.1: RV measurements, taken with the Hamilton Spectrograph at Lick Observatory

Julian date	RV [m/s]	σ_{RV} [m/s]	Julian date	RV [m/s]	σ_{RV} [m/s]
2451706.98438	460.85	4.83	2453613.74222	69.75	4.45
2451707.98438	451.37	4.33	2453649.67509	35.22	4.01
2451741.91016	403.48	7.50	2453855.01682	27.32	3.84
2451780.80664	407.99	4.78	2453911.92222	10.37	4.39
2451783.81445	378.32	4.75	2453933.84138	13.97	4.24
2451854.67676	407.88	4.90	2453967.84960	12.64	4.18
2452081.98535	321.66	4.44	2453981.75301	-36.60	3.75
2452083.88379	361.72	4.02	2453984.73620	-25.23	4.51
2452097.92188	378.95	4.53	2453986.73876	-51.93	4.55
2452412.96094	287.05	4.65	2454054.67229	-80.67	4.73
2452423.95996	293.37	5.42	2454207.04472	-44.79	3.95
2452425.94043	268.87	4.94	2454231.03096	-105.96	5.69
2452437.90332	290.17	4.90	2454266.92516	-102.01	3.72
2452439.91895	315.90	5.38	2454297.80744	-115.85	4.37
2452452.94629	317.19	5.98	2454314.81432	-164.07	4.26
2452453.84082	312.75	5.32	2454344.76833	-147.87	3.96
2452454.87695	320.42	4.94	2454348.78179	-135.68	4.31
2452464.93164	307.34	5.07	2454418.69310	-122.00	4.30
2452466.93555	304.43	5.03	2454443.58918	-85.67	3.98
2452483.86426	305.57	5.49	2454556.05647	-161.36	4.02
2452494.87402	328.13	5.39	2454582.99109	-175.77	4.17
2452496.87305	306.24	5.47	2454600.95348	-209.34	4.41
2452517.80273	286.98	4.96	2454603.96984	-204.89	4.45
2452519.81738	287.03	5.63	2454645.87328	-225.21	4.24
2452528.79688	298.46	5.62	2454667.84292	-244.28	4.80
2452530.75000	306.35	5.10	2454680.79087	-231.34	4.36
2452532.76855	311.60	5.91	2454684.82836	-228.37	4.74
2452541.76660	306.50	4.16	2454712.76665	-219.76	4.57
2452543.75195	304.72	4.83	2454755.75810	-201.52	4.62
2452559.72363	298.63	5.31	2454756.70367	-180.41	3.88
2452560.69336	292.10	5.49	2455028.87157	-297.24	6.85
2452561.66406	290.13	5.67	2455063.87089	-292.51	5.50
2452571.66406	295.23	5.54	2455099.77505	-306.75	5.27

TABLE B.1: continued.

Julian date	RV [m/s]	σ_{RV} [m/s]	Julian date	RV [m/s]	σ_{RV} [m/s]
2452603.61328	259.48	5.35	2455116.78794	-302.60	6.40
2452604.58105	267.44	4.87	2455155.60925	-344.84	8.17
2452616.60156	245.93	5.53	2455303.01287	-410.41	4.80
2452617.58105	235.56	5.75	2455329.00705	-429.26	4.92
2452765.99609	231.00	5.42	2455330.00259	-400.01	4.47
2452800.94922	257.25	4.21	2455360.96854	-396.61	4.14
2452802.92773	267.54	5.29	2455418.86449	-448.03	5.32
2452837.87012	227.85	5.23	2455446.76880	-469.53	5.16
2452838.87793	208.32	5.99	2455463.75022	-463.72	5.33
2452839.87891	230.79	6.14	2455466.71536	-473.39	4.63
2452861.81152	212.26	4.16	2455519.68550	-536.71	4.90
2452863.84277	218.17	5.86	2455651.05584	-491.11	8.26
2452879.83301	210.28	5.38	2455702.92398	-551.59	4.38
2452898.75391	210.09	4.92	2455732.91952	-610.73	4.14
2452933.65625	187.95	4.21	2455756.87251	-611.22	5.69
2452963.60645	154.50	6.08	2455761.87461	-622.09	5.40
2452966.60840	172.86	4.77	2455803.75976	-671.19	4.98
2453169.99707	172.97	3.74	2455805.77373	-680.70	5.00
2453171.85840	162.83	3.85	2455828.71590	-654.03	4.01
2453494.01918	58.84	3.98	2455861.70374	-681.09	4.70
2453548.94052	86.54	4.00	2455893.60677	-674.55	8.85
2453579.83708	59.90	4.15			

TABLE B.2: RV measurements of the SONG Tenerife node (for the high-cadence astero-seismic RV measurements only the nightly median is included)

Julian date	RV [m/s]	σ_{RV} [m/s]	Julian date	RV [m/s]	σ_{RV} [m/s]
2456876.66213	9956.54	2.39	2457828.74982	16016.96	2.45
2456893.56235	9932.62	2.36	2457833.75127	16296.52	1.46
2456900.42744	9940.72	2.24	2457834.75590	16377.74	3.17
2456911.49700	9887.80	2.17	2457835.75563	16417.01	1.59
2457128.75207	9633.20	2.64	2457836.75276	16457.51	1.80
2457134.70641	9637.64	2.82	2457837.75389	16492.14	1.84
2457147.71993	9621.57	2.46	2457838.75862	16551.40	2.34
2457148.68176	9629.92	2.51	2457839.75655	16577.12	1.62
2457165.66913	9575.35	2.46	2457840.75644	16640.81	1.39
2457184.62842	9498.94	3.50	2457841.75743	16659.81	1.67
2457201.67084	9466.58	2.43	2457843.75226	16703.60	1.82
2457215.73554	9407.17	2.22	2457846.75676	16834.92	1.45
2457230.35766	9407.52	1.33	2457847.75390	16851.21	1.21
2457275.67673	9317.86	2.41	2457848.75207	16906.71	1.30
2457296.31717	9313.50	2.23	2457849.75714	16925.57	1.42
2457307.31336	9299.20	2.59	2457850.75341	16960.57	1.46
2457337.29778	9231.60	3.95	2457851.75331	16975.96	1.47
2457344.42681	9220.74	3.27	2457852.73511	16995.09	1.50
2457353.30227	9238.63	2.78	2457853.75138	17036.18	1.27
2457355.32938	9195.68	3.00	2457854.70699	17034.46	1.36
2457368.35399	9187.28	2.81	2457855.67847	17063.25	1.31

TABLE B.2: continued.

Julian date	RV [m/s]	σ_{RV} [m/s]	Julian date	RV [m/s]	σ_{RV} [m/s]
2457370.31721	9189.62	2.72	2457860.73066	17166.22	1.49
2457371.34219	9180.70	4.69	2457867.69807	17284.30	2.89
2457373.30373	9168.94	2.74	2457868.69517	17283.01	1.59
2457379.29089	9172.09	2.92	2457869.72376	17310.66	2.15
2457385.29297	9141.15	2.47	2457872.74037	17332.24	1.88
2457390.28744	9113.12	2.34	2457874.68328	17337.38	1.68
2457394.28481	9132.69	2.24	2457875.68459	17356.59	2.76
2457395.30931	9113.91	2.23	2457876.65464	17385.91	1.62
2457396.31433	9104.33	2.26	2457877.65455	17371.84	2.15
2457397.30749	9095.82	2.80	2457878.65455	17374.18	2.69
2457399.34983	9103.37	3.19	2457879.67772	17380.26	2.09
2457400.30813	9086.09	2.44	2457880.65476	17385.22	1.91
2457401.29182	9088.16	3.07	2457881.64521	17396.75	2.14
2457402.29824	9059.76	5.89	2457883.64838	17439.50	5.31
2457404.29873	9081.41	2.29	2457884.64288	17423.62	3.04
2457406.29162	9068.70	2.26	2457885.72305	17418.30	2.07
2457407.29116	9087.69	2.01	2457886.70577	17404.01	2.04
2457453.78479	8892.88	2.61	2457887.74408	17408.31	1.91
2457458.78372	8877.58	2.20	2457888.62472	17418.73	2.15
2457465.77542	8823.23	2.90	2457889.62441	17404.34	2.73
2457467.76988	8831.18	2.90	2457890.62523	17426.41	2.17
2457477.72956	8792.34	2.76	2457899.61014	17433.62	2.04
2457488.74902	8754.96	3.39	2457900.59887	17430.98	2.18
2457492.71883	8731.70	3.51	2457901.60236	17431.37	2.01
2457498.66677	8706.91	3.41	2457906.57609	17440.78	2.21
2457503.65954	8696.50	2.78	2457907.57607	17425.46	2.32
2457507.61563	8671.26	2.74	2457908.56631	17415.66	1.92
2457508.65709	8670.33	2.60	2457910.56926	17426.02	2.13
2457521.67134	8637.01	2.76	2457911.56222	17449.66	2.30
2457528.65132	8589.04	2.82	2457912.56265	17435.39	2.06
2457534.64232	8589.76	2.59	2457913.56321	17406.59	2.10
2457535.64362	8584.11	3.01	2457914.55621	17425.98	2.00
2457538.59534	8579.41	2.68	2457915.55526	17424.71	1.89
2457546.59142	8544.96	2.58	2457916.55283	17410.69	1.81
2457557.57420	8517.30	2.93	2457917.55296	17414.37	1.98
2457577.51529	8459.87	2.55	2457938.51324	17309.43	1.31
2457579.73807	8433.51	2.52	2457948.72159	17249.88	2.22
2457583.73032	8437.60	2.27	2457961.75166	17205.69	2.65
2457592.74718	8392.62	1.73	2457990.65169	17016.09	2.33
2457594.75188	8393.96	2.33	2458002.49353	16914.85	1.85
2457622.68442	8359.69	2.89	2458027.41918	16769.17	2.41
2457625.67862	8339.92	2.55	2458042.37556	16693.05	2.31
2457626.67897	8342.04	2.37	2458056.34249	16602.83	1.75
2457633.41378	8333.12	2.78	2458074.49736	16484.48	2.79
2457642.37503	8295.42	2.36	2458095.32505	16359.57	2.20
2457643.37491	8314.78	2.36	2458108.31744	16288.57	2.19
2457648.38175	8280.11	2.39	2458121.30367	16253.26	3.07
2457659.52439	8276.09	2.64	2458211.76464	15967.45	1.39

TABLE B.2: continued.

Julian date	RV [m/s]	σ_{RV} [m/s]	Julian date	RV [m/s]	σ_{RV} [m/s]
2457683.38376	8251.49	2.74	2458215.76079	15997.44	1.67
2457710.29156	8380.45	2.68	2458234.69838	15923.49	1.88
2457711.29082	8370.53	2.77	2458248.67751	15895.13	1.91
2457717.35697	8443.63	2.84	2458262.58752	15867.21	1.99
2457722.34122	8523.71	3.43	2458278.61011	15825.06	1.72
2457723.29371	8556.30	2.75	2458290.50905	15772.83	1.43
2457725.30536	8573.93	3.10	2458291.49666	15763.27	1.25
2457728.30468	8627.07	2.82	2458292.63062	15767.91	1.10
2457729.29360	8650.22	2.75	2458293.44104	15769.16	1.54
2457730.30416	8672.82	2.65	2458294.58493	15760.53	1.07
2457731.29381	8675.86	2.71	2458295.49752	15773.74	1.16
2457732.29817	8707.91	2.95	2458296.53234	15755.39	1.21
2457733.29409	8734.55	2.70	2458297.56057	15748.97	0.89
2457738.33363	8857.51	2.72	2458298.54929	15754.48	0.87
2457745.34263	9054.82	4.94	2458299.70491	15744.46	0.86
2457752.31710	9367.17	2.89	2458300.49791	15751.42	1.05
2457753.30971	9397.21	3.29	2458301.58985	15753.58	0.86
2457758.31112	9683.94	3.29	2458302.54091	15758.19	1.10
2457761.29108	9857.82	3.64	2458303.49003	15732.94	1.53
2457764.31155	10040.42	2.66	2458304.49352	15732.44	1.07
2457765.30985	10095.42	2.67	2458305.39540	15729.80	2.28
2457766.28813	10168.91	4.42	2458306.56399	15709.50	1.21
2457767.28864	10277.07	3.89	2458307.63982	15723.92	1.09
2457768.28915	10310.56	2.45	2458311.55945	15709.01	1.27
2457769.28990	10396.51	2.68	2458312.65316	15716.98	1.21
2457770.29059	10488.73	2.63	2458313.43194	15709.90	2.32
2457773.29174	10729.36	2.43	2458314.71155	15709.75	1.51
2457778.29415	11176.75	2.04	2458315.55202	15713.12	0.98
2457781.29579	11466.26	2.49	2458316.59218	15709.41	1.14
2457782.29660	11554.08	2.46	2458317.70719	15706.30	1.28
2457783.29682	11647.21	2.35	2458321.41017	15677.77	1.63
2457788.79584	12241.47	2.28	2458322.53861	15682.02	1.61
2457789.79588	12329.03	2.28	2458323.46522	15682.35	1.89
2457790.79273	12485.57	2.60	2458324.45842	15678.36	1.90
2457791.79248	12543.80	2.61	2458325.53906	15669.93	1.84
2457792.79193	12664.89	2.33	2458326.67828	15669.99	1.90
2457799.78923	13458.14	5.10	2458334.61633	15652.50	2.07
2457800.78674	13527.54	2.81	2458346.76124	15594.36	2.88
2457819.78193	15387.50	2.41	2458359.68002	15566.62	2.03
2457820.77734	15447.16	3.02	2458380.58560	15508.34	1.95
2457822.77703	15595.58	2.53	2458405.39670	15449.90	2.52
2457823.77433	15658.11	2.42	2458424.34220	15394.80	2.18
2457824.77154	15733.44	2.33	2458449.46031	15338.05	2.93
2457826.75711	15870.96	2.74	2458467.28125	15321.43	2.20
2457827.73326	15958.18	2.17	2458480.33486	15317.27	2.73

Own publications

The following first-author publication was used in this thesis:

- **Heeren, P.**, Reffert, S., Trifonov, T., Wong, K. H., Lee, M. H., Lillo-Box, J., Quirrenbach, A., Arentoft, T., Albrecht, S., Grundahl, F., Andersen, M. F., Antoci, V., and Pallé, P. L. (2021). "Precise radial velocities of giant stars. XV. Mysterious nearly periodic radial velocity variations in the eccentric binary ε Cygni". In: *A&A* 647, A160. DOI: [10.1051/0004-6361/202040087](https://doi.org/10.1051/0004-6361/202040087).

I have also co-authored the following publications, which are not included in this thesis:

- Hjorth, M., Justesen, A. B., Hirano, T., Albrecht, S., Gandolfi, D., Dai, F., Alonso, R., Barragán, O., Esposito, M., Kuzuhara, M., Lam, K. W. F., Livingston, J. H., Montanes-Rodriguez, P., Narita, N., Nowak, G., Prieto-Arranz, J., Redfield, S., Rodler, F., Van Eylen, V., Winn, J. N. Antoniciello, G., Cabrera, J., Cochran, W. D., Csizmadia, Sz, de Leon, J., Deeg, H., Eigmüller, Ph, Endl, M., Erikson, A., Fridlund, M., Grziwa, S., Guenther, E., Hatzes, A. P., **Heeren, P.**, Hidalgo, D., Korth, J., Luque, R., Nespral, D., Palle, E., Pätzold, M., Persson, C. M., Rauer, H., Smith, A. M. S., and Trifonov, T. (2019). "K2-290: a warm Jupiter and a mini-Neptune in a triple-star system". In: *MNRAS* 484.3, pp. 3522–3536. DOI: [10.1093/mnras/stz139](https://doi.org/10.1093/mnras/stz139).
- Persson, C. M., Fridlund, M., Barragán, O., Dai, F., Gandolfi, D., Hatzes, A. P., Hirano, T., Grziwa, S., Korth, J., Prieto-Arranz, J., Fossati, L., Van Eylen, V., Justesen, A. B., Livingston, J., Kubyskhina, D., Deeg, H. J., Guenther, E. W., Nowak, G., Cabrera, J., Eigmüller, Ph. Csizmadia, Sz., Smith, A. M. S., Erikson, A., Albrecht, S., Sobrino, Alonso, Cochran, W. D., Endl, M., Esposito, M., Fukui, A., **Heeren, P.**, Hidalgo, D., Hjorth, M., Kuzuhara, M., Narita, N., Nespral, D., Palle, E., Pätzold, M., Rauer, H., Rodler, F., and Winn, J. N. (2018). "Super-Earth of $8 M_{\oplus}$ in a 2.2-day orbit around the K5V star K2-216". In: *A&A* 618, A33. DOI: [10.1051/0004-6361/201832867](https://doi.org/10.1051/0004-6361/201832867).
- Trifonov, T., Kürster, M., Reffert, S., Zechmeister, M., Endl, M., Rodler, F., Gandolfi, D., Barragán, O., Henning, T., Lee, M. H., Zakhozhay, O., Sarkis, P., **Heeren, P.**, Tala, M., Wolthoff, V., Brems, S. S., Stock, S., Hempel, A., and Kosakowski, D. (2018). "New HARPS and FEROS Observations of GJ 1046". In: *Research Notes of the American Astronomical Society*, Vol. 2, Issue 3, 180. DOI: [10.3847/2515-5172/aae437](https://doi.org/10.3847/2515-5172/aae437).
- Van Eylen, V., Dai, F., Mathur, S., Gandolfi, D., Albrecht, S., Fridlund, M., García, R. A., Guenther, E., Hjorth, M., Justesen, A. B., Livingston, J., Lund, M.

- N., Pérez Hernández, F., Prieto-Arranz, J., Regulo, C., Bugnet, L., Everett, M. E., Hirano, T., Nespral, D., Nowak, G. Palle, E., Silva Aguirre, V., Trifonov, T., Winn, J. N., Barragán, O., Beck, P. G., Chaplin, W. J., Cochran, W. D., Csizmadia, S., Deeg, H., Endl, M., **Heeren, P.**, Grziwa, S., Hatzes, A. P., Hidalgo, D., Korth, J., Mathis, S., Montañes Rodriguez, P., Narita, N., Patzold, M., Persson, C. M., Rodler, F., and Smith, A. M. S. (2018). "HD 89345: a bright oscillating star hosting a transiting warm Saturn-sized planet observed by K2". In: *MNRAS* 478.4, pp. 4866–4880. DOI: [10.1093/mnras/sty1390](https://doi.org/10.1093/mnras/sty1390).
- Tala, M., **Heeren, P.**, Grill, M., Harris, R. J., Stürmer, J., Schwab, C., Gutcke, T., Reffert, S., Quirrenbach, A., Seifert, W., Mandel, H., Geuer, L., Schäffner, L., Thimm, G., Seeman, U., Tietz, J., and Wagner, K. (2016). "A high-resolution spectrograph for the 72cm Waltz Telescope at Landessternwarte, Heidelberg". In: *Ground-based and Airborne Instrumentation for Astronomy VI. Proceedings of the SPIE* Vol. 9908, 99086O. DOI: [10.1117/12.2232730](https://doi.org/10.1117/12.2232730).

Bibliography

- Agard, D.A. et al. (1989). "Fluorescence microscopy in three dimensions". In: *Methods Cell Biol.* 30.353, p. 77. DOI: [10.1016/s0091-679x\(08\)60986-3](https://doi.org/10.1016/s0091-679x(08)60986-3).
- Andersen, M. F. et al. (2014). "Hardware and software for a robotic network of telescopes - SONG". In: *Revista Mexicana de Astronomia y Astrofisica Conference Series*. Vol. 45, p. 83. arXiv: [1901.08300](https://arxiv.org/abs/1901.08300) [astro-ph.IM].
- Andersen, M. F. et al. (2016). "The SONG prototype: Efficiency of a robotic telescope". In: *Revista Mexicana de Astronomia y Astrofisica Conference Series*. Vol. 48, pp. 54–58. arXiv: [1901.08293](https://arxiv.org/abs/1901.08293) [astro-ph.IM].
- Anglada-Escudé, Guillem and R. Paul Butler (2012). "The HARPS-TERRA Project. I. Description of the Algorithms, Performance, and New Measurements on a Few Remarkable Stars Observed by HARPS". In: *ApJS* 200.2, 15, p. 15. DOI: [10.1088/0067-0049/200/2/15](https://doi.org/10.1088/0067-0049/200/2/15). arXiv: [1202.2570](https://arxiv.org/abs/1202.2570) [astro-ph.EP].
- Antoci, V. et al. (2013). "Searching for solar-like oscillations in the δ Scuti star ρ Puppis". In: *MNRAS* 435.2, pp. 1563–1575. ISSN: 0035-8711. DOI: [10.1093/mnras/stt1397](https://doi.org/10.1093/mnras/stt1397). eprint: <https://academic.oup.com/mnras/article-pdf/435/2/1563/3524587/stt1397.pdf>. URL: <https://doi.org/10.1093/mnras/stt1397>.
- Arentoft, T. et al. (2019). "Astero-seismology of the Hyades red giant and planet host ϵ Tau". In: *A&A* 622, A190. DOI: [10.1051/0004-6361/201834690](https://doi.org/10.1051/0004-6361/201834690).
- Arras, Phil et al. (2012). "The radial velocity signature of tides raised in stars hosting exoplanets". In: *MNRAS* 422.2, pp. 1761–1766. DOI: [10.1111/j.1365-2966.2012.20756.x](https://doi.org/10.1111/j.1365-2966.2012.20756.x). arXiv: [1107.6005](https://arxiv.org/abs/1107.6005) [astro-ph.SR].
- Baranne, A. et al. (1996). "ELODIE: A spectrograph for accurate radial velocity measurements." In: *A&AS* 119, pp. 373–390.
- Barnes, J. R., S. V. Jeffers, and H. R. A. Jones (2011). "The effect of M dwarf starspot activity on low-mass planet detection thresholds". In: *MNRAS* 412.3, pp. 1599–1610. DOI: [10.1111/j.1365-2966.2010.17979.x](https://doi.org/10.1111/j.1365-2966.2010.17979.x).
- Beck, P.G. et al. (2014). "Pulsating red giant stars in eccentric binary systems discovered from Kepler space-based photometry". In: *A&A* 564, A36. DOI: [10.1051/0004-6361/201322477](https://doi.org/10.1051/0004-6361/201322477).
- Bedell, Megan et al. (2019). "WOBBLE: A Data-driven Analysis Technique for Time-series Stellar Spectra". In: *AJ* 158.4, 164, p. 164. DOI: [10.3847/1538-3881/ab40a7](https://doi.org/10.3847/1538-3881/ab40a7). arXiv: [1901.00503](https://arxiv.org/abs/1901.00503) [astro-ph.IM].
- Bowen, I. S. (1938). "The Image-Slicer a Device for Reducing Loss of Light at Slit of Stellar Spectrograph." In: *ApJ* 88, p. 113. DOI: [10.1086/143964](https://doi.org/10.1086/143964).
- Brahm, Rafael, Andrés Jordán, and Néstor Espinoza (2017). "CERES: A Set of Automated Routines for Echelle Spectra". In: *PASP* 129.973, p. 034002. DOI: [10.1088/1538-3873/aa5455](https://doi.org/10.1088/1538-3873/aa5455). arXiv: [1609.02279](https://arxiv.org/abs/1609.02279) [astro-ph.IM].
- Butler, R. P. et al. (1996). "Attaining Doppler Precision of 3 M s⁻¹". In: *PASP* 108, p. 500. DOI: [10.1086/133755](https://doi.org/10.1086/133755).

- Campbell, Bruce, G. A. H. Walker, and S. Yang (1988). "A Search for Substellar Companions to Solar-type Stars". In: *ApJ* 331, p. 902. DOI: [10.1086/166608](https://doi.org/10.1086/166608).
- Campbell, W. W. and J. H. Moore (1906). "Seven new spectroscopic binaries." In: *PASP* 18, p. 308. DOI: [10.1086/121705](https://doi.org/10.1086/121705).
- Cassan, A. et al. (2012). "One or more bound planets per Milky Way star from microlensing observations". In: *Nature* 481.7380, pp. 167–169. DOI: [10.1038/nature10684](https://doi.org/10.1038/nature10684). arXiv: [1202.0903](https://arxiv.org/abs/1202.0903) [astro-ph.EP].
- Chambers, J. E. (1999). "A symplectic integration scheme that allows close encounters between massive bodies." In: *Impact of Modern Dynamics in Astronomy*. Ed. by J. Henrard and S. Ferraz-Mello, p. 449.
- Crilly, P. B. et al. (2002). "Improving the convergence rate of Jansson's deconvolution method". In: *IEEE Transactions on Instrumentation and Measurement* 51.6, pp. 1142–1144. DOI: [10.1109/TIM.2002.807981](https://doi.org/10.1109/TIM.2002.807981).
- Cumming, Andrew et al. (2008). "The Keck Planet Search: Detectability and the Minimum Mass and Orbital Period Distribution of Extrasolar Planets". In: *PASP* 120.867, p. 531. DOI: [10.1086/588487](https://doi.org/10.1086/588487). arXiv: [0803.3357](https://arxiv.org/abs/0803.3357) [astro-ph].
- Díaz, Matías R. et al. (2019). "Deriving Iodine-free Spectra for High-resolution Echelle Spectrographs". In: *AJ* 157.5, 204, p. 204. DOI: [10.3847/1538-3881/ab14ed](https://doi.org/10.3847/1538-3881/ab14ed). arXiv: [1903.11718](https://arxiv.org/abs/1903.11718) [astro-ph.IM].
- D'Oodorico, S. et al. (1984). *An Atlas of the Thorium-Argon Spectrum for the ESO Echelle Spectrograph*. European Southern Observatory.
- Dumusque, X. (2016). "Radial velocity fitting challenge. I. Simulating the data set including realistic stellar radial-velocity signals". In: *A&A* 593, A5, A5. DOI: [10.1051/0004-6361/201628672](https://doi.org/10.1051/0004-6361/201628672). arXiv: [1607.06487](https://arxiv.org/abs/1607.06487) [astro-ph.EP].
- Dumusque, X., I. Boisse, and N.C. Santos (2014). "SOAP 2.0: A Tool to Estimate the Photometric and Radial Velocity Variations Induced by Stellar Spots and Plages". In: *ApJ* 796.2, p. 132. DOI: [10.1088/0004-637X/796/2/132](https://doi.org/10.1088/0004-637X/796/2/132).
- Dumusque, X. et al. (2017). "Radial-velocity fitting challenge. II. First results of the analysis of the data set". In: *A&A* 598, A133, A133. DOI: [10.1051/0004-6361/201628671](https://doi.org/10.1051/0004-6361/201628671). arXiv: [1609.03674](https://arxiv.org/abs/1609.03674) [astro-ph.EP].
- Duncan, Martin J., Harold F. Levison, and Man Hoi Lee (1998). "A Multiple Time Step Symplectic Algorithm for Integrating Close Encounters". In: *AJ* 116.4, pp. 2067–2077. DOI: [10.1086/300541](https://doi.org/10.1086/300541).
- Eastman, Jason, Robert Siverd, and B. Scott Gaudi (2010). "Achieving Better Than 1 Minute Accuracy in the Heliocentric and Barycentric Julian Dates". In: *PASP* 122.894, p. 935. DOI: [10.1086/655938](https://doi.org/10.1086/655938). arXiv: [1005.4415](https://arxiv.org/abs/1005.4415) [astro-ph.IM].
- Edlén, Bengt (1953). "The Dispersion of Standard Air". In: *J. Opt. Soc. Am.* 43.5, pp. 339–344. DOI: [10.1364/JOSA.43.000339](https://doi.org/10.1364/JOSA.43.000339). URL: <http://www.osapublishing.org/abstract.cfm?URI=josa-43-5-339>.
- Eversberg, Thomas and Klaus Vollmann (2015). *Spectroscopic Instrumentation - Fundamentals and Guidelines for Astronomers*. Springer-Verlag Berlin Heidelberg. ISBN: 978-3-662-44535-8. DOI: [10.1007/978-3-662-44535-8](https://doi.org/10.1007/978-3-662-44535-8).
- Fernandes, Rachel B. et al. (2019). "Hints for a Turnover at the Snow Line in the Giant Planet Occurrence Rate". In: *ApJ* 874.1, 81, p. 81. DOI: [10.3847/1538-4357/ab0300](https://doi.org/10.3847/1538-4357/ab0300). arXiv: [1812.05569](https://arxiv.org/abs/1812.05569) [astro-ph.SR].
- Figueira, P. et al. (2013). "Line-profile variations in radial-velocity measurements. Two alternative indicators for planetary searches". In: *A&A* 557, A93, A93. DOI: [10.1051/0004-6361/201220779](https://doi.org/10.1051/0004-6361/201220779). arXiv: [1307.7279](https://arxiv.org/abs/1307.7279) [astro-ph.EP].
- Fischer, D. A. et al. (2014). "Exoplanet Detection Techniques". In: *Protostars and Planets VI*. Ed. by Henrik Beuther et al., p. 715. DOI: [10.2458/azu_uapress_9780816531240-ch031](https://doi.org/10.2458/azu_uapress_9780816531240-ch031). arXiv: [1505.06869](https://arxiv.org/abs/1505.06869) [astro-ph.EP].

- Fischer, Debra A., Geoffrey W. Marcy, and Julien F. P. Spronck (2014). "The Twenty-five Year Lick Planet Search". In: *ApJS* 210.1, 5, p. 5. DOI: [10.1088/0067-0049/210/1/5](https://doi.org/10.1088/0067-0049/210/1/5). arXiv: [1310.7315](https://arxiv.org/abs/1310.7315) [astro-ph.EP].
- Fischer, Debra A. and Jeff Valenti (2005). "The Planet-Metallicity Correlation". In: *ApJ* 622.2, pp. 1102–1117. DOI: [10.1086/428383](https://doi.org/10.1086/428383).
- Foreman-Mackey, Daniel et al. (2013). "emcee: The MCMC Hammer". In: *PASP* 125.925, p. 306. DOI: [10.1086/670067](https://doi.org/10.1086/670067). arXiv: [1202.3665](https://arxiv.org/abs/1202.3665) [astro-ph.IM].
- Fredslund Andersen, M. et al. (2019). "Conducting the SONG: The Robotic Nature and Efficiency of a Fully Automated Telescope". In: *PASP* 131.998, p. 045003. DOI: [10.1088/1538-3873/aaff7c](https://doi.org/10.1088/1538-3873/aaff7c). arXiv: [1901.07560](https://arxiv.org/abs/1901.07560) [astro-ph.IM].
- Frink, Sabine et al. (2001). "A Strategy for Identifying the Grid Stars for the Space Interferometry Mission". In: *PASP* 113.780, pp. 173–187. DOI: [10.1086/318610](https://doi.org/10.1086/318610).
- Frink, Sabine et al. (2002). "Discovery of a Substellar Companion to the K2 III Giant ι Draconis". In: *ApJ* 576.1, pp. 478–484. DOI: [10.1086/341629](https://doi.org/10.1086/341629).
- Fuller, Jim (2017). "Heartbeat stars, tidally excited oscillations and resonance locking". In: *MNRAS* 472.2, pp. 1538–1564. DOI: [10.1093/mnras/stx2135](https://doi.org/10.1093/mnras/stx2135).
- Fulton, Benjamin J. et al. (2017). "The California-Kepler Survey. III. A Gap in the Radius Distribution of Small Planets". In: *AJ* 154.3, 109, p. 109. DOI: [10.3847/1538-3881/aa80eb](https://doi.org/10.3847/1538-3881/aa80eb). arXiv: [1703.10375](https://arxiv.org/abs/1703.10375) [astro-ph.EP].
- Gaia Collaboration et al. (2018). "Gaia Data Release 2. Summary of the contents and survey properties". In: *A&A* 616, A1, A1. DOI: [10.1051/0004-6361/201833051](https://doi.org/10.1051/0004-6361/201833051). arXiv: [1804.09365](https://arxiv.org/abs/1804.09365) [astro-ph.GA].
- Gilliland, Ronald L. et al. (1992). "Resolution and Noise Properties of the Goddard High Resolution Spectrograph". In: *PASP* 104, p. 367. DOI: [10.1086/133006](https://doi.org/10.1086/133006).
- Gladman, Brett (1993). "Dynamics of Systems of Two Close Planets". In: *Icarus* 106.1, pp. 247–263. DOI: [10.1006/icar.1993.1169](https://doi.org/10.1006/icar.1993.1169).
- Gray, D. F. (1982). "The temperature dependence of rotation and turbulence in giant stars". In: *ApJ* 262, p. 682. DOI: [10.1086/160461](https://doi.org/10.1086/160461).
- Gray, David F. (2015). "A spectroscopic analysis of the K0 III binary ϵ Cygni". In: *ApJ* 810, p. 117. DOI: [10.1088/0004-637x/810/2/117](https://doi.org/10.1088/0004-637x/810/2/117).
- Griffin, R. F. (1994). "Photoelectric radial velocities, Paper XIV. Variation of the radial velocity of ϵ Cygni". In: *MNRAS* 267, pp. 69–76.
- Grundahl, F. et al. (2017). "First Results from the Hertzprung SONG Telescope: Asteroseismology of the G5 Subgiant Star μ Herculis". In: *ApJ* 836.1, 142, p. 142. DOI: [10.3847/1538-4357/836/1/142](https://doi.org/10.3847/1538-4357/836/1/142). arXiv: [1701.03365](https://arxiv.org/abs/1701.03365) [astro-ph.SR].
- Halverson, Samuel et al. (2014). "Development of Fiber Fabry-Perot Interferometers as Stable Near-infrared Calibration Sources for High Resolution Spectrographs". In: *PASP* 126.939, p. 445. DOI: [10.1086/676649](https://doi.org/10.1086/676649). arXiv: [1403.6841](https://arxiv.org/abs/1403.6841) [astro-ph.IM].
- Hamilton, Douglas P. and Joseph A. Burns (1992). "Orbital stability zones about asteroids: II. The destabilizing effects of eccentric orbits and of solar radiation". In: *Icarus* 96.1, pp. 43–64. DOI: [10.1016/0019-1035\(92\)90005-R](https://doi.org/10.1016/0019-1035(92)90005-R).
- Han, Inwoo et al. (2010). "Detection of a planetary companion around the giant star γ 1 Leonis". In: *A&A* 509, A24. DOI: [10.1051/0004-6361/200912536](https://doi.org/10.1051/0004-6361/200912536).
- Hatzes, A. P. et al. (2006). "Confirmation of the planet hypothesis for the long-period radial velocity variations of β Geminorum". In: *A&A* 457.1, pp. 335–341. DOI: [10.1051/0004-6361:20065445](https://doi.org/10.1051/0004-6361:20065445). arXiv: [astro-ph/0606517](https://arxiv.org/abs/astro-ph/0606517) [astro-ph].
- Hatzes, A. P. et al. (2018). "The Radial Velocity Variability of the K-giant γ Draconis: Stellar Variability Masquerading as a Planet". In: *AJ* 155, p. 120. DOI: [10.3847/1538-3881/aaa8e1](https://doi.org/10.3847/1538-3881/aaa8e1).
- Hatzes, Artie P. and William D. Cochran (1993). "Long-Period Radial Velocity Variations in Three K Giants". In: *ApJ* 413, p. 339. DOI: [10.1086/173002](https://doi.org/10.1086/173002).

- Hatzes, Artie P. et al. (2003). "A Planetary Companion to γ Cephei A". In: *ApJ* 599.2, pp. 1383–1394. DOI: [10.1086/379281](https://doi.org/10.1086/379281). arXiv: [astro-ph/0305110](https://arxiv.org/abs/astro-ph/0305110) [astro-ph].
- Heeren, Paul et al. (2021). "Precise radial velocities of giant stars. XV. Mysterious nearly periodic radial velocity variations in the eccentric binary ϵ Cygni". In: *A&A* 647, A160. DOI: [10.1051/0004-6361/202040087](https://doi.org/10.1051/0004-6361/202040087).
- Heeren, Paul Phillip (2016). "Constructing an Echelle Spectrograph for Doppler Spectroscopy of G/K Giant Stars". Master thesis. Ruprecht-Karls Universität Heidelberg.
- Hinkle, Kenneth et al. (2000). *Visible and Near Infrared Atlas of the Arcturus Spectrum 3727-9300 Å*.
- Hinkle, K.H. et al. (2002). "Velocity Observations of Multiple-Mode Asymptotic Giant Branch Variable Stars". In: *AJ* 123.2, pp. 1002–1012. DOI: [10.1086/338314](https://doi.org/10.1086/338314).
- Hippke, M. and D. Angerhausen (2015). "A Statistical Search for a Population of Exo-Trojans in the Kepler Data Set". In: *ApJ* 811, 1, p. 1. DOI: [10.1088/0004-637X/811/1/1](https://doi.org/10.1088/0004-637X/811/1/1). arXiv: [1508.00427](https://arxiv.org/abs/1508.00427) [astro-ph.EP].
- Horne, K. (1986). "An optimal extraction algorithm for CCD spectroscopy". In: *PASP* 98, p. 609. DOI: [10.1086/131801](https://doi.org/10.1086/131801). URL: <https://doi.org/10.1086/131801>.
- Hsu, Danley C. et al. (2019). "Occurrence Rates of Planets Orbiting FGK Stars: Combining Kepler DR25, Gaia DR2, and Bayesian Inference". In: *AJ* 158.3, p. 109. ISSN: 1538-3881. DOI: [10.3847/1538-3881/ab31ab](https://doi.org/10.3847/1538-3881/ab31ab). URL: <http://dx.doi.org/10.3847/1538-3881/ab31ab>.
- Ivezić, Željko et al. (2014). *Statistics, Data Mining, and Machine Learning in Astronomy*. Princeton University Press.
- Janson (2013). "A Systematic Search for Trojan Planets in the Kepler Data". In: *ApJ* 774, 156, p. 156. arXiv: [1307.7161](https://arxiv.org/abs/1307.7161) [astro-ph.EP].
- Johnson, John Asher et al. (2010). "Giant Planet Occurrence in the Stellar Mass-Metallicity Plane". In: *PASP* 122.894, p. 905. DOI: [10.1086/655775](https://doi.org/10.1086/655775). arXiv: [1005.3084](https://arxiv.org/abs/1005.3084) [astro-ph.EP].
- Jones, M. I. et al. (2011). "Study of the impact of the post-MS evolution of the host star on the orbits of close-in planets. I. Sample definition and physical properties". In: *A&A* 536, A71, A71. DOI: [10.1051/0004-6361/201117887](https://doi.org/10.1051/0004-6361/201117887). arXiv: [1110.6459](https://arxiv.org/abs/1110.6459) [astro-ph.EP].
- Jones, M. I. et al. (2016). "Four new planets around giant stars and the mass-metallicity correlation of planet-hosting stars". In: *A&A* 590, A38, A38. DOI: [10.1051/0004-6361/201628067](https://doi.org/10.1051/0004-6361/201628067). arXiv: [1603.03738](https://arxiv.org/abs/1603.03738) [astro-ph.EP].
- Kanodia, Shubham and Jason Wright (2018). "Python Leap Second Management and Implementation of Precise Barycentric Correction (barycorrpy)". In: *Research Notes of the AAS* 2.1, p. 4. DOI: [10.3847/2515-5172/aaa4b7](https://doi.org/10.3847/2515-5172/aaa4b7). URL: <https://doi.org/10.3847/2515-5172/aaa4b7>.
- Keenan, Philip C. and Raymond C. McNeil (1989). "The Perkins Catalog of Revised MK Types for the Cooler Stars". In: *ApJS* 71, p. 245. DOI: [10.1086/191373](https://doi.org/10.1086/191373).
- Kennedy, Grant M. and Scott J. Kenyon (2008). "Planet Formation around Stars of Various Masses: The Snow Line and the Frequency of Giant Planets". In: *ApJ* 673.1, pp. 502–512. DOI: [10.1086/524130](https://doi.org/10.1086/524130). arXiv: [0710.1065](https://arxiv.org/abs/0710.1065) [astro-ph].
- Kirk, Brian et al. (2016). "Kepler Eclipsing Binary Stars. VII. The Catalog of Eclipsing Binaries Found in the Entire Kepler Data Set". In: *AJ* 151, p. 68. DOI: [10.3847/0004-6256/151/3/68](https://doi.org/10.3847/0004-6256/151/3/68).
- Kjeldsen, H. and T. R. Bedding (1995). "Amplitudes of stellar oscillations: the implications for asteroseismology." In: *A&A* 293, pp. 87–106. arXiv: [astro-ph/9403015](https://arxiv.org/abs/astro-ph/9403015) [astro-ph].

- Kjeldsen, H. and T. R. Bedding (2011). "Amplitudes of solar-like oscillations: a new scaling relation". In: *A&A* 529, p. L8. DOI: [10.1051/0004-6361/201116789](https://doi.org/10.1051/0004-6361/201116789).
- Kjeldsen, Hans et al. (2008). "The Amplitude of Solar Oscillations Using Stellar Techniques". In: *ApJ* 682.2, pp. 1370–1375. DOI: [10.1086/589142](https://doi.org/10.1086/589142). arXiv: [0804.1182](https://arxiv.org/abs/0804.1182) [astro-ph].
- Kopparapu, Ravi Kumar et al. (2013). "Habitable Zones around Main-sequence Stars: New Estimates". In: *ApJ* 765.2, 131, p. 131. DOI: [10.1088/0004-637X/765/2/131](https://doi.org/10.1088/0004-637X/765/2/131). arXiv: [1301.6674](https://arxiv.org/abs/1301.6674) [astro-ph.EP].
- Kozai, Y. (1962). "Secular perturbations of asteroids with high inclination and eccentricity". In: *AJ* 67.9, pp. 591–598. DOI: [10.1086/108790](https://doi.org/10.1086/108790).
- Kürster, M. et al. (2003). "The low-level radial velocity variability in Barnard's star (= GJ 699). Secular acceleration, indications for convective redshift, and planet mass limits". In: *A&A* 403, pp. 1077–1087. DOI: [10.1051/0004-6361:20030396](https://doi.org/10.1051/0004-6361:20030396). arXiv: [astro-ph/0303528](https://arxiv.org/abs/astro-ph/0303528) [astro-ph].
- Kustner, F. (1908). "Radial velocities of 99 stars of the second and third spectral classes observed at Bonn". In: *ApJ* 27.5, pp. 301–324.
- Laughlin, G. and J. E. Chambers (2002). "Extrasolar Trojans: The Viability and Detectability of Planets in the 1:1 Resonance". In: *AJ* 124, pp. 592–600. DOI: [10.1086/341173](https://doi.org/10.1086/341173). eprint: [astro-ph/0204091](https://arxiv.org/abs/astro-ph/0204091).
- Lee, B.-C. et al. (2015). "Search for exoplanet around northern circumpolar stars. Four planets around HD 11755, HD 12648, HD 24064, and 8 Ursae Minoris". In: *A&A* 584, A79. DOI: [10.1051/0004-6361/201527076](https://doi.org/10.1051/0004-6361/201527076).
- Leeuwen, F. van (2007). "Validation of the new Hipparcos reduction". In: *A&A* 474.2, pp. 653–664. DOI: [10.1051/0004-6361:20078357](https://doi.org/10.1051/0004-6361:20078357).
- Leleu, A. et al. (2019). "Co-orbital exoplanets from close-period candidates: the TOI-178 case". In: *A&A* 624, A46, A46. DOI: [10.1051/0004-6361/201834901](https://doi.org/10.1051/0004-6361/201834901). arXiv: [1901.07250](https://arxiv.org/abs/1901.07250) [astro-ph.EP].
- Leleu, Adrien, Philippe Robutel, and Alexandre C. M. Correia (2015). "Detectability of quasi-circular co-orbital planets. Application to the radial velocity technique". In: *A&A* 581, A128, A128. DOI: [10.1051/0004-6361/201526175](https://doi.org/10.1051/0004-6361/201526175). arXiv: [1509.02276](https://arxiv.org/abs/1509.02276) [astro-ph.EP].
- Li, Chih-Hao et al. (2008). "A laser frequency comb that enables radial velocity measurements with a precision of 1cm s^{-1} ". In: *Nature* 452.7187, pp. 610–612. DOI: [10.1038/nature06854](https://doi.org/10.1038/nature06854). arXiv: [0804.0955](https://arxiv.org/abs/0804.0955) [astro-ph].
- Lidov, M.L. (1962). "The evolution of orbits of artificial satellites of planets under the action of gravitational perturbations of external bodies". In: *Planetary and Space Science* 9.10, pp. 719–759. DOI: [10.1016/0032-0633\(62\)90129-0](https://doi.org/10.1016/0032-0633(62)90129-0).
- Lillo-Box, J. et al. (2018a). "The TROY project. II. Multi-technique constraints on extrojans in nine planetary systems". In: *A&A* 618, A42, A42. DOI: [10.1051/0004-6361/201833312](https://doi.org/10.1051/0004-6361/201833312). arXiv: [1807.00773](https://arxiv.org/abs/1807.00773) [astro-ph.EP].
- Lillo-Box, J. et al. (2018b). "The TROY project: Searching for co-orbital bodies to known planets. I. Project goals and first results from archival radial velocity". In: *A&A* 609, A96, A96. DOI: [10.1051/0004-6361/201730652](https://doi.org/10.1051/0004-6361/201730652). arXiv: [1710.01138](https://arxiv.org/abs/1710.01138) [astro-ph.EP].
- Lithwick, Y. and S. Naoz (2011). "The Eccentric Kozai Mechanism for a Test Particle". In: *ApJ* 742.2, 94. DOI: [10.1088/0004-637X/742/2/94](https://doi.org/10.1088/0004-637X/742/2/94).
- Liu, Y.-J. et al. (2008). "A Substellar Companion to the Intermediate-Mass Giant 11 Comae". In: *ApJ* 672.1, pp. 553–557. DOI: [10.1086/523297](https://doi.org/10.1086/523297).
- Lovis, C. and D. Fischer (2011). *Exoplanets: Radial Velocity Techniques for Exoplanets*. Vol. edited by S. Seager. University of Arizona Press, pp. 27–53. ISBN: 978-0-8165-2945-2.

- Luque, R. et al. (2019). "Precise radial velocities of giant stars. XIII. A second Jupiter orbiting in 4:3 resonance in the 7 CMa system". In: *A&A* 631, A136, A136. DOI: [10.1051/0004-6361/201936464](https://doi.org/10.1051/0004-6361/201936464). arXiv: [1910.05853](https://arxiv.org/abs/1910.05853) [astro-ph.EP].
- Maciejewski, Gracjan et al. (2020). "An Apparently Eccentric Orbit of the Exoplanet WASP-12 b as a Radial Velocity Signature of Planetary-induced Tides in the Host Star". In: *ApJ* 889.1, 54, p. 54. DOI: [10.3847/1538-4357/ab5e87](https://doi.org/10.3847/1538-4357/ab5e87). arXiv: [1912.01360](https://arxiv.org/abs/1912.01360) [astro-ph.EP].
- Marconi, A. et al. (2021). "HIRES, the High-resolution Spectrograph for the ELT". In: *The Messenger* 182, pp. 27–32. DOI: [10.18727/0722-6691/5219](https://doi.org/10.18727/0722-6691/5219). arXiv: [2011.12317](https://arxiv.org/abs/2011.12317) [astro-ph.IM].
- Marcy, Geoffrey W. and R. P. Butler (1992). "Precision Radial Velocities with an Iodine Absorption cell". In: *PASP* 104, p. 270. DOI: [10.1086/132989](https://doi.org/10.1086/132989).
- Marsh, T. R. (1989). "The Extraction of Highly Distorted Spectra". In: *PASP* 101, p. 1032. DOI: [10.1086/132570](https://doi.org/10.1086/132570).
- Massarotti, Alessandro et al. (2008). "Rotational and radial velocities for a sample of 761 Hipparcos giants and the role of binarity". In: *AJ* 135.1, pp. 209–231. DOI: [10.1088/0004-6256/135/1/209](https://doi.org/10.1088/0004-6256/135/1/209).
- Mayor, M. et al. (2003). "Setting New Standards with HARPS". In: *The Messenger* 114, pp. 20–24.
- Mayor, Michel and Didier Queloz (1995). "A Jupiter-mass companion to a solar-type star". In: *Nature* 378.6555, pp. 355–359. DOI: [10.1038/378355a0](https://doi.org/10.1038/378355a0).
- McMillan, R. S. et al. (1992). "Variation of the Radial and Velocity of Epsilon Cygni A". In: *PASP* 104, pp. 1173–1176.
- Mitchell, David S. et al. (2013). "Precise radial velocities of giant stars. V. A brown dwarf and a planet orbiting the K giant stars τ Geminorum and 91 Aquarii". In: *A&A* 555, A87. DOI: [10.1051/0004-6361/201321714](https://doi.org/10.1051/0004-6361/201321714).
- Montes, D. et al. (2018). "Calibrating the metallicity of M dwarfs in wide physical binaries with F-, G-, and K-primaries – I: High-resolution spectroscopy with HERMES: stellar parameters, abundances, and kinematics". In: *MNRAS* 479.1, pp. 1332–1382. DOI: [10.1093/mnras/sty1295](https://doi.org/10.1093/mnras/sty1295).
- Morais, M.H.M. and A.C.M. Correia (2008). "Stellar wobble caused by a binary system: Can it really be mistaken as an extra-solar planet?" In: *A&A* 491.3, pp. 899–906. DOI: [10.1051/0004-6361:200810741](https://doi.org/10.1051/0004-6361:200810741).
- (2011). "Stellar wobble caused by a nearby binary system: eccentric and inclined orbits". In: *A&A* 525, A152. DOI: [10.1051/0004-6361/201014812](https://doi.org/10.1051/0004-6361/201014812).
- (2012). "Precession due to a close binary system: an alternative explanation for ν -Octantis?" In: *MNRAS* 419.4, pp. 3447–3456. DOI: [10.1111/j.1365-2966.2011.19986.x](https://doi.org/10.1111/j.1365-2966.2011.19986.x).
- Mosser, B. and T. Appourchaux (2009). "On detecting the large separation in the autocorrelation of stellar oscillation times series". In: *A&A* 508, pp. 877–887. DOI: [10.1051/0004-6361/200912944](https://doi.org/10.1051/0004-6361/200912944). arXiv: [0909.0782](https://arxiv.org/abs/0909.0782) [astro-ph.SR].
- Nobel Prize Outreach AB (2019). "*The Nobel Prize in Physics 2019*". URL: <https://www.nobelprize.org/prizes/physics/2019/summary/> (visited on 05/09/2021).
- Oja, T. (1993). "UBV photometry of stars whose positions are accurately known. VII." In: *A&A* 100, pp. 591–592.
- Olling, Rob (2004). "Data Rate, S/N & Spectroscopy for the OBSS-A/B/C Concepts". In: *USNO/USRA, Washington, DC*. URL: https://www.astro.umd.edu/~olling/FAME/otm_spec_rpo_2004_01.pdf.
- Ortiz, Mauricio et al. (2016). "Precise radial velocities of giant stars. IX. HD 59686 Ab: a massive circumstellar planet orbiting a giant star in a 13.6 au eccentric binary system". In: *A&A* 595, A55. DOI: [10.1051/0004-6361/201628791](https://doi.org/10.1051/0004-6361/201628791).

- Pepe, F. et al. (2002). “The CORALIE survey for southern extra-solar planets VII. Two short-period Saturnian companions to HD 108147 and HD 168746”. In: *A&A* 388, pp. 632–638. DOI: [10.1051/0004-6361:20020433](https://doi.org/10.1051/0004-6361:20020433). arXiv: [astro-ph/0202457](https://arxiv.org/abs/astro-ph/0202457) [astro-ph].
- Pepe, F. et al. (2021). “ESPRESSO at VLT. On-sky performance and first results”. In: *A&A* 645, A96, A96. DOI: [10.1051/0004-6361/202038306](https://doi.org/10.1051/0004-6361/202038306). arXiv: [2010.00316](https://arxiv.org/abs/2010.00316) [astro-ph.IM].
- Petigura, Erik A., Andrew W. Howard, and Geoffrey W. Marcy (2013). “Prevalence of Earth-size planets orbiting Sun-like stars”. In: *Proceedings of the National Academy of Science* 110.48, pp. 19273–19278. DOI: [10.1073/pnas.1319909110](https://doi.org/10.1073/pnas.1319909110). arXiv: [1311.6806](https://arxiv.org/abs/1311.6806) [astro-ph.EP].
- Pierce, A. Keith (1965). “Construction of a Bowen Image Slicer”. In: *PASP* 77.456, p. 216. DOI: [10.1086/128199](https://doi.org/10.1086/128199).
- Piskunov, N. E. and J. A. Valenti (2002). “New algorithms for reducing cross-dispersed echelle spectra”. In: *A&A* 385, pp. 1095–1106. DOI: [10.1051/0004-6361:20020175](https://doi.org/10.1051/0004-6361:20020175).
- Quirrenbach, A. et al. (2016). “CARMENES: an overview six months after first light”. In: *Ground-based and Airborne Instrumentation for Astronomy VI*. Ed. by Christopher J. Evans, Luc Simard, and Hideki Takami. Vol. 9908. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, p. 990812. DOI: [10.1117/12.2231880](https://doi.org/10.1117/12.2231880).
- Quirrenbach, Andreas (2006). “Detection and Characterization of Extrasolar Planets”. In: *Saas-Fee Advanced Course 31: Extrasolar planets*. Ed. by Didier Queloz et al., pp. 1–242. DOI: [10.1007/978-3-540-31470-7_1](https://doi.org/10.1007/978-3-540-31470-7_1).
- Quirrenbach, Andreas et al. (2019). “Precise radial velocities of giant stars. XI. Two brown dwarfs in 6:1 mean motion resonance around the K giant star ν Ophiuchi”. In: *A&A* 624, A18, A18. DOI: [10.1051/0004-6361/201834423](https://doi.org/10.1051/0004-6361/201834423). arXiv: [1904.03557](https://arxiv.org/abs/1904.03557) [astro-ph.SR].
- Ramm, D. J. et al. (2009). “Spectroscopic orbits for K giants β Reticuli and ν Octantis: what is causing a low-amplitude radial velocity resonant perturbation in ν Oct?” In: *MNRAS* 394.3, pp. 1695–1710. DOI: [10.1111/j.1365-2966.2009.14459.x](https://doi.org/10.1111/j.1365-2966.2009.14459.x).
- Ramm, D. J. et al. (2016). “The conjectured S-type retrograde planet in ν Octantis: more evidence including four years of iodine-cell radial velocities”. In: *MNRAS* 460.4, pp. 3706–3719. DOI: [10.1093/mnras/stw1106](https://doi.org/10.1093/mnras/stw1106).
- Rasmussen, René Tronsgaard (2016). “Precise measurements of stellar radial velocities. A new, open-source pipeline for iodine cell doppler spectroscopy”. Master thesis. Department of Physics and Astronomy, Aarhus University.
- Reffert, Sabine et al. (2006). “Precise Radial Velocities of Giant Stars. II. Pollux and Its Planetary Companion”. In: *ApJ* 652.1, pp. 661–665. DOI: [10.1086/507516](https://doi.org/10.1086/507516).
- Reffert, Sabine et al. (2015). “Precise radial velocities of giant stars. VII. Occurrence rate of giant extrasolar planets as a function of mass and metallicity”. In: *A&A* 574, A116. DOI: [10.1051/0004-6361/201322360](https://doi.org/10.1051/0004-6361/201322360).
- Reichert, Katja et al. (2019). “Precise radial velocities of giant stars. XII. Evidence against the proposed planet Aldebaran b”. In: *A&A* 625, A22. DOI: [10.1051/0004-6361/201834028](https://doi.org/10.1051/0004-6361/201834028).
- Rein, H. and S. F. Liu (2012). “REBOUND: an open-source multi-purpose N-body code for collisional dynamics”. In: *A&A* 537, A128, A128. DOI: [10.1051/0004-6361/201118085](https://doi.org/10.1051/0004-6361/201118085). arXiv: [1110.4876](https://arxiv.org/abs/1110.4876) [astro-ph.EP].
- Reiners, A., R. K. Banyal, and R. G. Ulbrich (2014). “A laser-lock concept to reach cm s^{-1} -precision in Doppler experiments with Fabry-Pérot wavelength calibrators”. In: *A&A* 569, A77, A77. DOI: [10.1051/0004-6361/201424099](https://doi.org/10.1051/0004-6361/201424099). arXiv: [1408.6111](https://arxiv.org/abs/1408.6111) [astro-ph.IM].

- Reiners, A. et al. (2010). “Detecting Planets Around Very Low Mass Stars with the Radial Velocity Method”. In: *ApJ* 710.1, pp. 432–443. DOI: [10.1088/0004-637X/710/1/432](https://doi.org/10.1088/0004-637X/710/1/432).
- Roth, Kai Daniel (2019). “Dome Automation and Guiding Routine for the Waltz Telescope”. Master thesis. Ruprecht-Karls Universität Heidelberg.
- Saio, Hideyuki et al. (2015). “Oscillatory convective modes in red giants: a possible explanation of the long secondary periods”. In: *MNRAS* 452.4, pp. 3863–3868. DOI: [10.1093/mnras/stv1587](https://doi.org/10.1093/mnras/stv1587).
- Saio, Hideyuki et al. (2018). “Theory and evidence of global Rossby waves in upper main-sequence stars: r-mode oscillations in many Kepler stars”. In: *MNRAS* 474.2, pp. 2774–2786. DOI: [10.1093/mnras/stx2962](https://doi.org/10.1093/mnras/stx2962). arXiv: [1711.04908](https://arxiv.org/abs/1711.04908) [astro-ph.SR].
- Schneider, Klaus (2013). *DELTA CODE - Die ultimative Teleskopsteuerung*. URL: https://www.apm-telescopes.net/downloads/Handbuch_Deltacode.pdf.
- ScopeDome (2016). *Manual - ScopeDome USB Card ver. 2.1*. URL: https://www.scopedome.com/?page_id=5804&lang=en.
- Shporer, Avi et al. (2016). “Radial velocity monitoring of Kepler heartbeat stars”. In: *ApJ* 829.34. DOI: [10.3847/0004-637X/829/1/34](https://doi.org/10.3847/0004-637X/829/1/34).
- Soszyński, I. et al. (2009). “The Optical Gravitational Lensing Experiment. The OGLE-III Catalog of Variable Stars. IV. Long-Period Variables in the Large Magellanic Cloud”. In: *ACTAA* 59.3, pp. 239–253. arXiv: [0910.1354](https://arxiv.org/abs/0910.1354) [astro-ph.SR]. URL: <https://ui.adsabs.harvard.edu/abs/2009AcA...59..239S>.
- Späth, Dane Marvin (2019). “Software for Instrument- and Telescope Control in Support of the Waltz Radial Velocity Survey”. Master thesis. Ruprecht-Karls Universität Heidelberg.
- Stello, Dennis et al. (2017). “Asteroseismic masses of retired planet-hosting A-stars using SONG”. In: *MNRAS* 472.4, pp. 4110–4116. DOI: [10.1093/mnras/stx2295](https://doi.org/10.1093/mnras/stx2295).
- Stock, Stephan, Sabine Reffert, and Andreas Quirrenbach (2018). “Precise radial velocities of giant stars. X. Bayesian stellar parameters and evolutionary stages for 372 giant stars from the Lick planet search”. In: *A&A* 616, A33. DOI: [10.1051/0004-6361/201833111](https://doi.org/10.1051/0004-6361/201833111).
- Struve, O. (1952). “Proposal for a project of high-precision stellar radial velocity work”. In: *The Observatory* 72, pp. 199–200. URL: <https://ui.adsabs.harvard.edu/abs/1952Obs...72..199S>.
- Stürmer, Julian et al. (2016). “A rubidium traced white-light etalon calibrator for MAROON-X”. In: *Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation II*. Ed. by Ramón Navarro and James H. Burge. Vol. 9912. International Society for Optics and Photonics. SPIE, 777–785. DOI: [10.1117/12.2232865](https://doi.org/10.1117/12.2232865). URL: <https://doi.org/10.1117/12.2232865>.
- Tal-Or, L. et al. (2018). “The CARMENES search for exoplanets around M dwarfs: Radial-velocity variations of active stars in visual-channel spectra”. In: *A&A* 614.A122. DOI: [10.1051/0004-6361/201732362](https://doi.org/10.1051/0004-6361/201732362).
- Tala, M. et al. (2016). “A high-resolution spectrograph for the 72cm Waltz Telescope at Landessternwarte, Heidelberg”. In: *Ground-based and Airborne Instrumentation for Astronomy VI*. Ed. by Christopher J. Evans, Luc Simard, and Hideki Takami. Vol. 9908. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 99086O. DOI: [10.1117/12.2232730](https://doi.org/10.1117/12.2232730). arXiv: [1608.06090](https://arxiv.org/abs/1608.06090) [astro-ph.IM].
- Tala Pinto, Marcelo et al. (2020). “Precise radial velocities of giant stars. XIV. Evidence of planetary companions around HD 25723, 17 Sco, 3 Cnc, and 44 UMa”. In: *A&A* 644, A1, A1. DOI: [10.1051/0004-6361/202038285](https://doi.org/10.1051/0004-6361/202038285).
- Tala Pinto, Marcelo Said (2019). “Perspectives of the Radial Velocity Method: Physical Modeling of the Wavelength Solution & Exoplanetary Detections around Giant

- Stars". Dissertation. Ruprecht-Karls Universität Heidelberg. DOI: <https://doi.org/10.11588/heidok.00026834>.
- Tamayo, Daniel et al. (2020). "Predicting the long-term stability of compact multi-planet systems". In: *Proceedings of the National Academy of Science* 117.31, pp. 18194–18205. DOI: [10.1073/pnas.2001258117](https://doi.org/10.1073/pnas.2001258117). arXiv: [2007.06521](https://arxiv.org/abs/2007.06521) [astro-ph.EP].
- Tan, X. et al. (2013). "Characterizing the orbital and dynamical state of the HD 82943 planetary system with Keck radial velocity data". In: *ApJ* 777, p. 101. DOI: [10.1088/0004-637X/777/2/101](https://doi.org/10.1088/0004-637X/777/2/101).
- Trifonov, Trifon (2019). *The Exo-Striker: Transit and radial velocity interactive fitting tool for orbital analysis and N-body simulations*. ascl: [1906.004](https://ui.adsabs.harvard.edu/abs/2019ascl.soft06004T). URL: <https://ui.adsabs.harvard.edu/abs/2019ascl.soft06004T>.
- Trifonov, Trifon et al. (2014). "Precise radial velocities of giant stars. VI. A possible 2:1 resonant planet pair around the K giant star η Ceti". In: *A&A* 568, A64. DOI: [10.1051/0004-6361/201322885](https://doi.org/10.1051/0004-6361/201322885).
- Trifonov, Trifon et al. (2018). "Dynamical Analysis of the Circumprimary Planet in the Eccentric Binary System HD 59686". In: *AJ* 155.4, p. 174. DOI: [10.3847/1538-3881/aab439](https://doi.org/10.3847/1538-3881/aab439).
- Udalski, A., M. Kubiak, and M. Szymanski (1997). "Optical Gravitational Lensing Experiment. OGLE-2 – the Second Phase of the OGLE Project". In: *ACTAA* 47, pp. 319–344. arXiv: [astro-ph/9710091](https://arxiv.org/abs/astro-ph/9710091) [astro-ph].
- Urry, M. and G. Reichert (1988). "Gaussian Extraction Routine for IUE Data". In: *NASA IUE Newsletter* 34.95.
- Valenti, Jeff A., R. Paul Butler, and Geoffrey W. Marcy (1995). "Determining Spectrometer Instrumental Profiles Using FTS Reference Spectra". In: *PASP* 107, p. 966. DOI: [10.1086/133645](https://doi.org/10.1086/133645).
- van de Kamp, P. (1969). "Alternate dynamical analysis of Barnard's star." In: *AJ* 74, pp. 757–759. DOI: [10.1086/110852](https://doi.org/10.1086/110852).
- VanderPlas, Jacob T. (2018). "Understanding the Lomb-Scargle Periodogram". In: *ApJS* 236.1, 16, p. 16. DOI: [10.3847/1538-4365/aab766](https://doi.org/10.3847/1538-4365/aab766). arXiv: [1703.09824](https://arxiv.org/abs/1703.09824) [astro-ph.IM].
- Villaver, Eva and Mario Livio (2009). "The Orbital Evolution of Gas Giant Planets Around Giant Stars". In: *ApJ* 705.1, pp. L81–L85. DOI: [10.1088/0004-637X/705/1/L81](https://doi.org/10.1088/0004-637X/705/1/L81). arXiv: [0910.2396](https://arxiv.org/abs/0910.2396) [astro-ph.EP].
- Vogt, S. S. et al. (1994). "HIRES: the high-resolution echelle spectrometer on the Keck 10-m Telescope". In: *Instrumentation in Astronomy VIII*. Ed. by David L. Crawford and Eric R. Craine. Vol. 2198. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, p. 362. DOI: [10.1117/12.176725](https://doi.org/10.1117/12.176725).
- Vogt, Steven S. (1987). "The Lick Observatory Hamilton echelle spectrometer." In: *PASP* 99, pp. 1214–1228. DOI: [10.1086/132107](https://doi.org/10.1086/132107).
- Wallace, L. et al. (2011). "An Optical and Near-infrared (2958-9250 Å) Solar Flux Atlas". In: *ApJS* 195.1, 6, p. 6. DOI: [10.1088/0067-0049/195/1/6](https://doi.org/10.1088/0067-0049/195/1/6).
- Weiss, W. W. et al. (2014). "BRITe-Constellation: Nanosatellites for Precision Photometry of Bright Stars". In: *PASP* 126, pp. 573–585. DOI: [10.1086/677236](https://doi.org/10.1086/677236).
- Wicker, Mauritz Mats (2020). "Calibration Unit, Exposure Trigger and All-Sky Camera: Improvements to the Infrastructure of the Waltz Telescope". Bachelor thesis. Ruprecht-Karls Universität Heidelberg.
- Wittenmyer, Robert A. et al. (2011). "The Pan-Pacific Planet Search: A Southern Hemisphere Search for Planets Orbiting Evolved Massive Stars". In: *Planetary Systems Beyond the Main Sequence*. Ed. by Sonja Schuh, Horst Drechsel, and Ulrich Heber. Vol. 1331. American Institute of Physics Conference Series, pp. 117–121. DOI: [10.1063/1.3556191](https://doi.org/10.1063/1.3556191). arXiv: [1101.4052](https://arxiv.org/abs/1101.4052) [astro-ph.EP].

- Wittenmyer, Robert A. et al. (2020a). “Cool Jupiters greatly outnumber their toasty siblings: occurrence rates from the Anglo-Australian Planet Search”. In: *MNRAS* 492.1, pp. 377–383. DOI: [10.1093/mnras/stz3436](https://doi.org/10.1093/mnras/stz3436). arXiv: [1912.01821](https://arxiv.org/abs/1912.01821) [astro-ph.EP].
- Wittenmyer, Robert A. et al. (2020b). “The Pan-Pacific Planet Search - VIII. Complete results and the occurrence rate of planets around low-luminosity giants”. In: *MNRAS* 491.4, pp. 5248–5257. DOI: [10.1093/mnras/stz3378](https://doi.org/10.1093/mnras/stz3378). arXiv: [1911.11954](https://arxiv.org/abs/1911.11954) [astro-ph.EP].
- Wolszczan, A. and D. A. Frail (1992). “A planetary system around the millisecond pulsar PSR1257 + 12”. In: *Nature* 355.6356, pp. 145–147. DOI: [10.1038/355145a0](https://doi.org/10.1038/355145a0).
- Wood, P. R., E. A. Olivier, and S. D. Kawaler (2004). “Long Secondary Periods in Pulsating Asymptotic Giant Branch Stars: An Investigation of their Origin”. In: *ApJ* 604.2, pp. 800–816. DOI: [10.1086/382123](https://doi.org/10.1086/382123).
- Wood, P. R. et al. (1999). “MACHO observations of LMC red giants: Mira and semi-regular pulsators, and contact and semi-detached binaries”. In: *Asymptotic Giant Branch Stars, IAU Symposium #191* edited by T. Le Bertre, A. Lebre, and C. Waelkens, p. 151.
- Wright, J. T. and J. D. Eastman (2014). “Barycentric Corrections at 1 cm s^{-1} for Precise Doppler Velocities”. In: *PASP* 126.943, p. 838. DOI: [10.1086/678541](https://doi.org/10.1086/678541). arXiv: [1409.4774](https://arxiv.org/abs/1409.4774) [astro-ph.IM].
- Yu, Jie et al. (2018). “Predicting radial-velocity jitter induced by stellar oscillations based on Kepler data”. In: *MNRAS* 480.1, pp. L48–L53. DOI: [10.1093/mnrasl/sly123](https://doi.org/10.1093/mnrasl/sly123). arXiv: [1807.00096](https://arxiv.org/abs/1807.00096) [astro-ph.EP].
- Zechmeister, M. and M. Kürster (2009). “The generalised Lomb-Scargle periodogram”. In: *A&A* 496.2, pp. 577–584. DOI: [10.1051/0004-6361:200811296](https://doi.org/10.1051/0004-6361:200811296).
- Zechmeister, M. et al. (2018). “Spectrum radial velocity analyser (SERVAL). High-precision radial velocities and two alternative spectral indicators”. In: *A&A* 609, A12, A12. DOI: [10.1051/0004-6361/201731483](https://doi.org/10.1051/0004-6361/201731483). arXiv: [1710.10114](https://arxiv.org/abs/1710.10114) [astro-ph.IM].