

Universiteit
Leiden
The Netherlands

Painting with starlight : optical techniques for the high-contrast imaging of exoplanets

Wilby, M.J.

Citation
Wilby, M. J. (2018, November 27). *Painting with starlight : optical techniques for the high-contrast imaging of exoplanets*. Retrieved from https://hdl.handle.net/1887/67531

Version:	Not Applicable (or Unknown)
License:	Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden
Downloaded from:	https://hdl.handle.net/1887/67531

Note: To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The handle http://hdl.handle.net/1887/67531 holds various files of this Leiden University dissertation.

Author: Wilby, M.J.
Title: Painting with starlight : optical techniques for the high-contrast imaging of exoplanets
Issue Date: 2018-11-27

Painting with Starlight

Optical techniques for the high-contrast imaging of exoplanets

Proefschrift

ter verkrijging van
de graad van Doctor aan de Universiteit Leiden,
op gezag van Rector Magnificus prof. mr. C.J.J.M. Stolker,
volgens besluit van het College voor Promoties
te verdedigen op dinsdag 27 november 2018
klokke 12.30 uur

door

Michael James Wilby

geboren te
Hemel Hempstead, Verenigd Koninkrijk
in 1992

Promotor:	Prof. dr. Christoph Keller	
Co-promotor:	Dr. Matthew Kenworthy	
Promotiecommissie:	Prof. dr. Huub Röttgering	Universiteit Leiden (Sterrewacht)
	Prof. dr. Ignas Snellen	Universiteit Leiden (Sterrewacht)
	Prof. dr. Wolfgang Osten	Universität Stuttgart, Deutschland
	Dr. Silvania Pereira	Technische Universiteit Delft
	Prof. dr. Martin van Exter	Universiteit Leiden (LION)

ISBN: 978-94-028-1261-9

An electronic copy of this thesis can be found at https://openaccess.leidenuniv.nl

© Michael J. Wilby, 2018

Contents

1 Introduction **1**
- 1.1 Detecting and characterising exoplanets . . . 2
  - 1.1.1 Overview of exoplanet detection techniques . . . 2
  - 1.1.2 From detection to characterisation . . . 4
- 1.2 Planet formation . . . 7
  - 1.2.1 The planet-disk connection . . . 7
  - 1.2.2 Observational challenges to current planet formation theory . . . 9
  - 1.2.3 Observing proto-planetary disks . . . 10
- 1.3 Direct imaging of exoplanets . . . 11
  - 1.3.1 Challenges and rewards . . . 11
  - 1.3.2 Current state of the field . . . 13
  - 1.3.3 Outlook for next-generation instrumentation . . . 13
- 1.4 The high-contrast imager's toolbox . . . 16
  - 1.4.1 Coronagraphy . . . 16
  - 1.4.2 Adaptive optics . . . 18
  - 1.4.3 Image post-processing algorithms . . . 20
  - 1.4.4 Polarimetry . . . 22
  - 1.4.5 Spectroscopy . . . 24
- 1.5 This thesis: Focal-plane wavefront sensing techniques . . . 25
  - 1.5.1 Holographic wavefront sensing . . . 26
  - 1.5.2 Phase diversity wavefront sensing . . . 30
- 1.6 Thesis outline . . . 31

2 The coronagraphic Modal Wavefront Sensor: A hybrid focal-plane sensor for the high-contrast imaging of circumstellar environments **33**
- 2.1 Introduction . . . 34
  - 2.1.1 Scientific motivation . . . 34
  - 2.1.2 Holographic optics for focal-plane wavefront sensing . . . 35
  - 2.1.3 Content of paper . . . 36
- 2.2 Theory . . . 36
  - 2.2.1 The Holographic Modal Wavefront Sensor . . . 36
  - 2.2.2 Combination with an Apodizing Phase Plate coronagraph . . . 41
  - 2.2.3 Impact of multiplexing on mutual performance . . . 41
  - 2.2.4 Impact of structured telescope apertures . . . 43
- 2.3 Idealised performance simulations . . . 44
- 2.4 On-sky demonstration . . . 47
  - 2.4.1 Instrument design . . . 47
  - 2.4.2 Characterising HMWS on-sky response . . . 49
  - 2.4.3 Broadband wavefront sensing . . . 52
  - 2.4.4 Real-time atmospheric wavefront measurements . . . 52
- 2.5 Discussion & conclusions . . . 56

3 A 'Fast & Furious' solution to the low-wind effect for SPHERE at the VLT 59
3.1 Introduction . . . 60
3.1.1 The low-wind effect . . . 60
3.1.2 Proposal: Focal-plane wavefront sensing using the SPHERE DTTS 62
3.2 The Fast & Furious algorithm . . . 63
3.3 Simulating F&F performance in SPHERE . . . 66
3.3.1 Simulation parameters . . . 66
3.3.2 Simulation results . . . 68
3.4 Discussion . . . 71
3.5 Conclusions . . . 72

4 Laboratory verification of 'Fast & Furious' phase diversity: Towards controlling the low-wind effect in the SPHERE instrument 75
4.1 Introduction . . . 76
4.2 Methodology . . . 80
4.2.1 The Fast & Furious algorithm . . . 80
4.2.2 Implementing F&F in SPHERE-like simulations . . . 82
4.2.3 Laboratory verification of F&F on the MITHIC bench . . . 84
4.3 Results . . . 86
4.3.1 Comparison of MITHIC bench and simulation results . . . 86
4.3.2 Simulated low-S/N performance of F&F . . . 92
4.4 Discussion . . . 94
4.5 Conclusions . . . 98

5 These are not the voids you're looking for: Using Ceres to calibrate SPHERE-IRDIS coronagraphic DPI observations of TW Hydrae 101
5.1 Introduction . . . 102
5.2 Ceres observations and data reduction . . . 103
5.2.1 Observations . . . 103
5.2.2 Data reduction . . . 104
5.2.3 Analysis . . . 106
5.3 Modelling the SPHERE ALC Coronagraph System . . . 108
5.3.1 Description of the coronagraphic system model . . . 108
5.3.2 Diffraction-limited model . . . 109
5.3.3 Effects of residual atmospheric turbulence . . . 111
5.3.4 Additional considerations . . . 114
5.4 Calibrating Protoplanetary Disk Observations . . . 118
5.4.1 Polarimetric image formation and instrumental polarisation . . . 118
5.4.2 Convolutional depolarisation: an additional consideration for DPI-mode observations . . . 119
5.4.3 A recipe for disk calibration . . . 121
5.5 TW Hydrae: calibration of a real disk observation . . . 122
5.5.1 Background and observations . . . 122
5.5.2 A simple forward model of the TW Hydrae disk . . . 123
5.5.3 Assessing the accuracy of ALC extinction calibration by normalisation . . . 127
5.6 Conclusions . . . 129

Bibliography **132**

List of Abbreviations **141**

6 Nederlandse Samenvatting **145**
6.1 Hoe vinden we exoplaneten? . . . 145
6.1.1 Indirecte detectie methodes . . . 145
6.1.2 Directe detectie . . . 146
6.2 Hoe characteriseren we exoplaneten? . . . 146
6.3 Protoplanetaire schijven en planeet formatie . . . 147
6.4 Dit proefschrift . . . 148
6.5 Algemene conclusies . . . 150

7 English Summary **151**
7.1 How do we find exoplanets? . . . 151
7.1.1 Indirect detection methods . . . 151
7.1.2 Direct imaging . . . 151
7.2 How do we characterise exoplanets? . . . 152
7.3 Protoplanetary disks and planet formation . . . 153
7.4 This thesis . . . 154
7.5 Overall conclusions . . . 155

List of Publications **157**

Curriculum Vitae **159**

Acknowledgements **161**

Chapter 1

Introduction

"The history of astronomy is a history of receding horizons."

– Edwin Hubble

Prior to the 1990s, the study of planetary systems was exclusively confined to our own solar system. Theories on planet formation and evolution have therefore traditionally focussed on explaining all aspects of this single data point, which in turn have shaped our expectations regarding extra-solar planetary (exoplanetary) systems. A subsequent explosion in the number of detected exoplanets has since revolutionised the field; at the time of writing there are now known to be 3,791 confirmed planetary-mass companions in 2,836 planetary systems[1], with this number increasing on an almost-daily basis. This has generated a number of surprises which have forced us to re-evaluate many established aspects of planetary physics. The most prominent of these is the discovery of hot Jupiters; massive gas-giant planets occupying infernally small orbits well inside that of Mercury in our own solar system, which are not expected to be able to form at these locations.

The detailed study of planetary systems orbiting other stars now promises to provide significant insights into some of the most fundamental philosophical questions of *life, the universe and everything*. How unique is the solar system? How did the Earth form? And perhaps most importantly, are we alone in the universe?[2]

Dedicated exoplanet surveys such as the Kepler mission have recently revealed that planets are in fact abundant in our galactic neighbourhood; somewhere between 50-100% of sun-like stars host some form of planetary system (Winn & Fabrycky 2015), while approximately 10-20% are expected to host one or more rocky planets in the Earth-mass category (Fressin et al. 2013; Petigura et al. 2013). By observing young stellar systems at different stages in the process of formation, we can also build up a timeline of how our own solar system, and in turn the Earth, most likely formed.

The exoplanet community is now nearing the capability not only to detect nearby Earth-like planets, but also to determine the composition and dynamics of their atmospheres. This raises the intriguing possibility of detecting biosignatures, such as high atmospheric oxygen content or even industrial pollutants, which could indicate the presence of extraterrestrial life at varying levels of development.

This thesis is primarily concerned with helping to overcome some of the remaining technological hurdles in exoplanet-related instrumentation, to make these goals a reality within the next decade.

[1] Source: exoplanet.eu/, correct as of the 7th of June, 2018

[2] The answers to which, despite the claims of Adams (1986), are significantly more complex than the number *42*.

1.1 Detecting and characterising exoplanets

1.1.1 Overview of exoplanet detection techniques

The first confirmed detection of an exoplanetary system was made in 1992 by Aleksander Wolszczan and Dale Frail, who identified two planetary-mass bodies around the millisecond pulsar *PSR 1257+12* (Wolszczan & Frail 1992). The existence of these planets was revealed by the reflex motion of the host star orbiting around the common centre of mass of the star-planet system, which induces a measurable periodic delay in the arrival times of radio pulses emitted by the dead star.

Although this technique, dubbed *pulsar timing*, is limited to detecting planets around this exotic class of star, the same principle of reflex orbital motion also induces time-varying Doppler shifts in the light spectrum of all stars which host planetary systems. The *radial velocity (RV)* detection technique exploits this fact to look for the corresponding periodic shifts in wavelength of features in the stellar spectrum over the course of a planet's orbit. The first detection of an exoplanet around a main-sequence star was achieved using this technique just a few years later, with the discovery of 51 Pegasi b: a 0.5 Jupiter-mass planet orbiting a sun-like G2-type star (Mayor & Queloz 1995). Any similarity with the solar system ends here however: this object orbits its host star within just four days at a separation six times smaller than the orbit of Mercury, making it an archetypal member of the hot-Jupiter class of exoplanets. The fact that such an extreme object was one of the first detected exoplanets is however not surprising; the size of the RV signal is directly proportional to the mass of the planet, while a short orbital period helps to generate a statistically significant signal within a relatively short observing window.

Figure 1.1 provides an overview of currently known exoplanets by plotting their orbital semi-major axes[3] against the best estimate of planetary mass, with each point coloured according to the method with which it was originally detected. Three major populations can be seen in this way: the "Jupiters", gas-giants on temperate orbits beyond approximately one astronomical unit (AU); the aforementioned "hot Jupiters" on extremely tight orbits; and "super-Earths" with masses of the order of one hundredth of a Jupiter mass (three Earth masses). Planets in this last class are typically detected with orbital separations similar to Mercury (0.1 AU), although this is primarily due to the sensitivity limits of current surveys (e.g. Howard et al. 2010): planetary population synthesis simulations predict that this distribution should extend smoothly to encompass the parameter space in which the solar system planets exist (Mordasini et al. 2009).

It is apparent that the hot-Jupiter and super-Earth populations feature a large number of detections made using the *transit* method, where photometric monitoring of the star's light curve reveals characteristic, periodically repeating dips due to the planet passing in front of its host star. This technique is also heavily biased towards detecting giant close-in planets: the size of the transit signal scales with planet radius squared, while the likelihood of a planet's orbit being sufficiently well aligned with Earth for it to transit the host star becomes rapidly more improbable with increasing orbital distance. The source of the majority of these transit detections is the highly successful Kepler space observatory mission (Borucki et al. 2010), which provided continual high-precision photometric monitoring of stars in a 116 square degree

[3] Equivalent to the orbital radius for a circular orbit.

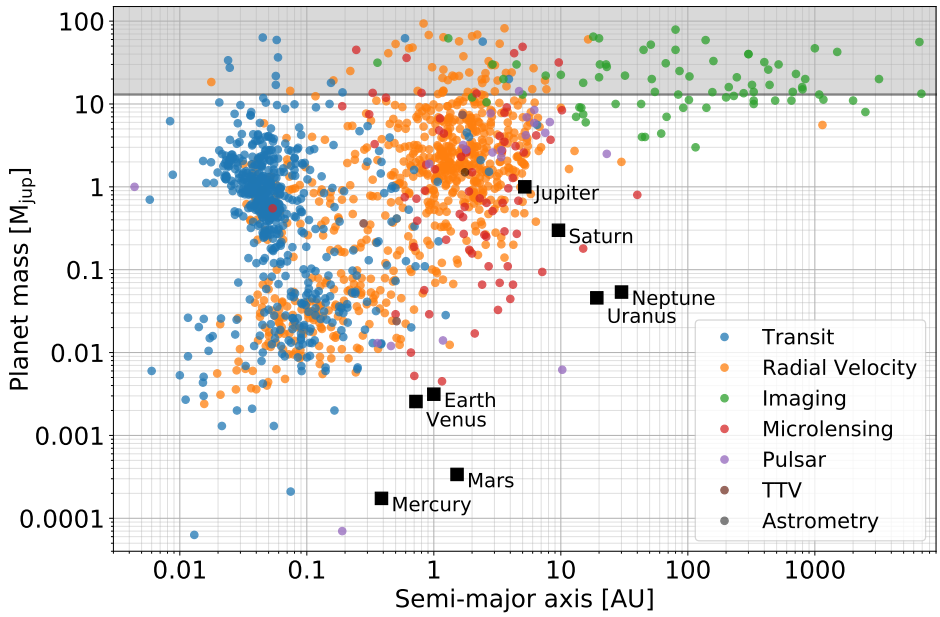


Figure 1.1: Logarithmic plot orbital semi-major axis against mass for all confirmed sub-stellar companions below 100 Jupiter masses (M_{jup}), coloured by discovery method. The horizontal grey line and shaded region denotes the 13 M_{jup} deuterium-burning mass limit, above which these objects begin to more closely resemble brown dwarf "failed stars" rather than planets (Spiegel et al. 2011). The properties of the solar system planets are also plotted for comparison. Data from http://exoplanet.eu/ (May 2018).

field of view in the constellation of Cygnus, between 2009 and 2013. This alone is currently responsible for over $1,500$ confirmed planet detections, with more than $2,000$ candidates still awaiting confirmation.

The transit and RV techniques are are by far the most productive methods to date in terms of number of discoveries, but are by no means the only methods by which exoplanets may be detected. High-precision *astrometry* can be used to look for the spatial rather than spectral modulation of starlight due to reflex orbital motion induced by a planetary system, achieved by imaging the star's changing position in the plane of the sky. Only a handful of planets have so far been detected via this method, however with the recent launch and first two data releases of the all-sky astrometry satellite GAIA (Gaia Collaboration et al. 2016a,b; Molnár et al. 2018), this number is expected to increase significantly in the coming years.

Microlensing events occur when a stellar system containing a planet passes in front of a background source, with both foreground objects acting as gravitational lenses and weakly amplifying its brightness. This method has detected some of the closest known analogues to solar system planets in terms of mass and orbital radius (e.g. Gould et al. 2014; Shvartzvald et al. 2017), but has the major drawback that lensing events are highly improbable, thus providing no opportunity for follow-up or further characterisation of the detected planet.

In systems with at least one known transiting exoplanet, the identification of *transit timing variations* (TTVs) in the signal, where the exact timing of the transit event is seen to vary from one orbit to the next, can be used to infer the existence of additional non-transiting planets in the system. This technique has proved capable of detecting

the gravitational perturbations of additional planets down to 7.5 Earth masses (Masuda 2014) and has also been proposed as a possible technique for detecting super-Galilean exomoons orbiting around giant transiting planets (Kipping 2009a,b).

All of the aforementioned detection methods are inherently indirect techniques, relying on effects induced by the planet on the light of the host star. *Direct imaging* techniques on the other hand attempt to resolve the light of the planet from that of the star, allowing us to study these planetary systems in great detail. As shown in Fig. 1.1, direct imaging probes a very different part of the exoplanet population to other techniques: currently it is only capable of detecting massive giant planets on extremely wide orbits, which are typically young and still hot from their formation and so particularly luminous at near-infrared wavelengths. This limitation is due to the severe technical challenge of teasing out faint planetary signals from a halo of starlight which is naturally many orders of magnitudes brighter. This task, also known as *high-contrast imaging* (HCI), only becomes more challenging when trying to detect signals closer in towards the host star. However, with rapid advances currently being made in instrumentation and data reduction techniques, the field is moving towards the capabilities required to image more conventional planetary systems, with an ultimate goal of directly resolving Earth-like planets around nearby stars.

The work contained in this thesis attempts to address some of the main instrumentation challenges associated with high-contrast imaging; the various benefits, difficulties and methodologies of this technique are therefore presented in significantly more detail in Sec. 1.3.

1.1.2 From detection to characterisation

Due to the limited amount of planetary information provided by most detection techniques, the exoplanet field has until recently broadly focussed on identifying the existence of planetary-mass companions and determining their orbital properties. A particular focus is made on the search for "Earth 2.0"; a planet with a sufficiently similar mass to Earth, orbiting in the habitable zone of a sun-like (F, G or K-type) star where liquid water may exist on the surface. For this reason the super-Earth class of exoplanets is of significant interest: not only are they currently the most similar known objects to Earth in terms of mass, but their nature is largely unknown. Are they rocky and Earth-like, or in fact more like low-density, atmosphere-dominated mini-Neptunes?

This question can already be partially addressed by determining the overall density of these objects, from which it is possible to constrain their bulk composition. This feat can be achieved via complementary transit and RV study, where the transit measurements constrain the planetary radius while RV provides the mass[4]. Fig. 1.2 shows the result of this characterisation for a subset of transiting exoplanets, directly compared to simulated mass-radius relationships for different idealised planetary compositions. The exoplanet population appears to undergo a smooth transition between gaseous (hydrogen-dominated) and solid (rock/iron-dominated) compositions at around ten Earth masses, which is consistent with the trend seen in the solar system. While the deduction of exact planetary composition in this way is inherently degener-

[4]Specifically, RV provides only a lower limit on the planetary mass, $m_{RV} = m_p \sin(i)$; this is degenerate with the (often unknown) system inclination i, since RV only measures the component of orbital motion along our line of sight. For transiting planets however the inclination must by necessity be very close to 90°, which in turn facilitates an accurate measurement of the true planetary mass m_p.

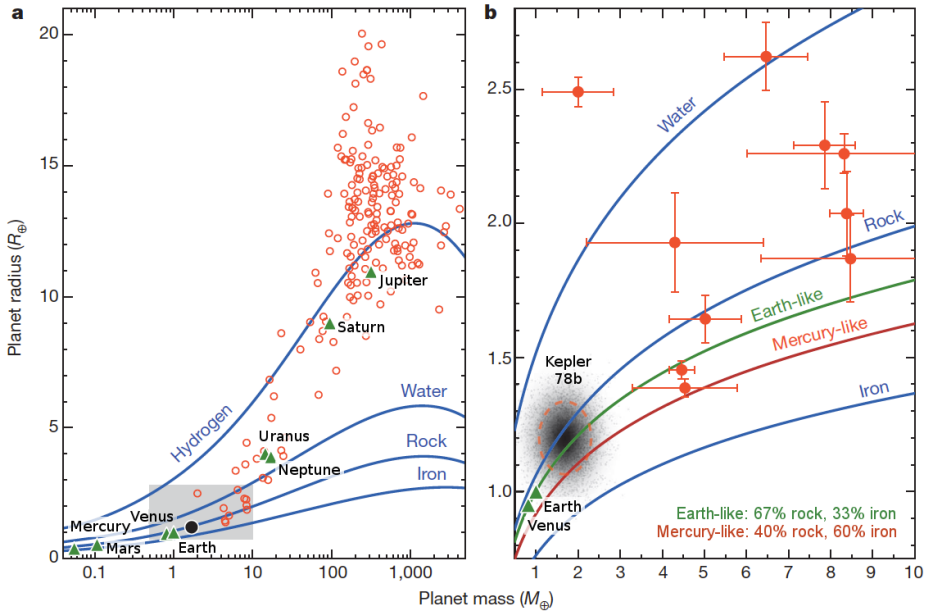


Figure 1.2: Planetary mass and radius (in terms of Earth mass and radii) of a subset of exoplanets which have been jointly characterised by transit and RV study (open circles). The over-plotted blue curves show the expected mass-radius relationships for idealised planetary models consisting entirely of the labelled constituent. The right panel shows a zoom-in on the non-gaseous super-Earth region, including error bars on detections and the then newly-discovered rocky Kepler-78b. Figure credit: Howard et al. (2013)

ate due to the sheer number of possible constituents, it nonetheless provides a useful indication of whether a given planet is likely to be most similar in nature to Earth, Neptune, or Jupiter.

Recent years have also seen the start of the next phase of characterisation of these exoplanetary systems: detailed single-object studies which are capable of determining planetary parameters beyond mass, radius and orbital separation, and hence can provide significant insights into the true nature of these bodies. The most productive way to achieve this is by using *high-resolution spectroscopy* (HRS) instruments such as CRIRES at the VLT (Käufl et al. 2004). Observing planetary light with spectral resolving powers of the order of 100,000 makes it possible to resolve individual spectral lines and molecular feature bands for simple molecules such as oxygen, water and methane in the atmospheres of these planets. Since the spectral "fingerprint" of these molecules can be well characterised via laboratory measurements, it is possible to construct high-resolution templates of their absorption and emission features. By cross-correlating observed spectra with such a template, it is possible to combine the signal from thousands of individual features in the planet spectrum and hence detect the presence of that molecular species in its atmosphere with much higher signal-to-noise than by looking for individual lines in the spectrum itself (e.g. Snellen et al. 2010; Brogi et al. 2012b; de Kok et al. 2013). By using such methods, the sub-field of astrobiology can hope to detect the spectral signatures of molecules such as oxygen and methane in exoplanetary atmospheres: since these molecules are not expected to be present in large quantities due to geological processes alone, they may constitute bio-marker signals which indicate the presence of life. In addition to the composition

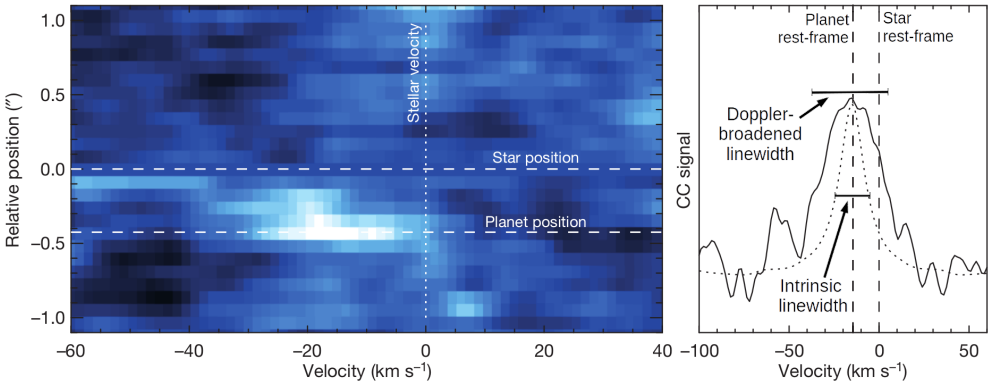


Figure 1.3: Determination of the spin rate of the directly-imaged planet β Pictoris b, using the molecular template cross-correlation method on data from the CRIRES instrument (Snellen et al. 2014). Left: a datacube of the cross-correlation signal from a water (H_2O) plus carbon monoxide (CO) spectral template, as a function of Doppler-shift velocity on the x-axis and angular separation from the star on the y-axis. The planet can be seen as a bright feature at -0.4". Right: A cut through the datacube at the planet's location, showing the average spectral line profile compared to the intrinsic linewidth expected for a non-rotating object. The additional broadening of the line reveals that the planet rotates with a day length of just eight hours, while the average velocity offset from the stellar rest frame measures the orbital velocity of the planet around its host star.

of planetary atmospheres, the identification of telltale atomic and molecular signals can also be used to provide detailed information about other planetary properties, including the atmospheric vertical temperature and pressure structure (Stevenson et al. 2014), and surface gravity (Martín & Zapatero Osorio 2003). It is even possible to detect the signatures of "disintegrating exoplanets", where material is being ejected from the atmosphere or surface of extremely close-orbiting planets due to the impact of powerful stellar winds (Brogi et al. 2012a; Rappaport et al. 2012; Ridden-Harper et al. 2016).

The molecular template cross-correlation technique can also be used to measure the planet's rotation speed and hence determine its day length, by detecting a slight broadening of the combined line profile with respect to the intrinsic width of the spectral features in question. This signal is caused by the rotating planet's limbs moving towards and away from us respectively, causing a corresponding Doppler shift in components of the molecular signal originating in these regions[5]. This measurement has now been achieved for a number of planets including the archetypal directly-imaged planet β Pictoris b (Snellen et al. 2014) shown in Fig. 1.3. This information can provide an insight into how the spin rates of planets change as they age (Schwarz 2017, PhD thesis), which is an important step towards understanding the planet formation process.

High-resolution spectroscopy observations of planetary atmospheres can be obtained in two main ways. The first is to place an optical fibre or spectrograph slit at the location of a directly-imaged planet to isolate its signal, as was done in the case of β Pictoris b. The other main method is the use of *transmission spectroscopy*, where the atmospheric signature of a transiting planet can be measured from the small fraction of starlight which passes through the planet's atmosphere during the transit. The planetary radius appears fractionally larger at wavelengths corresponding to the ab-

[5]This perhaps gives a new, and quite literal, meaning to the quote at the beginning of this introduction.

sorption features of molecular species present in the atmosphere, making the transit signal marginally deeper in these wavelengths. An extension to this is possible for hot Jupiters, where the planetary signal is often bright enough with respect to that of the star to be a detectable component of the total, making it observable even when out of transit. Monitoring the changing spectrum of this planetary component at different points along its orbit allows us to build up a *phase curve* for the planet. Since these planets are most often tidally locked (always pointing the same face towards the star), it is possible to build up a one-dimensional longitude map of the temperature and composition of the unresolved planet's upper atmospheric layers (Knutson et al. 2007; Stevenson et al. 2014). The hot Jupiter population is therefore becoming increasingly well characterised: the challenge for the immediate future is therefore to push towards achieving the same level of detailed characterisation for temperate Earth-sized planets. It is in this effort where an inter-optimised combination of high-resolution spectroscopy with high-contrast imaging techniques (Snellen et al. 2015) stands to make a major impact, by facilitating the detailed characterisation of directly-resolved planets, including those not detectable by transit or RV methods (e.g. Schwarz et al. 2016a; Hoeijmakers et al. 2018).

1.2 Planet formation

1.2.1 The planet-disk connection

Following the description of Williams & Cieza (2011), stars and planetary systems are born from the gravitational collapse and fragmentation of over-dense regions of *giant molecular clouds* (GMCs), triggered by external factors such as gravitational interactions or shock waves from nearby supernovae. Any amount of net angular momentum in the collapsing cloud causes it to spin up as it contracts, resulting in a thick disk of material surrounding a central protostar, which is not yet dense enough to initiate nuclear fusion. Over the course of typically 1 million years, the protostar contracts and accretes material from the surrounding gas- and dust-rich *proto-planetary disk* (PPD) before initiating nuclear fusion and transitioning into either a T-Tauri or Herbig Be/Ae class pre-main-sequence star, depending on its mass (Herbig 1960). This star- and planet-formation process from GMC collapse through to a mature solar system is illustrated in Fig. 1.4, along with examples of each class of object.

As the radiation output of the newly-formed central star intensifies, it begins to clear the gas and small dust grains in the proto-planetary disk from the inside-out through a combination of photo-evaporation and radiation pressure. These so-called *transition disks* are defined by the presence of a central cavity devoid of material, but often also display numerous interesting features in the remaining disk material such as narrow gaps, spirals and local over-densities, which may be a signature of ongoing planet formation.

It is during this transition process that planets are indeed expected to form out of the dust- and gas-rich material of the proto-planetary disk. There are three main competing theories as to how this occurs, and all may be responsible for forming planets under different conditions: no single theory is currently capable of explaining the full population of known exoplanets. The first is *core accretion* (Pollack et al. 1996; Hubickyj et al. 2005), whereby planets form slowly in a bottom-up manner from the collisional accretion of dust grains into rocky *planetesimals* of increasing size. In the

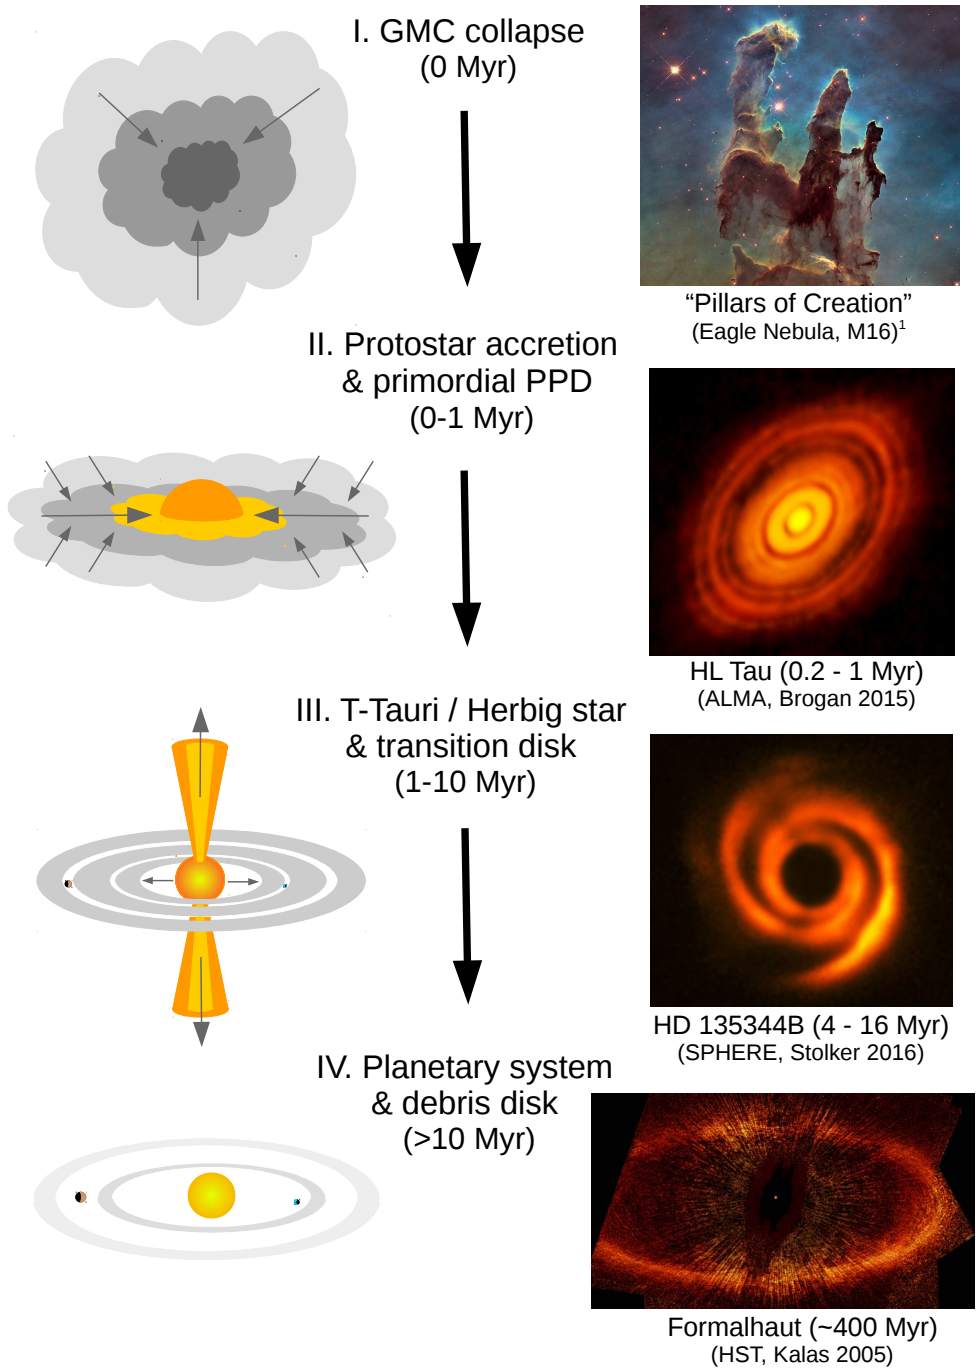


Figure 1.4: Schematic diagram of the four main phases of star and planet formation.
[1] www.nasa.gov/image-feature/the-pillars-of-creation

case of giant planets, these cores eventually grow massive enough to accrete gas from the surrounding disk, thus forming their massive atmospheres. *Gravitational instability* (Boss 2001; Mayer et al. 2002) on the other hand proposes a top-down approach whereby natural turbulence or the rapid cooling of regions in the disk can result in significant gas and dust overdensities, which accumulate matter until gravity takes over: a planet then forms rapidly by runaway accretion on timescales of just a few orbits. Since both gas and dust are accreted at the same time, gas-giant cores then form later by differentiation. In the case of very massive exoplanets or brown dwarf companions, a third method of *binary fragmentation* (Chabrier et al. 2014) is also a possibility; here the initial GMC can fragment as it collapses, forming a secondary protostellar core in the same manner as binary stars form, although this secondary object never becomes massive enough to achieve fusion.

Once the proto-planetary disk has been fully evaporated by the central star after approximately 10 million years (Myr), the formation of gaseous planets must come to an end, since there is no longer any free gas available to accrete and thereby form their atmospheres. Only the newly-formed planetary system and a residual gas-poor *debris disk* (Wyatt 2008) remain, the latter consisting of rocky and icy grains large enough such that the effect of radiation pressure is negligible, plus smaller components from the subsequent collisional fragmentation of proto-planetary bodies. This is the final product of the planet formation process, and the direct equivalent to our solar system including the asteroid and Kuiper belts, and an Oort cloud of material ejected onto extremely eccentric orbits by close gravitational encounters.

1.2.2 Observational challenges to current planet formation theory

The first major issue with the core accretion formation mechanism is that it is difficult to form metre-sized bodies by the process of collisional build-up from mm-scale dust. This is firstly due to the dynamics typically present in these disks; cm-sized bodies tend to approach each other at velocities which will more often shatter the nascent planetesimals under typical conditions, rather than having them stick together (Ros & Johansen 2013). A potentially larger issue is that cm to m-sized bodies begin to decouple from the gas component of the disk, which orbits at sub-Keplerian velocities due to the effect of radiation pressure. These bodies then spiral in towards the central star due to gas drag and disappear on timescales of less than a hundred years (Brauer et al. 2007); far too quickly to assemble planets. It has been proposed that pressure traps in the gas could be able to halt this inward drift and help accumulate planet-forming material (Meheut et al. 2012). The existence of such a pressure trap typically requires a planet to already exist in the system however, which leads to a chicken-and-egg problem. This process of bottom-up planet formation is also comparatively slow in forming gas-giant planets: building up a 10 M_{Earth} planetary core, which is thought to be needed in order to accrete a significant atmosphere, is expected to take a comparable time to the transition disk clearing timescale (Rice & Armitage 2003). This makes it especially challenging to explain the formation of extremely wide-orbit giant planets such as Uranus and Neptune, where collision rates between accreting bodies are low and hence the growth rate is especially slow. Gravitational instability on the other hand is able to form planets significantly faster, but is only thought to be feasible in extremely massive disks such as those around Herbig stars (Mayer et al.

2005): it is therefore unlikely to be responsible for giant planets around solar-type stars.

The discovery of a sizeable population of hot Jupiters was initially surprising, not just because no such object exists in our solar system but because accepted planet formation theories hold that large gas-giant planets can only form further out in the solar system. Bias-corrected statistics from the Kepler mission in fact show that only approximately 0.5% of stars host hot Jupiters, making them relatively rare, however it is still necessary to determine the mechanism by which this population came to exist. It is most likely that some form of migration, whether through slow inward migration due to interaction with the disk (Chambers 2006) or from gravitational encounters between planets (Weidenschilling & Marzari 1996) is responsible for the existence of these objects. This allows them to form in the same manner as conventional gas-giants (although this formation mechanism is itself uncertain) out beyond the snow-line of the proto-planetary disk where it is cold enough for ice species to condense onto dust grains, providing significantly more solid material for the accretion process.

At the opposite extreme, the wide-orbit giant planets to which direct imaging is most sensitive are often massive enough to lie in the grey area in terms of what may be considered a planet. The fusion of deuterium becomes possible inside the cores of objects larger than approximately 13 M_{jup}, beyond which point these bodies become increasingly similar to low-mass stars rather than planets. Studies of these brown dwarf "failed stars" is a diverse field in its own right, and there is much ongoing discussion regarding the exact mass limit cut-offs (Spiegel et al. 2011). To add to this uncertainty, accurately determining the masses of these directly-imaged exoplanets is also a significant challenge: no estimate of this quantity is provided directly by observations, which give only the relative photometry and hence luminosity of the target. Models which constrain the mass-to-light ratio of exoplanets must therefore be used to convert between the two. Poor constraints on the planet formation mechanism and hence the amount of internal heat planets start with, and how fast this is emitted as the planet ages, therefore lead to significant systematic uncertainties in mass estimates when trying to characterise these systems. The two main model classes used to estimate directly-imaged planet masses are referred to as *hot-start* (e.g. Baraffe et al. 2003) and *cold-start* models (Fortney et al. 2008) respectively, which loosely correspond to the gravitational instability and core accretion planet formation methods mentioned earlier; the slow accretion of rocky bodies allows heat to radiate away more effectively during the formation process, resulting in dimmer planets, while gravitational instability forms planets significantly faster, making them hotter for any given age. Understanding planetary formation mechanisms by observing proto-planetary disk environments will therefore ultimately allow for significantly more robust mass estimates of these planets.

1.2.3 Observing proto-planetary disks

Young T-Tauri and Herbig stars are excellent sites for studying the formation of planetary systems, since the proto-planetary disks they host are typically at some stage of the transition process to a mature system. By combining detailed single-object studies with comparative analyses of many disks of different ages, it is possible to begin constraining an overall picture of the planet-formation process.

A good rule of thumb when observing proto-planetary disks is that the observ-

ing wavelength is more or less equal to the size of dust particles that dominate the observed signal. The reasoning behind this is non-trivial, stemming from a combination of the wavelength-dependent photon scattering and emission cross-sections of dust particles and number-density arguments (Draine 2006). Near-infrared (NIR) observations with wavelengths of the order of microns therefore trace the distribution of micron-sized dust, specifically from the scattering of stellar photons by the protoplanetary disk. Due to the high opacity of the disk at these wavelengths, this signal typically corresponds to only the outermost layers of the disk. These "scattered-light" observations can therefore be used to probe the surface features of the disk (such as gaps, spirals and cavities), measure the relative scale height at different radii in the disk, and constrain the dust grain composition by making series of observations at different wavelengths.

By contrast, observations at (sub-)mm wavelengths are primarily sensitive to the thermal emission of mm-sized dust grains. These typically lie close to the mid-plane of the disk, due to dynamical processes which stratify the disk and allow larger grains to settle towards the ecliptic plane. Sub-mm observations using large interferometric arrays, such as the Atacama Large Millimeter Array (ALMA), are also able to measure the presence of certain gas species (such as carbon monoxide, CO) which produce molecular emission features in the sub-mm. These gas species are often used as tracers for more interesting but less visible molecules such as hydrogen, which makes up the majority of the gas mass of the disk (Schwarz et al. 2016b). Complementary NIR scattered-light and sub-mm emission observations of the same target (e.g. de Juan Ovelar et al. 2016; Muro-Arena et al. 2018) can therefore provide a much more complete picture of the disk in question, including constraints on the dust-to-gas ratio, vertical structure and the total disk mass. This is crucial for understanding what material is available for planet formation as a function of distance from the star, and hence what types of object are most likely to form at different locations in the disk.

1.3 Direct imaging of exoplanets

1.3.1 Challenges and rewards

While the premise of direct imaging is a straightforward one, achieving the goal of resolving Earth-like planets around even the nearest stars is anything but simple. This is primarily due to the extreme flux ratio between planetary signals and the host star, which is typically of the order of 10^{-4} to 10^{-6} at near-infrared wavelengths for young, thermally luminous giant planets. This drops to contrasts of 10^{-8} to 10^{-10} at visible wavelengths for mature Jupiter-like and Earth-like planets respectively, whose brightest signal is from reflected starlight. Coupled with the fact that these planets are separated by a fraction of an arcsecond (1/3600th of a degree) from the star in the plane of the sky, any instrument hoping to disentangle the two signals has to make use of advanced starlight rejection techniques and have exceedingly precise control over all of its constituent optical components.

A common analogy for this task is trying to detect a firefly (which emits typically 10^{-6} W of power[6]) fluttering around a lighthouse of approximately 1000 W, which is equivalent to a contrast ratio of the order of 10^{-9}. However, in this particular analogy

[6]As adopted in https://what-if.xkcd.com/151/

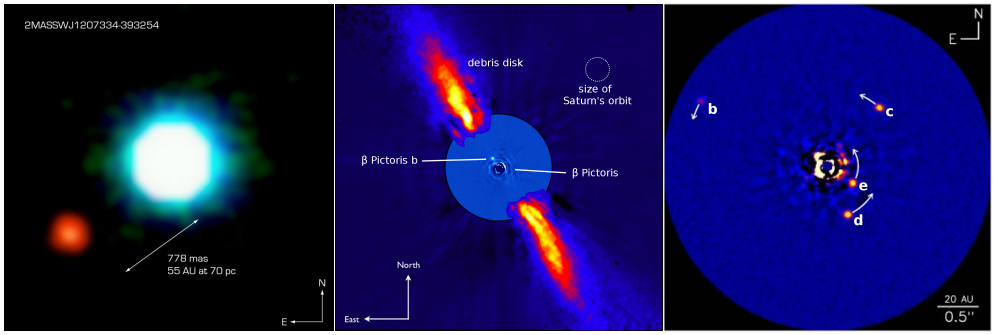


Figure 1.5: Three of the most notable directly-imaged planetary systems. Left: first directly-imaged planetary-mass companion 2M1207 b (Chauvin et al. 2004). Middle: β Pictoris b and debris disk (Lagrange et al. 2010). Right: HR 8799 four-planet system (Marois et al. 2008, 2010b).

we are looking through a telescope at Leiden Observatory at Dover lighthouse 240 km away[7], and trying to detect the firefly when it is only 12 cm away from the lamp, equivalent to an angular separation 100 milli-arcseconds (mas).

If this extreme technical challenge can be overcome however, direct imaging presents major advantages over other techniques for detecting solar system analogues. Aside from the opportunities for detailed atmospheric characterisation already outlined in Sec. 1.1.2, a major advantage is that it does not require planets to complete a significant fraction of an orbit in order to produce a statistically significant detection. This is a major limitation for transit and RV studies, and makes it extremely difficult to detect solar-system like bodies with these methods without running observing campaigns spanning years or even decades. By capturing the entire planetary system in one go, direct imaging is also a versatile tool for studying multi-planet systems and the interplay between nascent proto-planetary disks and their planetary progeny in young, still-forming systems.

The first directly-imaged planetary-mass companion was the 3-10 Jupiter-mass (M_{jup}) 2M1207 b (Chauvin et al. 2004), for which the discovery image is shown in the left-hand panel of Fig 1.5. This object has a contrast ratio with its host of only 10^{-2}: since the host is itself not a main sequence star but a brown dwarf approximately five times more massive than 2M1207 b itself, this system more closely resembles a very low-mass binary than a conventional star-planet system. Images of two other notable directly-resolved planetary systems are also shown in Fig. 1.5; β Pictoris b (Lagrange et al. 2010) has already been introduced in Sec. 1.1.2, and is shown in the middle panel to be surrounded by a significant debris disk. β Pictoris b orbits in a nearly edge-on but non-transiting configuration (Wang et al. 2016), which made its most recent closest approach in 2017. This raised hopes for the observation of circum-planetary material transiting the host star (Stuik et al. 2017), but no significant detections have yet been reported. The final panel shows the multi-planet system HR 8799 (Marois et al. 2008, 2010b), which consists of four giant planets between 4-10 M_{jup} orbiting in a nearly face-on configuration with periods of 45-460 years. During the last decade, repeated observations of these systems has allowed us to follow the orbital motion of these planets[8] and perform detailed characterisation studies (e.g. Marley et al. 2012;

[7]Ignoring for now the fact that obstruction to our line of sight by the Dutch barrier dunes, let alone the Earth's curvature, would in reality make this feat impossible.

[8]Time-lapse videos of which can be found at https://jasonwang.space/orbits.html

Bonnefoy et al. 2016; van Holstein et al. 2017).

1.3.2 Current state of the field

Figure 1.6 provides a quantitative overview of the field of direct imaging, in terms of its angular separation and planet-star contrast capabilities. Here a selection of the most notable directly-imaged planets to date are shown as orange named circles, in addition to reflected-light estimates for the solar system planets when seen in an edge-on configuration from a distance of 10 parsecs, with grey lines denoting how their detectability will vary over the course of an orbit. Overlaid on this are the performance curves of various high-contrast imaging instruments, which in general show a steady improvement between subsequent generations. NIRC2 on the W.M. Keck telescope and VLT-NACO (Lenzen et al. 2003) are examples of "zeroth-generation" instruments which began operation in the early 2000s; they were not originally designed with the direct imaging of exoplanets in mind but have since been extensively used for this task by the astronomical community. The Gemini Planet Imager (GPI, Macintosh et al. 2008) and VLT-SPHERE (Beuzit et al. 2008) came online in 2013-2014 and represent the first (and current) generation of dedicated ground-based high-contrast imagers. These instruments feature extreme adaptive optics systems along with dedicated spectroscopic and polarimetric capabilities designed to significantly enhance the ultimate contrast performance these instruments can obtain. Using advanced data reduction techniques it is now possible to achieve a *5σ contrast*[9] of approximately 10^{-6} at angular separations of 200 mas (Zurlo et al. 2016).

The minimum angular separation these instruments can achieve has however not improved greatly between these two generations. This is primarily due to the performance limits of the adaptive optics and coronagraphic systems they employ, but is also fundamentally constrained by the resolution limit of the 8 m-diameter telescopes on which they are mounted. The resolving power of any telescope is ultimately set by the diffraction limit: the *Rayleigh criterion* states that the minimum angular separation of two point sources at which they may still be resolved is $\theta_{\text{min}} = 1.22\ \lambda/D$, where λ is the observing wavelength and D is the telescope primary mirror diameter. This means that the *inner-working angle* (IWA) of an instrument is ultimately a trade-off between telescope primary mirror diameter D and the observing wavelength λ. Space missions such as the Hubble Space Telescope (HST) and the upcoming James Webb Space Telescope (JWST) are therefore limited in terms of IWA since the largest single-piece mirror which can fit into launch vehicles is currently around 2.5 m; this is a significant disadvantage compared to larger ground-based facilities when it comes to imaging close-in planets.

1.3.3 Outlook for next-generation instrumentation

The next generation of HCI instrumentation for 30-m-class extremely large telescopes (ELTs), such as EPICS (Kasper et al. 2010) for the European ELT and the Planet Formation Imager (PFI, Monnier et al. 2014) for the Thirty-Meter Telescope (TMT) as shown in Fig. 1.6, are due to see first light in the 2020s and early 2030s. This will lead to a significant improvement in both IWA and contrast capabilities, but only if

[9] A common contrast metric in high-contrast imaging, where a point source must be at least five times higher than the noise and/or variability of the residual starlight field in order to count as a detection.

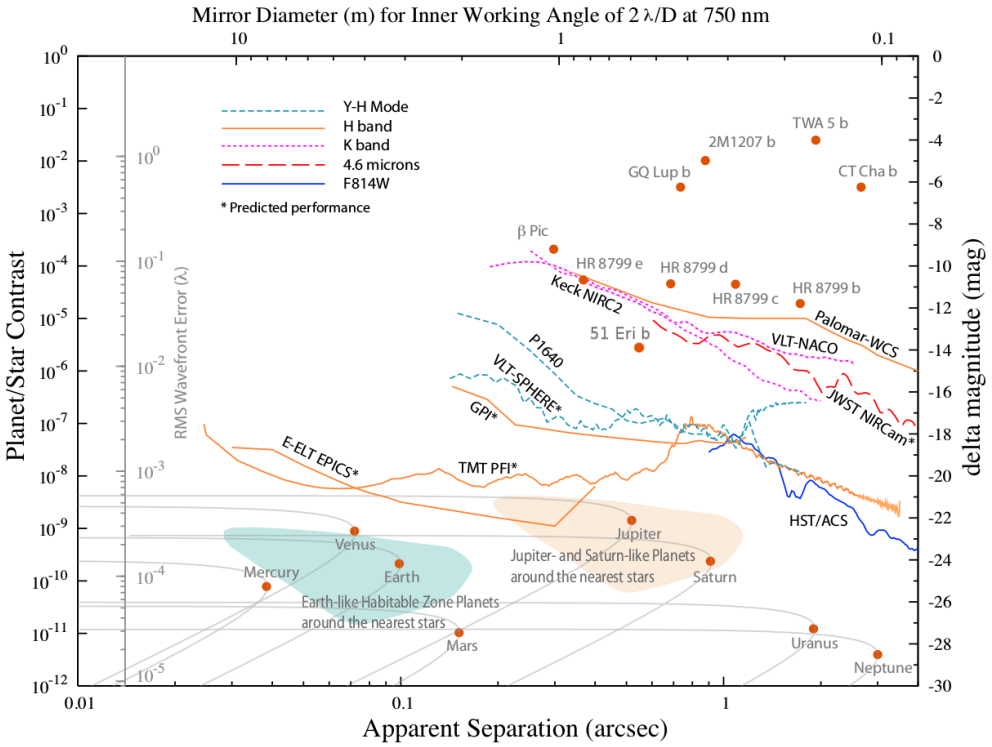


Figure 1.6: Ultimate contrast performance of current and future high-contrast imaging instruments, as a function of apparent (angular) separation from the host star. Also plotted are the solar system planets as seen in reflected light from a distance of 10 parsecs (32 light years), and a selection of the most notable giant exoplanets discovered to-date by direct imaging. Figure credit: Lawson et al. (2012); Mawet et al. (2012), updated with the subsequently discovered 51 Eridani b (Macintosh et al. 2015).

the imaging instruments concerned can actually reach the diffraction limit. In order to achieve the extreme contrast goals, efficient coronagraphs or starlight suppression techniques must first be developed; on current instruments these have functional IWAs that are much larger than the diffraction limit. The contrast of raw images taken with ground-based instruments is also significantly degraded by the distorting effects of atmospheric turbulence and optical instabilities inside the instrument. Extremely high-performance adaptive optics systems will therefore be required, even by the standards of current high-contrast imagers: the difficulty of accurate atmospheric wavefront correction only increases with increasing telescope diameter, and the accuracy of instrumental wavefront control is already the main limiting factor of current HCI instruments. These factors are discussed explicitly in Sec. 1.4.

The ultimate goal of high-contrast imaging is to detect solar system analogues, however Fig 1.7 also illustrates the interesting parameter space which contains a large number of currently known exoplanets which have not yet been directly imaged. This exoplanet population is predicted to form a broad trend with logarithmically increasing contrast requirements as a function of angular separation, since exoplanets on wider orbits will receive significantly less flux from their stars and so are correspondingly fainter when observed in reflected light at visible wavelengths. By comparing this with the predicted contrast performance of upcoming ELT instruments in Fig. 1.6,

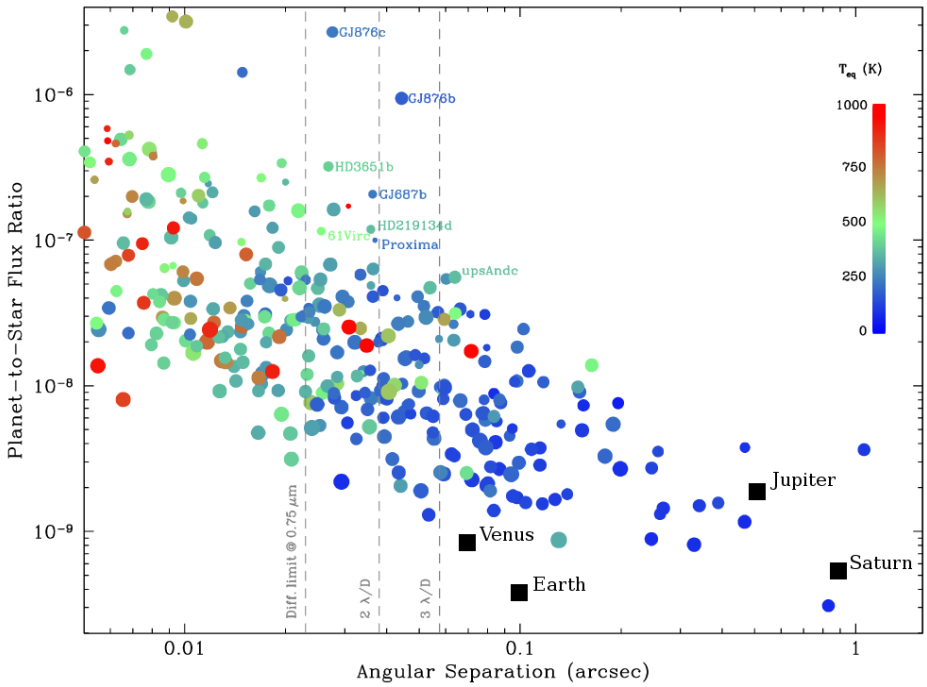


Figure 1.7: Estimated reflected-light contrast ratios of known non-directly-imaged exoplanets at an observing wavelength of 750 nm (R/I-band), calculated from known orbital parameters and estimates of likely planetary albedos, and colour coded by expected surface temperature. Vertical dashed lines indicate the diffraction-limited resolution of the 8-m VLT, plus the two and three diffraction-width cut-offs which are the main performance targets for next-generation coronagraphic imaging systems. Figure credit: Lovis et al. (2017).

we can expect that a large number of these planets should be resolvable within the next two decades. This includes the recently-discovered Proxima Centauri b (Anglada-Escudé et al. 2016), which at its point of greatest elongation sits at an angular separation of 37 mas, with an estimated contrast ratio of 10^{-7}. As the closest exoplanet to Earth, which is also potentially rocky in nature and sits on the inner edge of the nominal liquid-water habitable zone of Proxima Centauri, this object is a target of extreme interest for further characterisation studies. It is worth noting that the spatial resolution of a 30-m-class telescope operating in the near-infrared ($\lambda \approx 3\ \mu$m), such as the upcoming ELT-METIS instrument (Brandl et al. 2014), is roughly equal to that fundamentally achievable by 8-m-class instruments operating at visible wavelengths ($\lambda \approx 750$ nm), such as the current-generation VLT-SPHERE-ZIMPOL subsystem (Thalmann et al. 2008): in principle, both are therefore capable of resolving Proxima Centauri b. This goal has motivated an ongoing upgrade initiative for the SPHERE instrument, which will install cutting-edge coronagraphs, wavefront control techniques and high-resolution spectroscopy capabilities in order to bring the performance of the instrument close to the fundamental limits set by the telescope (Lovis et al. 2017).

1.4 The high-contrast imager's toolbox

The following section provides an overview of the various technologies and techniques which can be used to overcome the twin challenges of contrast and angular separation in current-generation high-contrast imaging instruments.

1.4.1 Coronagraphy

The fundamental purpose of a coronagraph is to act as an angular filter; it should suppress as much of the starlight contribution as possible at the location of the planetary signal, whilst simultaneously not degrading the strength of the planetary signal itself. The first coronagraph was invented by Bernard Lyot (Lyot 1939) in order to study the solar corona by masking out the disk of the sun, and this name has been carried over into the exoplanet field where it is instead used to observe circumstellar material by filtering out the light of the central star. There now exists a whole suite of different coronagraph designs which achieve this filtering in a variety of ways, and are comprehensively summarised by Mawet et al. (2012). These are broadly split into two main families, the difference between which is illustrated in Fig. 1.8. Focal-plane coronagraphs such as the Lyot coronagraph use an opaque mask to block out the starlight in an intermediate image plane, although they also typically use upstream and downstream pupil stops in order to optimally suppress diffraction effects (e.g. Martinez et al. 2007). Pupil-plane coronagraphs, on the other hand, exclusively use apodising optics in the telescope pupil to modify the point-spread function (PSF) of the telescope, creating regions of destructive interference, or dark zones, where the natural diffraction pattern of the telescope is significantly suppressed. This can be done either via amplitude such as with the shaped pupil coronagraph (Kasdin et al. 2004), or phase as with the Apodizing Phase Plate (APP) coronagraph (Kenworthy et al. 2010c). Interferometric nulling techniques such as Baudoz et al. (2000) can also be used for starlight suppression, while sparse aperture masking techniques (e.g. Cheetham et al. 2016) can be used to achieve angular resolutions below the conventional 1.22 λ/D resolution limit, albeit at the cost of contrast performance.

In order to make the most of the telescope's diffraction limit, it is important to develop coronagraphs which have an IWA which is as small as possible; the Apodised-Pupil Lyot Coronagraph (APLCs) is currently the workhorse coronagraph of VLT-SPHERE (Guerri et al. 2011), but is limited to inner-working angles of 3-4 λ/D by the size of the obscuring focal-plane mask. Small IWA focal-plane coronagraphs such as the vector vortex (Mawet et al. 2010) and Phase-Induced Amplitude Apodisation (PIAA, Guyon 2003) coronagraphs have been developed to push this significantly further by using phase rather than amplitude masks in the focal plane: this results in higher off-axis transmission whilst nulling out any on-axis starlight with greater precision. In laboratory environments these have been shown to reach extinction ratios of 10^{-8} to 10^{-10} at angular separations down to 1.5 λ/D, meeting the requirements for imaging solar-system analogues. The fundamental challenge with focal-plane coronagraphs however is that as the IWA becomes smaller, the coronagraph becomes increasingly sensitive to errors in instrumental alignment: any shift to the central star will cause it to move out of the nulling region and re-appear in the image, significantly degrading the contrast performance of the coronagraph. This is most often caused by high-frequency vibrations, due to cooling systems or resonances in the telescope and

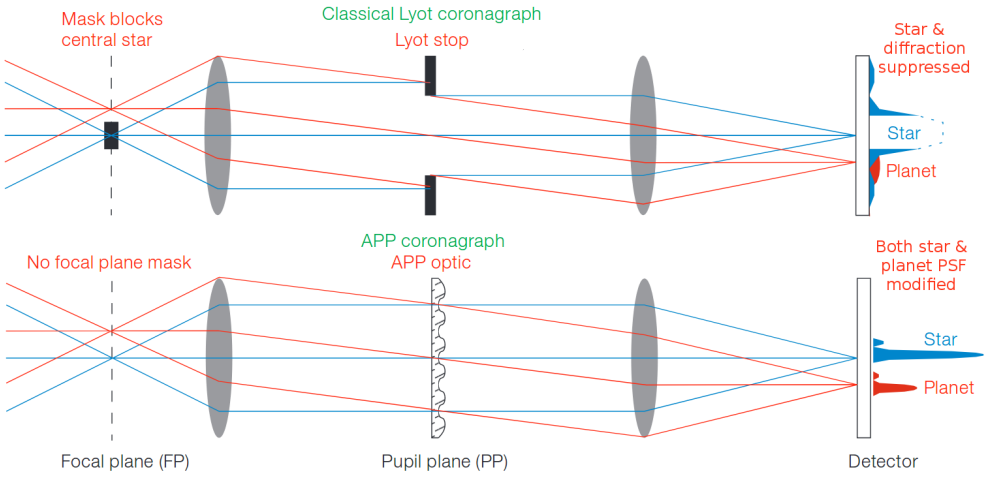


Figure 1.8: Optical layout of two example coronagraphs: Top: the traditional Lyot focal-plane coronagraph, and Bottom: the Apodizing Phase Plate (APP) pupil-plane coronagraph. Here the focal plane (FP) denotes an intermediate image plane, while the pupil plane (PP) corresponds to the re-imaged telescope aperture. In both diagrams blue rays trace the light path of on-axis starlight (which should be rejected), while the red rays illustrate how an off-axis source (i.e. the planet) is transmitted through the system. Figure credit: Kenworthy et al. (2010b).

instrument structure. Extremely fine alignment tolerances as small as a hundredth of a λ/D need to be achieved in order to reproduce laboratory performance at the telescope, meaning that these systems often do not perform at the expected level during on-sky observations.

Pupil-plane coronagraphs such as the APP sidestep this vibration-sensitivity problem, since they work by modifying the PSF of the instrument. This means that every point-source in the image is modified in the same way regardless of its exact position in the image, such that small shifts have minimal impact on the contrast in the dark zone as a whole. Such coronagraphs can also be designed to allow the high-contrast observation of multiple targets simultaneously, enabling planet searches in binary systems. It does however mean that the peak flux of the planetary signal is lowered, since this apodisation process comes at the cost of a lower *Strehl ratio*[10] for all point sources in the image. The vector Apodizing Phase Plate coronagraph (vAPP, Snik et al. 2012) is a powerful extension to the APP concept, and is discussed further in Sec. 1.5.1. vAPP optics have been demonstrated on-sky to achieve contrast levels of better than 10^{-5} over wide areas (Otten et al. 2017), and can be designed to produce contrasts of 10^{-10} in small dark zones of the order of a few square λ/D in size (Keller 2016). This optimisation parameter space is complex, with many trade-offs between contrast, IWA and image Strehl ratio (Por 2017): very different designs are therefore favoured for the detection versus the subsequent characterisation of exoplanets. The APP is straightforward to install and calibrate since it requires only a single pupil-plane optic, which has contributed to it becoming widespread in recent years. APPs are available for science observations with VLT-NACO (Quanz et al. 2010), MagAO (Otten et al. 2017) and the Large Binocular Telescope (LBT, Kenworthy et al. 2010a), with a number of new designs being produced for VLT-ERIS (Boehle et al. 2018) and Subaru-ScExAO (Lozi

[10] A common measure of image quality which is defined by the peak signal of the non-coronagraphic instrument PSF, where a perfect, diffraction-limited image has a Strehl ratio of unity.

et al. 2018), amongst others.

1.4.2 Adaptive optics

Adaptive optics (AO) is an essential component of any ground-based direct imaging instrument (Milli et al. 2016). This is because turbulent layers in the atmosphere heavily distort the phase of incoming light, which breaks up the stellar image into a widely distributed halo of speckles. This phase distortion changes on extremely short timescales on the order of milliseconds, which results in long-exposure images averaging out into a smooth, extended *seeing halo* which has a Strehl ratio of just a few percent and contains very little spatial information about the target. Such a seeing-limited image is shown in the top inset panel of Fig. 1.9.

AO systems aim to correct for this atmospheric wavefront distortion in real time, by using a *wavefront sensor* (WFS) to measure the instantaneous distorting phase aberration. The appropriate correction can then be applied via a *deformable mirror* (DM), which can change its shape in order to induce an equal and opposite phase shift into the light beam, as shown in the main diagram of Fig. 1.9. This AO correction not only increases the spatial resolution of the instrument but also dramatically enhances the signal-to-noise ratio of all point sources in the image, since the majority of the light is re-concentrated into a diffraction-limited core. Deformable mirrors can only provide an approximate correction however; the precision to which they can correct the distorting wavefront is limited by the density of actuators which control the mirror shape. This *fitting error* leads to a specific cut-off in terms of the spatial scales which it can correct, resulting in AO-corrected images displaying a typically circular *control region* inside which the PSF is nearly diffraction-limited, while regions outside see no correction whatsoever. Extreme adaptive optics (XAO) systems are now included in almost every high-contrast imaging instrument, defined by their ability to produce Strehl ratios on the order of 90% at near-infrared wavelengths thanks to their use of high-resolution deformable mirrors containing thousands of individual actuators. An example science-quality XAO-corrected image is shown in the lower inset panel of Fig. 1.9, showing the substantial difference between this and the original seeing-limited image.

Wavefront sensing is most commonly achieved by splitting off a fraction of the light from the main science beam of the instrument into a wavefront-sensing arm, which measures some form of spatial derivative of the incoming phase of the wavefront in a pupil plane of the instrument, such as the local tilt. This is mainly done because it allows the wavefront to be spatially resolved in a way that makes it more straightforward to control the individual actuators of the DM, and does not interfere with science observation. These so-called pupil-plane wavefront sensors have a major disadvantage however, in that they create regions of *non-common path* (NCP) between the wavefront sensor and the science camera; these are highlighted in red in Fig. 1.9. Any deformation to the wavefront which occurs in these regions, for example due to vibrations, internal turbulence or the slow flexing of optics due to temperature changes, will lead to an erroneous measurement of the correction needed to form a perfect image at the science focal plane. As shown in Fig. 1.10, these so-called *non-common path aberrations* (NCPAs), also known as *non-common path errors* (NCPEs), result in a field of *quasi-static speckles* (QSS) which typically evolve on timescales of seconds to hours (Fitzgerald & Graham 2006; Soummer et al. 2007a).

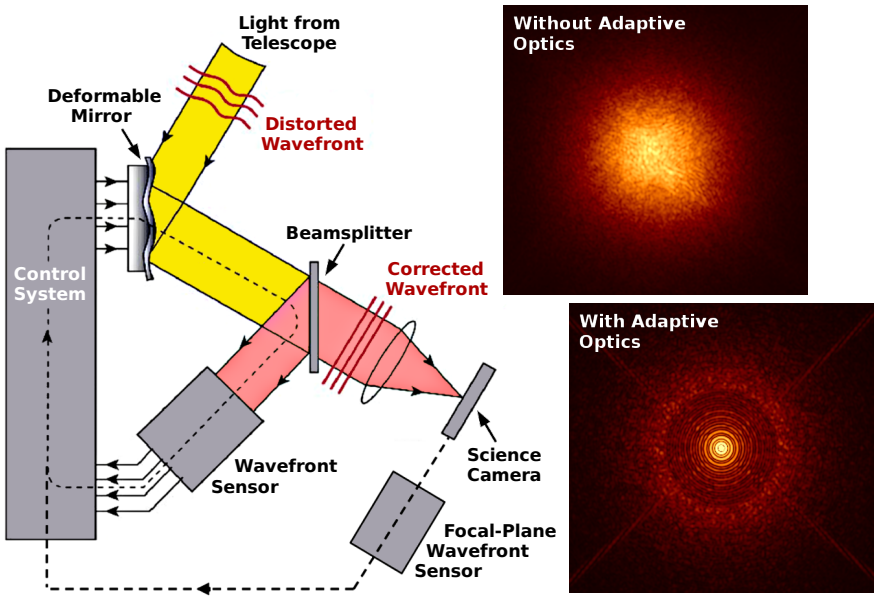


Figure 1.9: Schematic of an adaptive optics (AO) system, with image panels showing the stellar image before (top) and after (bottom) simulated AO correction. The Strehl ratio (SR) is an indicator of image quality, where a perfectly diffraction-limited image has $\mathbf{SR} = 1$. Non-common path (NCP) regions of the main pupil-plane wavefront sensor are shaded in red. A focal-plane wavefront sensor has been included, which operates on science-camera images to provide additional corrections for these NCP regions. Figure adapted from http://lyot.org.

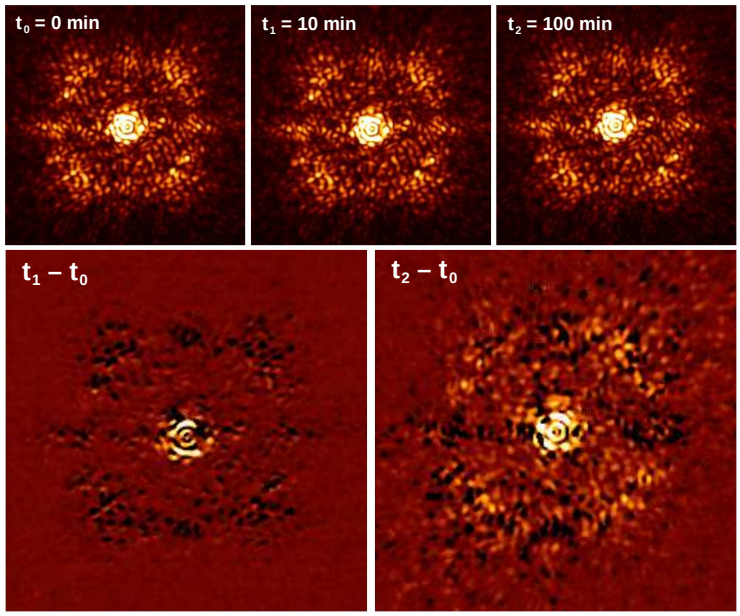


Figure 1.10: Example of a temporally evolving quasi-static speckle field in SPHERE-IRDIS (Dohlen et al. 2008) coronagraphic calibration images, due to non-common path aberrations in the instrument. Lower panels show the subtraction of very similar raw images (upper panels) taken at different times, revealing contrast-limiting residuals of the changing speckle field. Figure credit: Martinez et al. (2013).

The impact of these speckles becomes apparent when subtracting two images taken at different times during an observation window, with the raw image contrast currently limited to 10^{-3} to 10^{-4} at best (Fusco et al. 2016). A typical speckle then only has to change in brightness by 0.01% in order to produce a residual signal of the same amplitude as that expected from Proxima Centauri b.

The most straightforward data reduction technique of *reference differential imaging* (RDI) uses more or less the same approach as Fig. 1.10, by making separate observations of a planet-less star and subtracting this from the data in order to remove residual starlight. The fact that this approach is severely limited by quasi-static speckles due to the time lag between data and reference image acquisition means that NCPAs are currently the main contrast-limiting factor of high-contrast imagers. Maintaining a stable PSF is therefore in most cases more crucial to large, complex ground-based instruments than designing efficient coronagraphs capable of reaching contrasts of 10^{-10} under diffraction-limited conditions. The use of more aggressive coronagraph designs is in general only necessary once the residual QSS noise has been reduced below the photon noise limit of the static speckle background. That said, coronagraphs do help to reduce the phenomenon of *speckle pinning*, where local maxima in the diffraction-limited PSF tend to display worse speckle behaviour than regions where the starlight is suppressed (Aime & Soummer 2004).

The ideal solution to these NCPAs is to use a *focal-plane wavefront sensor*, which analyses the science image in order to determine the optimal correction without the complication of NCP regions. This is shown in Fig. 1.9 as a secondary wavefront sensing loop joined to the science camera, which provides an update to the correction of a primary pupil-plane wavefront sensor. Such a two-sensor setup is optimal because the many existing types of pupil-plane wavefront sensors have already been optimised with fast, high-resolution and high-amplitude atmospheric correction in mind, so can remove the vast majority of wavefront error from the system. Focal-plane correction schemes can then be designed with a much higher sensitivity but lower dynamic range suitable for controlling the residual QSS aberrations (which are typically less than a radian RMS[11] in total), and operate at speeds more in line with the typical coherence timescales of NCPAs and science imaging exposure times.

The development and implementation of focal-plane wavefront sensing techniques is however non-trivial; this task forms a major focus of this thesis and is therefore discussed in significantly more detail in Sec. 1.5 and Chapters 2, 3 & 4 of this thesis.

1.4.3 Image post-processing algorithms

Quasi-static speckles limit the raw contrast of coronagraphic, XAO-corrected images delivered by current HCI instruments to orders of magnitude above the planetary signatures we wish to detect. A wide variety of advanced data reduction techniques have therefore been developed in order to partially overcome this barrier, which are essential in order to reach the final 5σ contrast performance limits described in Sec. 1.3.2 and Fig. 1.6. These techniques operate by exploiting a difference in the behaviour of the planetary signal of interest compared to the halo of unwanted starlight in which it is embedded, which provides a source of diversity by which to distinguish the two components. This can then be used to construct a reference model of the stellar speckle

[11] A root-mean-square error of $1/2\pi$ wavelengths across the full telescope aperture.

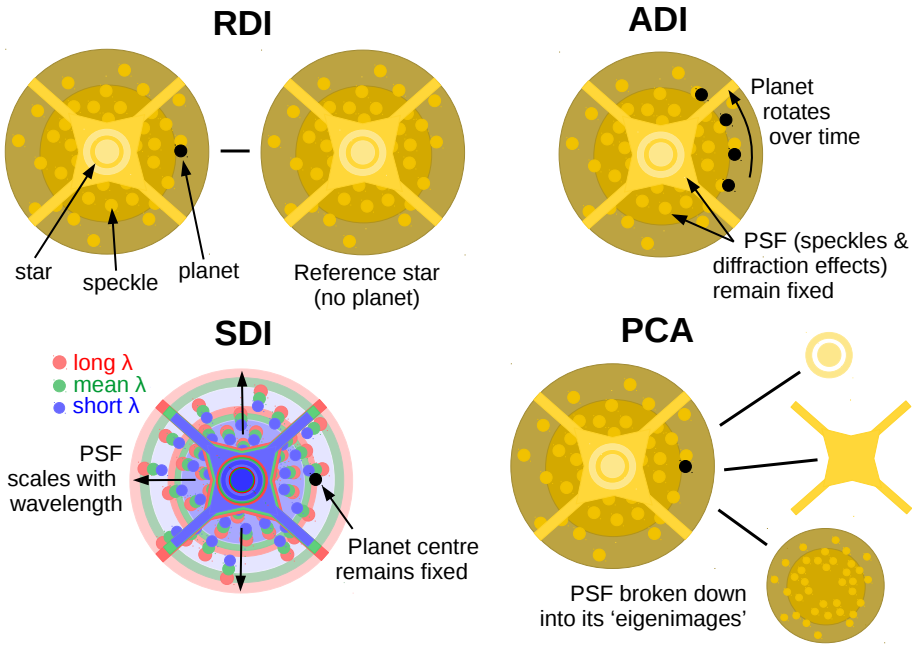


Figure 1.11: Diagrams of the various common post-processing techniques used for contrast enhancement, illustrating the form of diversity between the star and planet signal that each exploits.

field using the scientific dataset itself, thus minimising the time difference between science and reference data, over which the speckle field has a chance to evolve.

The most widely used of these techniques are illustrated in Fig. 1.11. As already mentioned, *reference differential imaging* (RDI) uses a star without any circumstellar material as the PSF reference, but is limited by the time delay (and additional NCPAs generated by changing the telescope pointing) between science and reference image acquisition. *Angular differential imaging* (ADI, Marois et al. 2006) takes advantage of the fact that altitude-azimuth mounted telescopes can be configured to allow the sky to rotate during the observation window, whilst keeping the pupil of the instrument stabilised. This means that any astronomical sources will slowly rotate around the centre of the image over time, while all instrumental speckles remain fixed in place. Many variants of ADI such as LOCI (Lafrenière et al. 2007) then use knowledge of the sky rotation rate to construct an optimal reference image in which the planet is not present, thereby minimising *self-subtraction* of the useful signal. *Principal component analysis* (PCA, Soummer et al. 2012) is an alternative method to LOCI for reducing ADI datasets, where in each frame the stellar image is decomposed into a limited set of basis functions from which a PSF reference image can be directly formed. A form of *Spectral differential imaging* (SDI, Crepp et al. 2011), on the other hand, is possible using low- and medium-spectral resolution integral field unit (IFU) imaging spectrographs, which simultaneously create a low-resolution image of the target at a number of neighbouring wavelengths. Since the diffraction limit scales with wavelength, all instrumental speckles shift radially between the bands while any true companion remains centred on the same location.

These advanced techniques are currently only capable of gaining approximately one to two orders of magnitude in contrast over a basic RDI reduction (e.g. Zurlo et al.

2016), which is primarily limited by the amount of diversity that can be introduced by these imaging modes within the typical coherence timescales of NCPAs. This is especially true at the smallest angular separations, where the amount of sky rotation and spectral dispersion is correspondingly smaller than in the outer regions of the image. ADI algorithms are also carefully optimised to detect point sources with minimal self-subtraction, but by doing so heavily remove or distort the signals of extended objects such as proto-planetary disks, and can also significantly modify the relative photometry of the final image (Milli et al. 2012).

The following subsections therefore present two major ways in which HCI instruments can be augmented, to make them sensitive to intrinsic differences between star and planet light. This is in contrast to the instrumentally-induced diversities which have been presented in this section: the following methods can be used to much more efficiently disentangle and remove unwanted stellar signal, while producing fewer data reduction artefacts.

1.4.4 Polarimetry

Light can be described as a superposition of linearly polarised components, where the electric field of the electromagnetic wave oscillates in only one direction. This fact can be exploited in order to image circumstellar material, since light is linearly polarised by scattering interactions with dust particles in the same way that the blue sky is polarised by Rayleigh scattering. By contrast, un-scattered starlight is un-polarised to an extremely high degree. At its core, the technique of *differential polarimetric imaging* (DPI, Kuhn et al. 2001) uses linear polarisation filters (analysers) to filter out all but a single polarisation state. By subtracting images of the target made with these polarisers in different configurations, an image of the total linearly polarised intensity can be built up which contains no un-scattered signal, providing an efficient means of starlight rejection.

Mathematically, the polarisation state of light can be described using the Stokes vector $\boldsymbol{S} = [I, Q, U, V]$, where I is total intensity, Q is the difference between linear polarisation along the horizontal (positive) and vertical (negative) axes, U is the equivalent linear polarisation vector rotated by 45° with respect to Q, and V is the circular polarisation. Circular polarisation is rarely seen in scattered light but is fundamental to the operation of liquid crystal-based optics, which are introduced in Sec. 1.5.1. Polarimeters are therefore designed to be sensitive to Q and U, from which the fractional degree of linear polarisation (P_I) and angle of linear polarisation (P_θ) of the astronomical signal can be calculated according to

$$P_I = \sqrt{Q^2 + U^2}/I, \tag{1.1}$$

$$P_\theta = \arctan(U/Q)/2. \tag{1.2}$$

The power of this technique when compared to intensity imaging is illustrated by Fig. 1.12, which shows a thin, high-inclination ring structure around the star HR 4796 A (Perrin et al. 2015). Unfortunately the efficiency of polarisation for light scattered by circumstellar disk material is typically observed to be of the order of only 5-10% (e.g. Krist et al. 2010; Mulders et al. 2013): DPI observations therefore throw away a significant fraction of photons from the target of interest, resulting in elevated photon noise. This is however more than outweighed by the gains in contrast that this technique can achieve over intensity imaging at small angular separations, and most

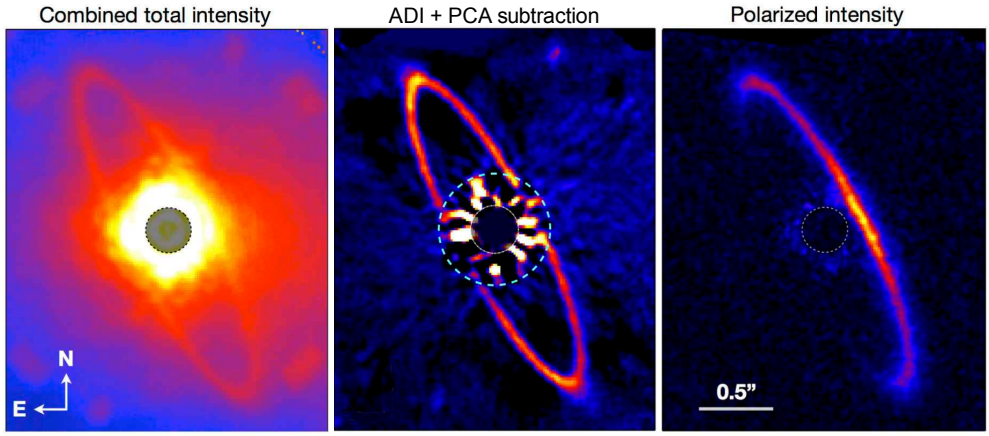


Figure 1.12: First-light images with the Gemini Planet Imager (GPI) of HR 4796 A, a massive A-type star hosting a narrow ring-like debris disk. Polarimetric imaging (right) displays significantly fewer residual stellar artefacts than intensity imaging, shown in the left panel with only minimal processing and in the middle panel following a combined ADI and PCA reduction. Only one side of the disk is strongly polarised due to the angular dependence of polarisation by scattering, revealing that the north-west side of the disk is closest to Earth. Figure credit: Perrin et al. (2015).

importantly results in significantly less distortion of extended objects than equivalent processing with ADI-like techniques.

Self-luminous planets are expected to have only a very small net polarisation signal at NIR infrared wavelengths (Stolker et al. 2017), although this rises significantly for mature planets observed in reflected light (Stam et al. 2004). The degree of linear polarisation of the HR 8799 planets has recently been constrained to less than 1% in the H-band, or a 1σ contrast of almost 10^{-7} (van Holstein et al. 2017), and while the first detection of a strongly polarised, spatially-resolved companion was recently made around CS Cha (Ginski et al. 2018), this is expected to be due to the existence of a circum-planetary disk rather than the companion itself. This technique has nonetheless had great success in imaging proto-planetary disks in scattered light; the SPHERE instrument is extremely productive for this purpose at both visible (e.g. Roelfsema et al. 2016) and near-infrared wavelengths (e.g. Benisty et al. 2015; de Boer et al. 2016; van Boekel et al. 2017). As discussed in Sec. 1.2.3, these scattered-light observations allow the determination of a number of key properties of the surface layers of these disks, including the potential identification of ongoing planet formation (Keppler et al. 2018). The interpretation of polarimetric images is often non-trivial however: reflective surfaces in the instrument often provide a strong source of artificial linear polarisation which hampers the exact determination of P_I and must first be calibrated out. Self-shadowing effects can often make it challenging to determine whether observed features in the disk are due to a dearth of scattering material or to geometric effects, while the strong angular dependence of polarisation by scattering (Pinte et al. 2009) induces a further modulation to the signal in highly inclined disks such as HR 4796 A, which does not reflect the underlying density of material.

1.4.5 Spectroscopy

As already introduced in Sec. 1.1.2, the combination of high-contrast imaging with high-resolution spectroscopy (HRS) is an extremely powerful characterisation approach, since planetary spectra can provide a wealth of information when sampled at a sufficiently high resolution. The significant diversity between planetary and stellar spectra at these resolutions also facilitates a variety of contrast-enhancing techniques.

The simplest method of using spectral information as a contrast-enhancing technique is *simultaneous differential imaging* (SDI), which originally referred to the approach of imaging the target in parallel through two narrowband filters of very slightly different wavelengths (Smith 1987; Biller et al. 2006). These filters are chosen such that one lies on top of a strong emission or absorption feature of the planet which is not present in the star, while the other lies just outside this feature. By subtracting the two images, only sources containing the chosen spectral feature remain and the stellar speckle field should be efficiently removed. This technique was employed in a number of exoplanet surveys (e.g. Biller et al. 2007; Nielsen et al. 2008) which focussed on searching for a 1.62 μm methane absorption feature, since it was predicted from observations of free-floating field brown dwarfs that exoplanet atmospheres should contain a significant methane component. Unfortunately, the dearth of new detections reported by these surveys implies that this prediction of exoplanet and brown-dwarf spectral similarity is not necessarily true.

1.4.5.1 Future contrast-enhancing applications of HRS

In addition to its value as a characterisation tool, the HRS spectral template cross-correlation method can itself be used as a contrast-enhancing technique, capable of gaining an estimated additional factor of 10^{-4} over raw coronagraphic images (Snellen et al. 2015). This is possible since both the differences in spectral features between the star and planet light and the overall Doppler shift from the orbital motion of the planet with respect to the star provide sources of diversity which can be exploited during data reduction. The process of coupling HCI and HRS instrumentation also holds the potential to provide additional starlight nulling. This coupling is typically achieved by feeding light from the HCI focal plane into an optical fibre, through which it can be passed to the spectrograph. While wide-diameter (multi-mode) optical fibres are often used in order to transmit as much light as possible, the newly developed SCAR coronagraph concept (Por & Haffert 2018; Haffert et al. 2018a) shows that single-mode fibres can be used to filter out non-point-source-like light to increase the contrast performance of a fibre-fed integral field unit (IFU).

In order to meet the performance requirements of the next generation of ELT-class exoplanet characterisation instruments, it is necessary to combine as many of the the contrast-enhancing techniques described in this section as is possible. In this way, a fully system-engineered instrument which optimally combines high-contrast imaging with high resolution spectroscopy, and ideally also with polarimetric and IFU capabilities, is the most promising approach to optimise the ultimate contrast performance of instruments such as ELT-EPICS (Kasper et al. 2010).

1.5 This thesis: Focal-plane wavefront sensing techniques

As discussed in Sec. 1.4.2, focal-plane wavefront sensing is the most direct method for eliminating the quasi-static speckles which currently limit high-contrast imaging instruments: it is however a fundamentally non-trivial task when compared to pupil-plane wavefront sensing. When combined with the need to avoid interfering with the science observations, focal-plane sensing with the main science camera constitutes an unattractive solution for most AO applications and is therefore comparatively underdeveloped at the present time.

The two main reasons why most wavefront sensors operate in the pupil plane are that it allows the wavefront to be spatially resolved, and that there are numerous ways to straightforwardly measure some derivative of the phase (e.g. local gradient) for each resolution element. This is illustrated in the top row of Fig. 1.13 for the example of the ubiquitous Shack-Hartmann pupil-plane wavefront sensor (SH-WFS). The SH-WFS uses a lenslet array to create a grid of subsidiary PSFs on the sensor camera, each of which samples only a small region of the pupil. The location of each PSF is linearly related to the local gradient (tip-tilt phase error) of the incoming wavefront in each resolution element, allowing for a direct reconstruction of the full phase map by assuming continuity of the wavefront.

Conversely, the science focal-plane image $p(x, y)$ is formed by the absolute value

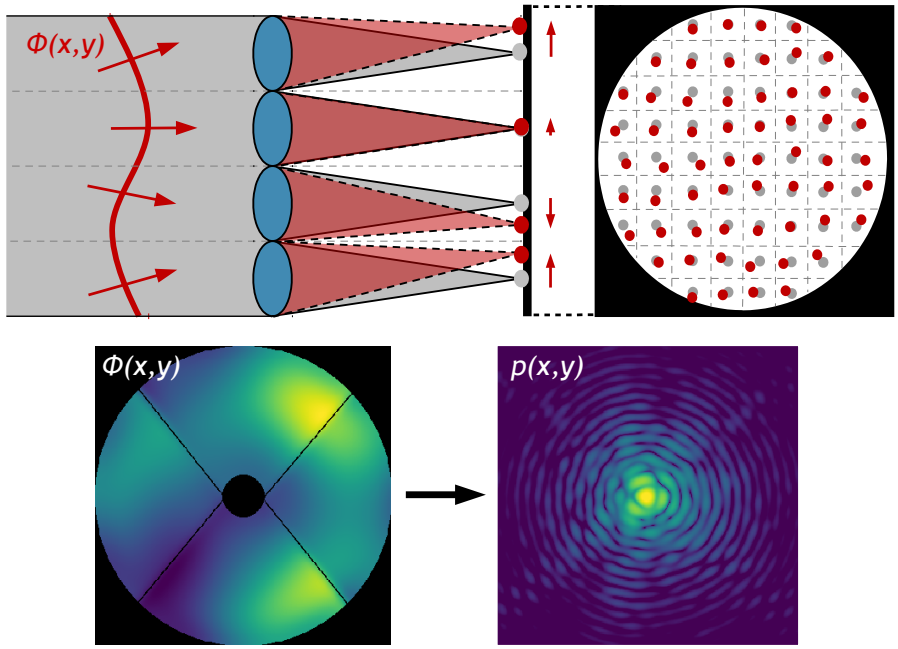


Figure 1.13: Illustration of the information content available to pupil-plane and focal-plane wavefront sensors. Top: Diagram of the Shack-Hartmann pupil-plane wavefront sensor, showing how the position of each subsidiary PSF is sensitive to the local gradient of the aberrating wavefront $\phi(x, y)$ (red spots, with respect to the grey nominal locations). Bottom: An example wavefront phase map $\phi(x, y)$ and the corresponding aberrated focal-plane PSF $p(x, y)$, which by comparison contains no easily retrievable phase information.

squared of the complex Fourier transform of the complex electric field in the pupil plane, expressed as

$$p(x,y) = \left|\mathcal{F}\left[A(x,y)e^{i\phi(x,y)}\right]\right|^2, \tag{1.3}$$

where $A(x,y)$ is the telescope aperture function and $\phi(x,y)$ is the wavefront phase to be sensed. This is illustrated in the bottom row of Fig. 1.13: not only does this mean that the wavefront is no longer spatially resolved, but since imaging cameras are only capable of making intensity measurements there is also a fundamental loss of information about the complex focal-plane electric field, $\mathcal{F}\left[A(x,y)e^{i\phi(x,y)}\right]$. The PSF therefore by default contains no easily retrievable phase information, and in general it is not theoretically possible to reconstruct the full wavefront phase from a single PSF intensity image.

A variety of focal-plane wavefront sensing techniques are now under development which can overcome these challenges, many of which are summarised by Jovanovic et al. (2018). Two specific approaches which are directly relevant to this thesis are briefly outlined in the following subsections.

1.5.1 Holographic wavefront sensing

1.5.1.1 Principles of holography

Since the science image does not contain easily retrievable phase information by default, holographic optics can effectively be used to engineer the PSF of the instrument in such a way as to make direct phase retrieval possible. This works by splitting off small amounts of light from the main science PSF to different locations in the focal plane, which can each be modified with independent wavefront aberrations whilst maintaining a dominant, un-affected central image for science use.

This approach is heavily based on the Holographic Modal Wavefront Sensor (HMWS), originally developed for confocal microscopy (Neil et al. 2000; Booth 2003). The principle behind this is illustrated in Fig. 1.14, where a phase-only holographic optic is used to create secondary copies of the science PSF at different locations in the focal plane. Since these secondary copies can be biased with independent wavefront aberrations, it is possible to encode phase information about a finite set of important aberration modes into the focal plane, in a manner which is easily measurable. The coronagraphic Modal Wavefront Sensor (cMWS) uses this principle in conjunction with the APP coronagraph to provide simultaneous high-contrast imaging and focal-plane wavefront sensing capabilities with the same detector: the development and testing of this technique is the subject of Chapter 2.

These *computer-generated holograms* may be designed via three main encoding methods, which form a direct analogy with the various common types of *diffraction grating* (Brown & Lohmann 1969):

$$H_{cosine}(x,y) = \sum_k a_k \cos\left(O_k(x,y) + R_k(x,y)\right), \tag{1.4}$$

$$H_{binary}(x,y) = \arg\left(H_{cosine}(x,y)\right), \text{and} \tag{1.5}$$

$$H_{blazed}(x,y) = \arg\left(\sum_k a_k e^{i(O_k(x,y)+R_k(x,y))}\right). \tag{1.6}$$

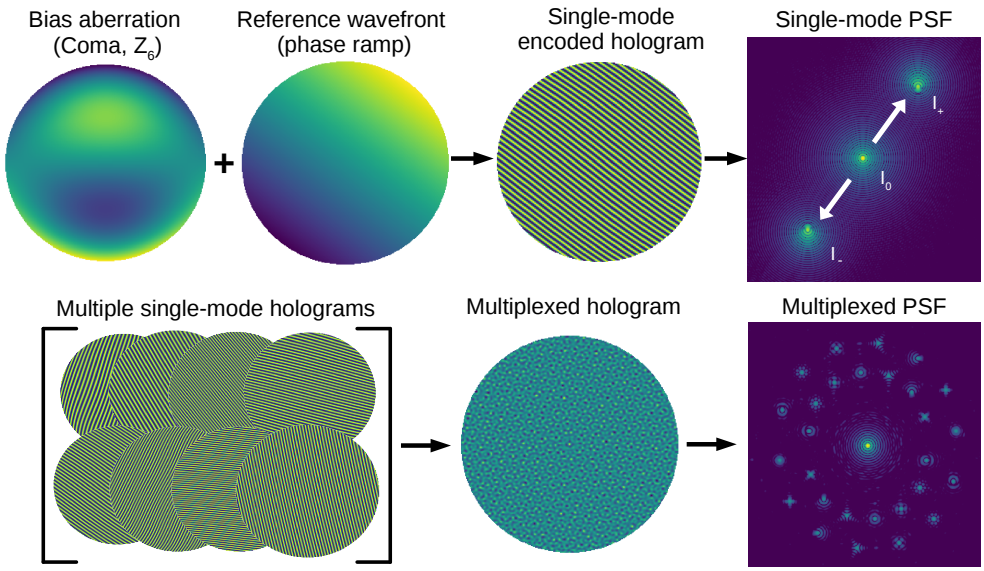


Figure 1.14: Illustration of the design process of a Holographic Modal Wavefront Sensor (HMWS). Top row: A single phase aberration (in this case coma) forms the bias wavefront (O_k), which corresponds to the wavefront mode to be sensed, while the amplitude and direction of the tilted reference wavefront (R_k) defines the location of the biased PSF copies in the focal plane. This is encoded into a phase-only hologram according to one of Eqs. 1.4 - 1.6 (in this case Eq. 1.4), which functions as a form of spatially modulated diffraction grating. The resulting instrument PSF then contains secondary PSF copies in the ±1 diffraction orders of the grating, biased with opposite signs of the wavefront aberration. Bottom row: The summation of multiple single-mode holograms (H_k) allows for many pairs of PSF copies to be created simultaneously, with independent positions, wavefront biases, and relative brightnesses with respect to the central science image.

Here $O_k(x, y)$ is the object (bias) wavefront and $R_k(x, y)$ is the reference (tilted) wavefront as illustrated in Fig. 1.14: this terminology is conventional for traditional forms of optically-exposed hologram (Dong et al. 2012). arg() denotes the complex argument function: multiplexed holograms can be created by summing the individual single-mode components before taking the argument, with differing amplitude a_k controlling the relative brightness of each resulting copy.

These different holographic encodings have various strengths and weaknesses in terms of performance: the cosine hologram creates significantly less speckle noise in the science image plane than either binary or blazed holograms, since it does not suffer from phase information loss (due to taking the complex argument in the latter two cases), however it is not capable of 100% diffraction efficiency and is limited in the number of modes that can be multiplexed (Changhai et al. 2011). This trade-off is explicitly discussed for the case of binary and cosine holograms in Wilby et al. (2016a). The more recent development of optimised cMWS optics based on multiplexed blazed holograms is published in Haffert et al. (2018b), as part of further on-sky cMWS demonstration work with the *Leiden EXoplanet Instrument* (LEXI). These blazed grating holograms provide an excellent trade-off between the binary and cosine approaches, outperforming both in the case of highly multiplexed holograms.

1.5.1.2 The vector-APP coronagraph

The vector-APP (vAPP) coronagraph was developed in Snik et al. (2012) and Otten et al. (2014a,b) as a work-around for some of the limitations of the classical APP coronagraph. The most notable of these is that only one side of the PSF could be made dark at any one time, meaning that for survey work two sets of observations would need to be made for each target, rotating the optic in between. The vAPP is created using a modified polarisation grating to create two PSFs simultaneously, each containing an APP dark zone on the opposite side of the star, as illustrated by Fig. 1.15. This is fundamentally the same principle as the blazed HMWS, simply using a single wavefront aberration which is customised to function as an APP coronagraph rather than as a wavefront sensor: this makes it perfect for incorporation into the cMWS.

Since this work it has now become possible to design annular (360 degree) dark-hole-generating APPs with useful throughput. However, the 180 degree vAPP concept remains useful: it can achieve inner-working angles as small as 1.6 λ/D whilst retaining a useful level of throughput and contrast performance, whereas the design parameter space of 360 degree APPs is significantly more restricted (Por 2017).

1.5.1.3 Liquid crystal manufacturing technologies

The development of the vAPP coronagraph is crucial to the work in this thesis, since many advances in its manufacture have been made in recent years. APP optics were originally produced as diamond-turned glass plates, with the thickness at each location in the pupil inducing the desired phase delay into the wavefront (Codona et al. 2006; Kenworthy et al. 2007): unfortunately this manufacturing technique has relatively limited spatial resolution and is not capable of producing sharp jumps in phase. This rules out a large fraction of the design parameter space of both the vAPP and the cMWS, which typically produce extreme phase patterns which are dominated by discontinuities and high-spatial frequency phase content.

Direct-write liquid crystal technologies (Miskiewicz & Escuti 2014), on the other hand, can be used to create phase optics with spatial resolutions of the order of 1 μm and with artefact-free phase wrapping, allowing almost any phase design to be accurately manufactured. This is achieved by using a scanning UV laser to expose a photo-alignment layer, which in turn controls the orientation of subsequently-deposited liquid crystal layers at each point on the substrate. The resulting optic is effectively a *spatially-modulated half-wave retarder*, which acts on incident circular polarisation, flipping its handedness while inducing a phase shift which is proportional to the angle between the crystal fast axis and the incident electric field. The resulting phase pattern is also achromatic since the phase modulation effect of liquid crystals is produced via so-called *geometric phase*, which is independent of wavelength (Komanduri et al. 2013). For a significantly more detailed description of this manufacturing technique and its implications, see the PhD thesis of Gilles Otten (Otten 2016).

1.5.1.4 Future applications for holographic optics

The success of liquid crystal phase optic manufacturing techniques has led to a recent explosion in the number of vAPP coronagraphs being implemented in high-contrast imaging instruments around the world: many of these are listed in Sec 1.4.1. This also applies to the cMWS, with optics now being implemented in LEXI at the William

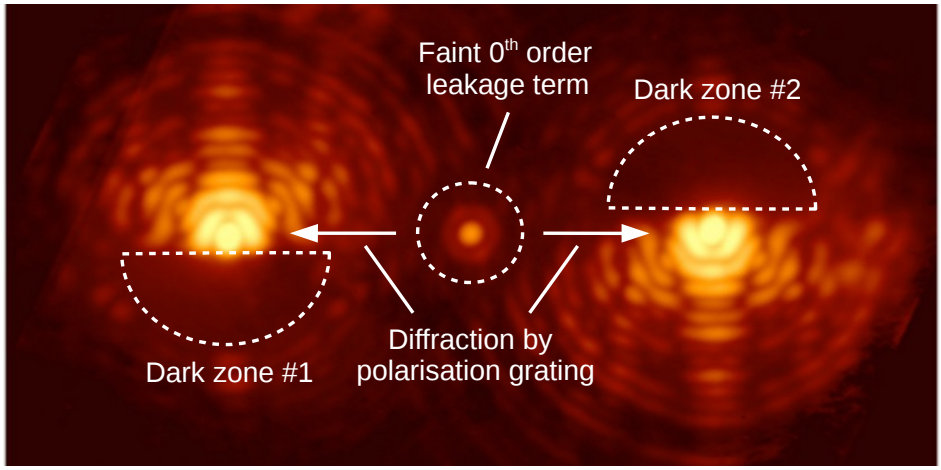


Figure 1.15: On-sky image of a single point source using the vector-APP coronagraph installed in the MagAO instrument at the Magellan Clay Telescope, Chile (Morzinski et al. 2014). Each of the two science PSF copies consists of light of a single circular polarisation, artificially split off from the stellar position by means of a diffraction grating. A small fraction of the stellar light passes through the optic unaltered, forming the zeroth-order leakage term at the original stellar location. Image credit: Gilles Otten.

Herschel Telescope (Haffert et al. 2018a), the HiCIBaS high-altitude balloon mission (Côté et al. 2018), the MagAO-X upgrade for the Magellan Telescope (Miller et al. 2018) and SCExAO at the Subaru Telescope (Doelman et al. 2017).

The development of holographic optics for the cMWS has also already led to a number of spin-off applications. There exists a wide variety of other focal-plane wavefront sensing strategies which, like the original (non-holographic) concept of modal wavefront sensing, are designed to be performed in multiple sequential steps which each degrade the science image. The main potential of the multiplexed holographic approach is that it can transform almost any such temporally-modulated technique into a simultaneous, spatially-modulated equivalent, thereby enabling it to operate in real time and in parallel with science observations. A notable example of this is the electric field conjugation technique (EFC, Give'on et al. 2007) which, like the cMWS, uses phase probes to reconstruct the wavefront. Instead of using a set of low-order modes however, EFC simultaneously adds many high-frequency, low-amplitude aberrations into the dark zone of an APP coronagraph. By making four or more holographic copies of the coronagraphic PSF (Por & Keller 2016) to which these EFC phase probes have been added, the full aberrated electric field and hence the necessary wavefront correction can be retrieved without degrading the science image. An additional high-contrast imaging concept which has been enabled by the development of multiplexed blazed grating holograms is the holographic aperture masking (HAM) technique (Doelman et al. 2018). This is an extension to sparse aperture masking (SAM, Tuthill et al. 2010), which effectively divides up the telescope aperture into multiple independent SAMs which each produce interference fringes on different parts of the image, thereby maximising useful throughput of the system. SAM exploits the fact that interferometric techniques can beat the conventional diffraction limit and so gain significantly in effective inner-working angle performance, while HAM can also be engineered to provide simultaneous coverage of multiple position angles around the star.

In summary, the power of holographic optics to grant unprecedented control over

the instrument PSF structure and hence the information content delivered in science images is now opening up a large window of opportunity for new high-contrast imaging techniques, both for wavefront control and coronagraphy applications.

1.5.2 Phase diversity wavefront sensing

An alternative class of focal-plane wavefront sensing techniques focus on overcoming the over-arching PSF phase degeneracy issue via the use of *phase diversity*. The core principle behind this approach is that, while full phase reconstruction is in general not possible using a single science image, it is still possible to probe the underlying complex focal-plane electric field by deliberately modifying the total phase aberration content of the system in a controlled manner. When using two images which differ only by a known, deliberately injected phase aberration, it becomes possible to reconstruct the full electric field and hence any unknown component of the aberrating wavefront. Fig. 1.16 illustrates how this may be achieved in practice by defocussing the science camera; a commonly-used source of phase diversity since it is typically reproducible and easy to implement.

There are many possible reconstruction algorithms which make use of such phase diversity information (Carrano et al. 1998; Gonsalves 2001; Lamb et al. 2016), some of which are also capable of working with coronagraphic images (Paul et al. 2014b). However, if this compensation scheme is to be operated in real time (as is required for accurate NCPA correction) it is necessary to interleave defocussed images into the acquisition process, thereby degrading the imaging performance of the instrument for 50% of science frames. In most cases this means that the frequency of wave-

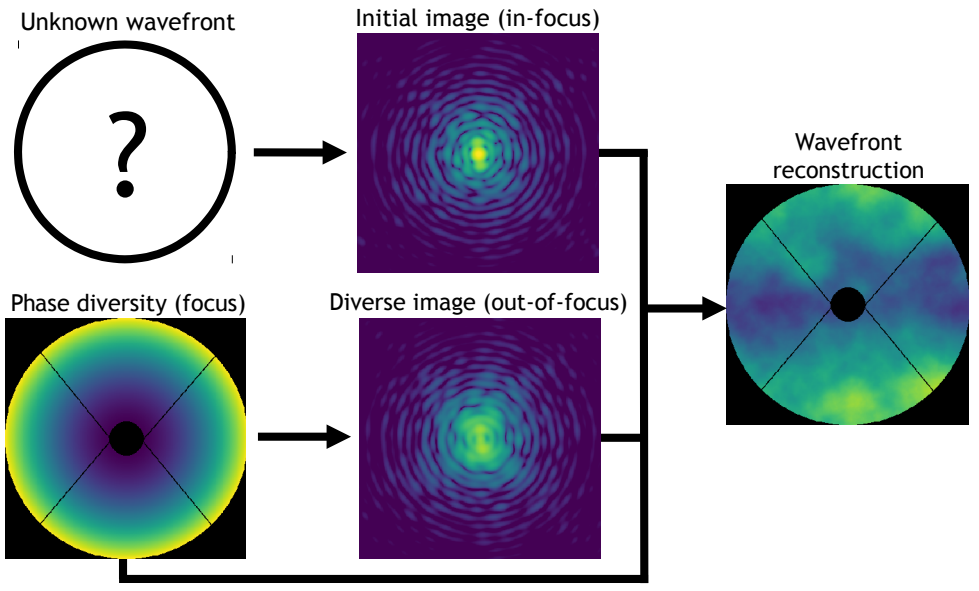


Figure 1.16: Illustration of the "classical" phase diversity wavefront reconstruction technique. A controlled phase aberration, in most cases strong defocus, is added to the system in order to modify the complex electric field in the focal plane in a manner which breaks the phase degeneracy of the PSF. By using the in- and out-of-focus intensity images and the known phase aberration which separates them, it is possible to reconstruct the full electric field in the focal plane and hence the unknown aberrating wavefront.

front correction updates will instead be significantly lowered, in order to maximise the yield of science-quality data: quasi-static wavefront aberrations will then not be fully addressed, allowing contrast-limiting speckle residuals to creep back into the final dataset.

One way around this limitation is to utilise the phase diversity information which is naturally present in any system with active wavefront control: after its first iteration, a focal-plane correction algorithm can in principle use its own previous correction update along with the current and previous image as the necessary input for phase diversity reconstruction, by direct analogy with the original approach in Fig. 1.16. By avoiding the need for destructive phase diversity probes this ensures that the image quality always improves from one iteration to the next, making it possible to operate at appropriately fast correction cadences (typically 10 Hz for NCPA control) without impacting science observations. An example of such a real-time phase diversity routine is the so-called "Fast & Furious" (F&F) sequential phase diversity algorithm, originally developed by Keller et al. (2012) and Korkiakoski et al. (2014). F&F is proposed in Chapters 3 & 4 as a software-only phase reconstruction solution for wavefront control issues seen in the SPHERE current-generation high-contrast imager at the VLT, with wider applications to other current and future instruments.

1.6 Thesis outline

The goal of this thesis is twofold: to address some of the outstanding issues limiting the performance of existing high-contrast imagers, and to develop new technologies which will ultimately facilitate the characterisation of Earth-like exoplanets with ELT-class telescopes. To this end, Chapters 2, 3 and 4 of this work focus on the wavefront control aspects of high-contrast imaging, since this is where the biggest performance gains currently stand to be made. Chapter 5 also illustrates how detailed instrumental characterisation work can be used to optimise the data reduction procedure, and ultimately improve science yields at the smallest angular separations.

Chapter 2: Introducing the coronagraphic Modal Wavefront Sensor

As outlined in Sec. 1.5.1, holographic focal-plane wavefront sensing techniques are a promising approach for the real-time correction of non-common path aberrations (NCPAs), which limit the contrast performance of current high-contrast imaging instruments. This chapter presents the theory, laboratory implementation and first on-sky validation of the *coronagraphic Modal Wavefront Sensor* (cMWS), which is capable of simultaneous wavefront retrieval and coronagraphy using the science focal plane. We show that this sensor is capable of the real-time on-sky recovery of known aberrations for low-order wavefront modes at a cadence of 50 Hz, for a large observing bandwidth of 50 % in the R-band (600 nm). This is important since many focal-plane wavefront retrieval algorithms rely on a quasi-monochromatic image, making the cMWS a versatile focal-plane sensor which is capable of controlling those aberrations which contribute most to the wavefront error budget of NCPAs. The cMWS has since been further validated as part of the *Leiden EXoplanet Instrument* (LEXI, Haffert et al. 2016): more recent on-sky results including on-sky closed-loop operation can be found in Wilby et al. (2016a) and Haffert et al. (2018a).

Chapters 3 & 4: Tackling the low-wind effect at VLT-SPHERE

NCPAs are not always the limiting factor in current high-contrast instruments. In these two chapters a potential real-time wavefront control solution is developed and tested for the so-called *low-wind effect* (LWE, Sauvage et al. 2016), which is seen to consistently degrade the contrast performance of the SPHERE current-generation high-contrast imager under otherwise optimal seeing conditions. Chapter 3 presents closed-loop simulation results for the *"Fast & Furious"* (F&F, Korkiakoski et al. 2014) phase diversity algorithm under SPHERE-like observing conditions, in order to assess the suitability of the control scheme as an immediately implementable, software-only solution for the LWE. Chapter 4 then reports on the subsequent laboratory testing of F&F at the MITHIC high-contrast testbench (Vigan et al. 2016b), located at the Laboratoire d'Astrophysique de Marseille, in order to verify the results of Chapter 3 and compare with other focal-plane wavefront control methods. F&F was found in both cases to be capable of robust real-time wavefront control under LWE-like conditions without degrading the science image, making it a promising algorithm for implementation both on SPHERE and other instruments suffering from similar wavefront control issues.

Chapter 5: Calibrating the SPHERE-IRDIS APLC coronagraph

Optimal reduction techniques are equally crucial when it comes to making the most of the data produced by current high-contrast imaging facilities. This in turn requires the thorough calibration of all instrumentation behind the observations, the results of which are not always easily available to observers. This chapter presents a characterisation effort of the apodised Lyot coronagraph system of the SPHERE-IRDIS near-infrared imager (Dohlen et al. 2008), in order to develop a suitable calibration algorithm for coronagraphic, polarimetric imaging data. This is important since the innermost regions of circumstellar disk observations, which are typically the most valuable to study, are often dominated by artefacts of the imaging system. An improved calibration routine which can correctly restore the relative photometry of these innermost regions is therefore capable of significantly improving the effective inner-working angle of the instrument. This approach is particularly crucial for the identification of central cavities in transitional proto-planetary disks; one of the main science yields of the SPHERE instrument. Calibration observations were made using the minor planet Ceres in order to accurately determine the extinction profile of the coronagraphic system. This was combined with extensive optical modelling in order to fully understand the observed signal, before re-reducing polarimetric observations of TW Hydrae (van Boekel et al. 2017) in order to validate the accuracy of the resulting calibration routine.

Chapter 2

The coronagraphic Modal Wavefront Sensor
A hybrid focal-plane sensor for the high-contrast imaging of circumstellar environments

Michael J. Wilby, Christoph U. Keller, Frans Snik, Visa Korkiakoski, and Alexander G.M. Pietrow
A&A 597, A112 (2017)

Abstract

The raw coronagraphic performance of current high-contrast imaging instruments is limited by the presence of a quasi-static speckle (QSS) background, resulting from instrumental non-common path errors (NCPEs). Rapid development of efficient speckle subtraction techniques in data reduction has enabled final contrasts of up to 10^{-6} to be obtained, however it remains preferable to eliminate the underlying NCPEs at the source. In this work we introduce the coronagraphic Modal Wavefront Sensor (cMWS), a new wavefront sensor suitable for real-time NCPE correction. This combines the Apodizing Phase Plate (APP) coronagraph with a holographic modal wavefront sensor to provide simultaneous coronagraphic imaging and focal-plane wavefront sensing with the science point-spread function. We first characterise the baseline performance of the cMWS via idealised closed-loop simulations, showing that the sensor is able to successfully recover diffraction-limited coronagraph performance over an effective dynamic range of ± 2.5 radians root-mean-square (RMS) wavefront error within 2-10 iterations, with performance independent of the specific choice of mode basis. We then present the results of initial on-sky testing at the William Herschel Telescope, which demonstrate that the sensor is capable of NCPE sensing under realistic seeing conditions via the recovery of known static aberrations to an accuracy of 10 nm (0.1 radians) RMS error in the presence of a dominant atmospheric speckle foreground. We also find that the sensor is capable of real-time measurement of broadband atmospheric wavefront variance (50% bandwidth, 158 nm RMS wavefront error) at a cadence of 50 Hz over an uncorrected telescope sub-aperture. When combined with a suitable closed-loop adaptive optics system, the cMWS holds the potential to deliver an improvement of up to two orders of magnitude over the uncorrected QSS floor. Such a sensor would be eminently suitable for the direct imaging and spectroscopy of exoplanets with both existing and future instruments, including EPICS and METIS for the E-ELT.

2.1 Introduction

2.1.1 Scientific motivation

Since the first direct image of a planetary mass companion around a nearby star was obtained in 2004 (Chauvin et al. 2004), the field of high-contrast imaging has undergone rapid development with the advent of advanced coronagraphic techniques (Mawet et al. 2012) and eXtreme Adaptive Optics (XAO) systems (e.g. Sauvage et al. 2010). This progress continues with the recent first light science of the high-contrast imaging instruments GPI (Macintosh et al. 2014), SPHERE (Beuzit et al. 2008) and ScExAO (Jovanovic et al. 2015), which are detecting and characterising young gaseous exoplanets with ever lower masses approaching that of Jupiter (Macintosh et al. 2015; Bonnefoy et al. 2016) and comprehensively studying planet-disk interactions and the planet formation process (e.g. Avenhaus et al. 2014b; Benisty et al. 2015). Such work is also informing the design parameters of the next generation of ground-based ELT-class instruments which aim to characterise rocky exoplanets in the habitable zones of nearby stars. This challenging goal requires final contrast ratios of better than 10^{-7} at inner-working angles of the order 10 milli-arcseconds Guyon et al. (2012), starting with planets orbiting M-dwarf host stars such as the newly discovered Proxima Centauri b (Anglada-Escudé et al. 2016). The expected limit on achievable raw imaging contrast with ground-based, coronagraph-enabled XAO systems is of the order of 10^{-5} for large field-of-view starlight suppression regions (Kasper et al. 2010; Guyon et al. 2012), hence this must be combined with complementary high-contrast techniques such as polarimetric differential imaging (Keller et al. 2010; Perrin et al. 2015) and high-dispersion spectroscopy (Snellen et al. 2015), which are already expanding the toolkit of the exoplanet imaging community.

Of the diverse approaches to high-contrast imaging and specifically coronagraphy, the Apodizing Phase Plate (APP) coronagraph (Codona et al. 2006; Kenworthy et al. 2010c; Quanz et al. 2010) is of particular relevance to this paper. This technique uses a pupil-plane phase mask to modify the point-spread function (PSF) of the instrument, thereby using destructive interference to create a "dark hole" in the diffracted stellar halo at the location of the planet. This approach makes the APP an extremely versatile coronagraph, allowing simultaneous coronagraphic observation of multiple targets in the same field, providing insensitivity to tip-tilt errors, and reducing the pointing tolerances on chopping offsets required for accurate background subtraction at the near-infrared wavelengths most favourable for observation of young, thermally luminous exoplanets. The recent development of the vector-Apodizing Phase Plate (vAPP, Otten et al. 2014a), which provides simultaneous 360° coverage around the host star by using circular polarisation beam-splitting to create duplicate copies of the classical APP pattern, has also accompanied significant gains in inner-working angle, with a vAPP operating at radial separations of 1.2-6 λ/D (where λ is the observing wavelength and D is the telescope diameter) installed and available for science observations in MagAO (Morzinski et al. 2014) at the Magellan Clay Telescope (Otten et al., submitted).

These ground-based, XAO-corrected high-contrast imagers are limited by ever-present Non-Common Path Errors (NCPEs); these wavefront aberrations arise due to the presence of differential optics between the AO wavefront sensor and the science focal plane, which may be influenced by slow thermal or mechanical fluctuations. The resulting focal-plane quasi-static speckle (QSS) field is coherent on timescales of min-

utes to hours, and limits the raw performance of most coronagraphs to 10^{-4} – 10^{-5} in contrast, defined here as the 5σ companion detectability limit, over an entire observation period (Martinez et al. 2012). Advanced observation and data reduction algorithms such as the Locally Optimised Combination of Images (LOCI) (Lafrenière et al. 2007) and Principle Component Analysis (Soummer et al. 2012; Amara & Quanz 2012) have been used to surpass this limit and achieve detection thresholds of 10^{-6} at separations larger than 7 λ/D with SPHERE and GPI (Zurlo et al. 2016; Macintosh et al. 2014). However, due to the impact of quasi-static speckles on the ultimate photon noise limit, in addition to ongoing uncertainties surrounding the influence of post-observation NCPE suppression algorithms on the derived properties of subsequently detected companions (e.g. Marois et al. 2010a), it remains preferable to correct these non-common path errors in real time and thereby return coronagraphic performance to the diffraction-limited regime.

The complete elimination of NCPEs ultimately relies on the principle of focal-plane wavefront sensing; only by using the science camera as a sensor can the AO loop have a truly common path with observations. Existing focal-plane wavefront reconstruction techniques use artificially induced phase diversity (Keller et al. 2012; Korkiakoski et al. 2013) or properties of the speckle field itself (Codona & Kenworthy 2013) to overcome the degeneracies associated with a loss of wavefront spatial resolution and incomplete knowledge of the focal-plane electric field. Although there have been some successful on-sky demonstrations of these techniques (e.g. Martinache et al. 2014), factors such as computational complexity, invasive modification of the science PSF, and limited dynamic or chromatic range mean that such reconstruction methods have not yet been widely adopted for science observations. To avoid these limitations many high-contrast imaging instruments instead perform periodic offline NCPE calibrations, such as the COFFEE coronagraphic phase diversity algorithm proposed for use in SPHERE (Sauvage et al. 2011), at the cost of temporal resolution and the loss of simultaneity with science observations.

There is therefore an ongoing drive to develop a coronagraphic focal-plane wavefront sensor which is able to operate in parallel with science imaging in a non-invasive manner, and provide unbiased real-time compensation of the low spatial frequency NCPEs which correspond to small angular separations in the observed stellar image.

2.1.2 Holographic optics for focal-plane wavefront sensing

The use of computer-generated holograms as a method of focal-plane wavefront sensing has been extensively explored in the literature, with specific focus on applications in confocal microscopy (Neil et al. 2000; Booth 2003) and laser collimation (Changhai et al. 2011). This approach is used to generate secondary PSF copies in the science focal plane, which are spatially separated from the main science PSF to avoid mutual interference. In the so-called Holographic Modal Wavefront Sensor (HMWS) these wavefront-sensing PSFs are artificially biased with a set of chosen aberration modes drawn from a suitable basis set (for example the Zernike modes), such that the Strehl ratio of each PSF copy responds linearly to the corresponding aberration mode present in the input wavefront. In this way the sensor performs a modal decomposition of the incoming wavefront into the chosen basis, which may be reconstructed in real time with the intensity measurement of two focal-plane photometric apertures per mode.

A modal approach to wavefront sensing has multiple advantages over traditional

wavefront sensors as well as other focal-plane wavefront sensing techniques, most notably in terms of reduced computational complexity, the fact that the resolution of the reconstructed wavefront is not limited to the spatial resolution of the sensor's pupil element as with a Shack-Hartmann sensor, and that a modal wavefront is simple to implement on many current deformable elements. For the science case of high-contrast imaging of exoplanets and circumstellar environments, the HMWS should operate simultaneously with a coronagraph in the science focal plane, to directly retrieve the aberrations that are seen by the starlight suppression system in the instrument. We therefore consider here the promising combination of the HMWS with the APP coronagraph: for the purposes of this paper we shall refer to the resulting optic as the coronagraphic Modal Wavefront Sensor (cMWS). This hybrid approach can be easily implemented since both concepts are phase-only pupil plane optics, which may be easily multiplexed into a single physical element. The HMWS is however not limited to use with pupil plane phase-only coronagraphs, provided that the hologram is positioned upstream of any focal-plane masking elements in order to transmit the central diffraction core of all holographic PSF copies.

2.1.3 Content of paper

This paper is divided into the following sections: in Sec. 2.2 we summarise the underlying mathematics behind holographic modal wavefront sensing, and present the critical factors which must be considered when multiplexing the HMWS with an APP coronagraph. Sec. 2.3 shows the results of idealised closed-loop simulations and outlines the baseline performance of the sensor for the case of a clear circular aperture. Sec. 2.4 presents results from the first on-sky implementation of a cMWS sensor at the William Herschel Telescope (WHT) located in La Palma, Spain, including a demonstration of sensitivity to both static and dynamic wavefront errors. In Sec. 2.5 we draw final conclusions and present goals for ongoing and future work.

2.2 Theory

2.2.1 The Holographic Modal Wavefront Sensor

The principle of the HMWS relies on the fact that the phase component $\phi(x,y)$ of an arbitrary wavefront may be decomposed into coefficients of a chosen 2D mode basis describing the telescope aperture, for which the complex electric field $\Psi(x,y)$ may be written as

$$\Psi(x,y) = A(x,y)e^{i\phi(x,y)} = A(x,y)e^{i\sum_j a_j M_j(x,y)}, \tag{2.1}$$

where $A(x,y)$ is the telescope aperture function, $M_j(x,y)$ is some complete, and ideally orthonormal, mode basis with RMS coefficients a_j (in radians) and x, y are coordinates in the pupil plane. In this paper we focus exclusively on phase-only aberrations as these are simpler to implement and correct for, and dominate the total wavefront error in almost all practical cases.

In order to provide full phase-aberration information in a single focal-plane intensity image, the sensor uses a computer-generated holographic element to perform an instantaneous modal decomposition and extract the set of coefficients a_j, albeit up to a truncated mode order. It is then trivial to reconstruct the wavefront using the set

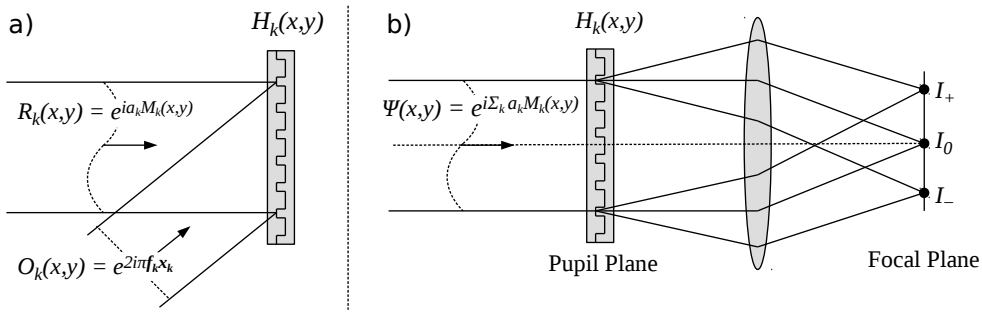


Figure 2.1: Diagram showing the principle of HMWS operation. *a*) Visual representation of the creation process of a single-mode computer-generated hologram, by analogy with optical exposure. *b*) Operation of a single-mode hologram in the presence of an aberrating wavefront. Figure adapted from Dong et al. (2012).

of template modes using Eq. 2.1, which may then be passed to an adaptive optics system for correction either as a direct command or via the adjustment of reference slope offsets.

2.2.1.1 Generating holograms

The purpose of the hologram in a HMWS is twofold: firstly it creates secondary PSF copies which are spatially separated from the zero-order PSF in the science focal plane. Secondly, it adds an artificial bias wavefront independently to each of these PSF copies, such that each responds differently to the input wavefront Ψ. This can therefore be thought of as a system of 2N simultaneous phase diversities chosen to span the desired mode basis, but instead of the normal approach to focal-plane phase diversity reconstruction (which typically uses only one diversity and the intensities of all pixels in the PSF), the modal content of the wavefront is extracted in a more direct fashion by measuring only the relative core intensities of all PSF copies.

As illustrated in Fig. 2.1a, the holographic element is constructed numerically from two independent components which perform the functions described above. Adopting the notation of Dong et al. (2012), the reference wave R_k

$$R_k(x,y) = e^{ib_k M_k(x,y)} \tag{2.2}$$

contains a single bias mode M_k with an RMS aberration strength (in radians) set by the bias strength b_k. The object wave O_k is given by

$$O_k(x,y) = e^{2i\pi(f_{kx}x + f_{ky}y)}, \tag{2.3}$$

where the spatial frequencies $f_{kx,y} = x'_k, y'_k / f\lambda$ specify the desired tilted plane wave and thus the coordinates (x'_k, y'_k) in the focal plane. The holographic phase pattern $H_k(x,y)$ for this particular mode is then the interferogram between these two waves,

$$\begin{aligned} H_k(x,y) &= \left|O_k(x,y) + R_k(x,y)\right|^2 && (2.4)\\ &= |O_k|^2 + |R_k|^2 + O_k^* R_k + O_k R_k^* && (2.5)\\ &= 2 + 2\mathbb{R}\left[O_k^* R_k\right], && (2.6) \end{aligned}$$

where * is the complex conjugate operator and $\mathbb{R}\,[]$ denotes the real component of the complex argument.

It follows from this that the two conjugate terms naturally result in the creation of two wavefront sensing spots which may be treated as the ±1 orders of a diffraction grating, containing equal and opposite bias aberrations $\pm b_k$. The first two terms in Eq. 5 are equal to unity and are discarded such that $\langle H_k \rangle = 0$. The behaviour of this hologram in the presence of an aberrated wavefront Ψ is shown graphically in Fig. 2.1b. The total focal-plane intensity is then given by $I = \left|\mathcal{F}\left[H_k \Psi\right]\right|^2$, where $\mathcal{F}$ is the Fourier Transform operator in the Fraunhofer diffraction regime. Following from this and Eqs. 2.1, 2.2, 2.3 and 2.6, the local intensity distribution $I_{k\pm}$ of the pair of biased PSF copies is given by

$$\begin{aligned} I_{k\pm}(x', y') = & \underbrace{\delta(\boldsymbol{x'} \pm \boldsymbol{x'_k})}_{\text{Carrier Frequency}} * \underbrace{\left|\mathcal{F}\left[A(x,y)\right]\right|^2}_{\text{Telescope PSF}} \\ & * \left|\mathcal{F}\left[\underbrace{e^{i(a_k \pm b_k) M_k(x,y)}}_{\text{Wavefront Bias}} \underbrace{e^{i \sum_{j \neq k} a_j M_j(x,y)}}_{\text{Inter-Modal Crosstalk}} \right]\right|^2, \end{aligned} \tag{2.7}$$

where $I_{k\pm}$ correspond to the positively and negatively biased wavefront sensing spots respectively, and a_j is the RMS error present in the incident wavefront corresponding to mode M_j. The term $\delta(\boldsymbol{x'})$ is the 2D delta function, with focal-plane coordinates $\boldsymbol{x'_k} = (x'_k, y'_k)$ deriving directly from the frequency of the carrier wave O_k. The second term encompasses the desired sensor response to the aberrated wavefront, with net aberration $a_k \pm b_k$. The final term represents a fundamental source of inter-modal crosstalk as a convolution with all other modes present in the input wavefront, which acts equally on both $I_{k\pm}$; see Section 2.2.1.3 for a full discussion of the impact of this term.

An arbitrary number of holograms may be multiplexed into a single element, allowing the generation of multiple pairs of independently biased PSF copies and hence the simultaneous coverage of many wavefront modes. For simplicity of implementation we now create a phase-only hologram $\phi_h(x, y)$ by taking the argument of the multiplexed hologram

$$\phi_h(x,y) = \frac{s}{\pi} \arg\left[H(x,y)\right] = \frac{s}{\pi} \arg\left[\sum_k^N H_k(x,y))\right], \tag{2.8}$$

which is by definition binary as all H_k are real from Eq. 2.6, and is normalised to have a grating amplitude of s radians. Scaling down the amplitude from $(0, \pi)$ allows direct control over the fractional transmission to the zeroth order, which forms the science PSF. It is assumed here that the holographic PSF copies are located sufficiently far from each other and the zeroth-order in the focal plane that there is negligible overlap; if this is not the case there will be additional inter-modal crosstalk in the sensor response due to mutual interference, which is independent of that arising from the final term of Eq. 2.7.

The optimal positioning of WFS copies for minimal inter-modal crosstalk is a significant optimisation problem in itself, which will be investigated in future work. As a rule of thumb, each spot should be positioned at least 5-6 λ/D from not only all other first order PSF copies, but also from the locations of all corresponding higher-order diffraction copies and cross-terms; see the treatment in Changhai et al. (2011) for full details. In the general case this requires the computation of an appropriate

non-redundant pattern, which is outside the scope of this paper, however a circular or "sawtooth" geometry (the latter is shown in Fig. 2.3) was found to be a suitable alternative geometry for the prototype cMWS.

2.2.1.2 Sensor response

Following the approach of Booth (2003) it is possible to approximate the sensor response for $a_k \ll b_k$ as the Taylor expansion of Eq. 2.7 about $a_k = 0$, where the on-axis intensity of each PSF copy can this way be expressed as

$$I_{k\pm} = I_0 \left[f(b_k) \pm a_k f'(b_k) + \frac{a_k^2}{2} f''(b_k) + O(a_k^3) \right] \tag{2.9}$$

where I_0 is a multiplicative factor proportional to total spot intensity and $f(b_k) = \left| 1/\pi \iint e^{ib_k M_k(x,y)} dxdy \right|^2$ is the Fourier integral for an on-axis detector of infinitesimal size. Throughout this paper we adopt the normalised intensity difference between spot pairs as the metric for sensor measurement, equivalent to the "Type B" sensor of Booth (2003). In this case, the sensor response per mode I_k is given by

$$I_k = \frac{I_{k+} - I_{k-}}{I_{k+} + I_{k-}} = \frac{2a_k f'(b_k) + O(a_k^3)}{2f(b_k) + a_k^2 f''(b_k) + O(a_k^5)}. \tag{2.10}$$

If b_k can be chosen such that $f''(b_k) = 0$, this expression becomes linear to 3rd order: for a Zernike basis Booth (2003) find that this occurs for values of $\langle b_k \rangle = 1.1$ rad, while values of $\langle b_k \rangle = 0.7$ rad resulted in maximal sensitivity; we adopt the latter value throughout this work. In principle the improved "Type C" sensor also suggested by Booth (2003), which uses the metric $I_k = (I_{k+} - I_{k-})/(I_{k+} + \gamma I_0 + I_{k-})$, can yield further improved linearity and suppression of intermodal crosstalk, however the inclusion of additional measurement requirements of an unbiased PSF copy I_0 and free parameter γ (which must be determined empirically) make this unnecessary for use in a first implementation of the sensor.

2.2.1.3 Wavefront reconstruction

Final estimates of the mode coefficients a_k of the incoming wavefront must then be obtained by calibrating intensity measurements with a cMWS response matrix $\hat{\boldsymbol{G}}$, which provides the nominal scaling factors between I_k and a_k but is also capable of providing a first-order correction for inter-modal crosstalk via its off-diagonal terms. This matrix is formed from the gradients of the characteristic response curves $I_i(a_j)$ response curves about $I = 0$,

$$G_{ij} = \left. \frac{\delta I_i}{\delta a_j} \right|_{I=0}, \tag{2.11}$$

where G_{ij} is the response of sensor mode i to input wavefront error of mode j. The solution for the set of mode coefficients $\boldsymbol{a}$ of the incoming wavefront is then

$$\boldsymbol{a} = \hat{\boldsymbol{G}}^{-1} \boldsymbol{I}, \tag{2.12}$$

where $\boldsymbol{a}$ and $\boldsymbol{I}$ are the column vectors comprising sensor response I_k and the corresponding wavefront coefficient estimates a_k respectively. Note that the standard multiplicative inverse $\hat{\boldsymbol{G}}^{-1}$ of the interaction matrix is used here, since the interaction matrix is square, highly diagonal and with on-diagonal elements defined so as to have

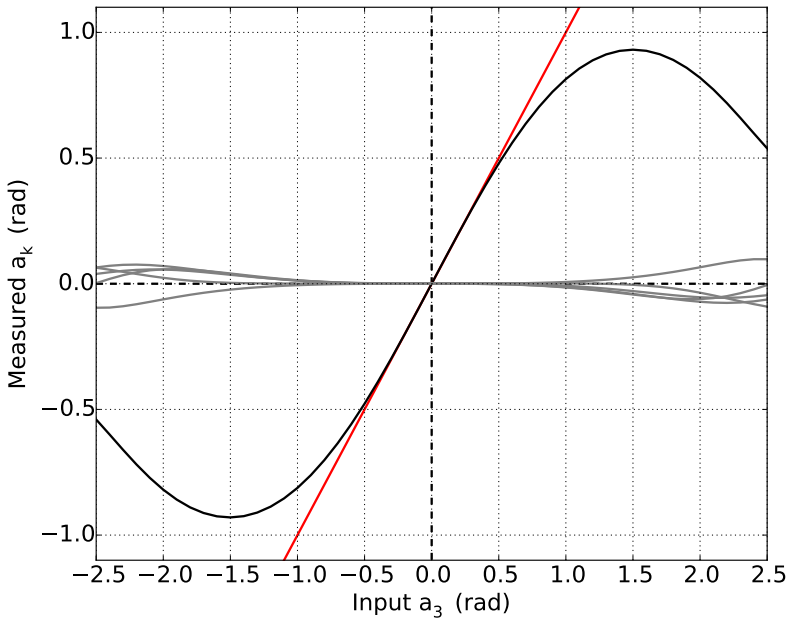


Figure 2.2: Response curve of a 6-mode Zernike HMWS ($b_k = 0.7$ rad) to defocus error $a_3 Z_3$ (all $a_{k\neq 3} = 0$), using photometric apertures of radius $r = 1.22\lambda/D$. The diagonal line (red) denotes a perfect sensor with 1:1 correspondence, which is achieved by the HMWS for $|a_3| \lesssim 0.5$. The response of the remaining sensor modes (Z_4-Z_8, grey) are well constrained about zero over the linear sensing regime, with residual non-linear inter-modal crosstalk behaviour manifesting for $|a_3| \gtrsim 1$.

the same sign. It is therefore extremely unlikely that this matrix is degenerate and thus non-invertible, but in such a case the Moore-Penrose pseudo-inverse $\hat{\boldsymbol{G}}^{+}$ (see e.g. Barata & Hussein 2012) may be used as an alternative. Fig. 2.2 shows an illustrative response curve to which this calibration has been applied, showing that the sensor response is linear over the range $|a_k| \lesssim b_k$ with negligible inter-modal crosstalk, beyond which wavefront error is increasingly underestimated as the main assumption of Eq. 2.9 begins to break down. A turnover in sensitivity occurs at the point $a_k = 2b_k$ since beyond this the input wavefront error dominates over the differential bias $\pm b_k$.

In addition to calibrating sensor measurements to physical units, $\hat{\boldsymbol{G}}$ also performs a linear correction for inter-modal crosstalk; this allows the knowledge of the responses of all other spot pairs to be used to infer the correct mode measurement of one particular pair. As denoted by the final term in Eq.2.7, this effect occurs via a convolution of the WFS spot I_k with all remaining wavefront aberrations $M_{j\neq k}$ present in the input wavefront. This effect was neglected in the previous section as the convolution term is reduced to a constant multiplicative factor under the on-axis assumption, factoring out in Eq. 2.10. The theoretical response matrix for any set of orthogonal modes is therefore diagonally dominated and sparse (Booth 2003), but in practice many factors such as use of photometric apertures of non-zeros size, alignment errors or overlap with the wings of other PSF copies or the zeroth order, may result in significantly elevated crosstalk behaviour.

Empirical determination of a full response matrix for each cMWS design is therefore the most robust method of compensating these effects to first order. This process is straightforward and once automated takes only a few minutes to perform: each

column of the interaction matrix requires a minimum of two measurements of the normalised intensity vector $\boldsymbol{I}$, each for different known coefficients a_k of the corresponding input wavefront mode applied on the corrective element, in order to fit the gradients of each response curve. This procedure is in principle required only once for any given instrument configuration, however performing regular re-calibration before each observation night is feasible and allows the elimination of slow drifts in actuator response or instrument alignment quality.

2.2.2 Combination with an Apodizing Phase Plate coronagraph

The APP is an optimal coronagraph for use in the cMWS as not only is it a pupil plane phase only optic and thus simple to multiplex with the HMWS, but is also preserves an Airy-like PSF core required for production of holographic copies. By contrast, focal-plane or hybrid coronagraphs would require the hologram to be located in pupil upstream of the focal-plane mask in order to create the off-axis PSF copies before rejection of on-axis stellar light occurs. The resulting optic may be implemented using the same techniques as for the APP; as either a transmissive optic such as a turned glass phase plate (Kenworthy et al. 2010c) or achromatic liquid crystal retarder (Snik et al. 2012), or via a phase-apodizing Spatial Light Modulator (SLM) (Otten et al. 2014a).

Consider now the combination of the HMWS presented above with an APP coronagraph into a single optic such that the modification to the complex wavefront $\Psi_{cMWS}(x,y)$ may be described as

$$\Psi_{cMWS}(x,y) = A(x,y)e^{i[\phi_c(x,y)+\phi_h(x,y)]}, \tag{2.13}$$

where $\phi_c(x,y)$ and $\phi_h(x,y)$ correspond to the coronagraph and normalised hologram (Eq. 2.8) phase patterns respectively.

Fig 2.3 shows the simulated pupil optic and corresponding PSF of a cMWS coded for the 14 lowest order non-trivial Zernike modes, including an APP with a 180 degree dark hole extending from $2.7 - 6\lambda/D$, generated using a Gerchberg-Saxton style iterative optimisation algorithm. The hologram pattern is seen in the pupil as an irregular binary grating overlaid on top of the smooth phase variations of the APP. The wavefront sensing spots can clearly be seen surrounding the dominant central science PSF, with the PSF of each copy formed by the convolution of the characteristic Zernike mode PSF with that of the APP. For illustration purposes a grating amplitude of $s = \pi/2$ here results in an average normalised intensity difference of -1.8dex between the peak flux of each WFS copy and the zeroth order PSF, with an effective transmission to the science PSF of 50%. It is however possible to operate the sensor with significantly fainter PSF copies in practice, making $80-90\%$ transmission achievable with respect to the APP alone.

2.2.3 Impact of multiplexing on mutual performance

As the zeroth order PSF may be considered a 'leakage' term of the binary hologram grating, the APP pattern is in principle independent of all wavefront biases which appear in the ± 1 diffraction orders. However there are two notable effects which must be considered when multiplexing these two optics, the first of which is that any stray light scattered by the HMWS will fill in the coronagraphic dark hole. As shown in Fig. 2.4,

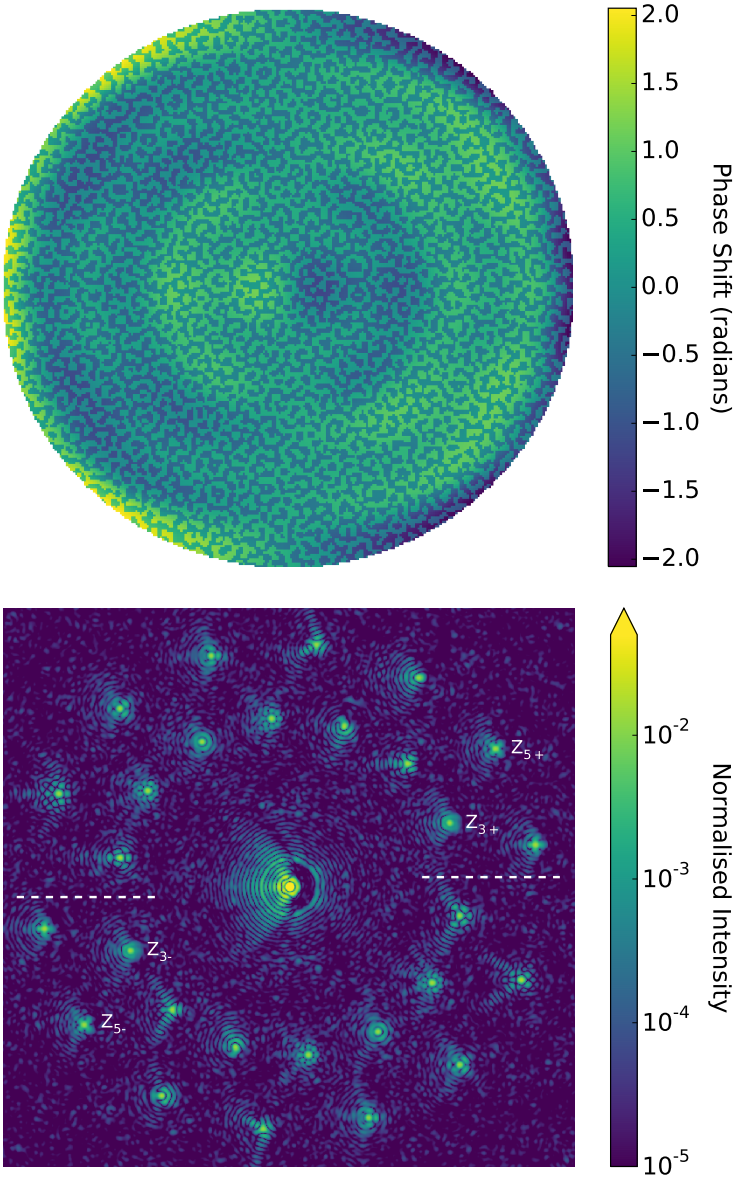


Figure 2.3: Simulation of a 14-mode Zernike cMWS (modes $Z_3 - Z_{16}$) combined with an APP optimised for a 10^{-6} dark hole with a 180° opening angle. ***Top:*** Multiplexed pupil-plane phase design Ψ_{cMWS} containing the high spatial frequency HMWS binary grating overlaid on the smoother APP design. ***Bottom:*** Corresponding focal-plane PSF: positively and negatively biased PSF copies are located in the top and bottom half of the image respectively, separated by the white dashed lines. Two example pairs are labelled (defocus, Z_3 and 45° astigmatism, Z_5), illustrating the symmetry of the ±1 orders about the zeroth order PSF.

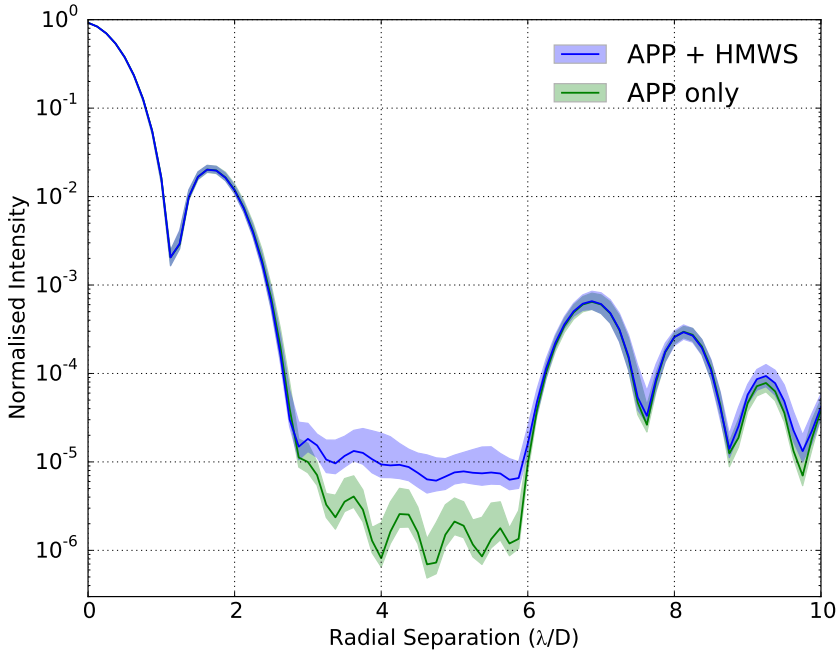


Figure 2.4: Contrast curves for the APP coronagraph presented in Fig. 2.3, with (blue) and without (green) the HMWS hologram. The shaded regions denote the 1-sigma variance limit of residual structure at each radius, azimuthally averaged over a 170 degree region corresponding to the dark hole contrast floor.

it was found that the binary holograms generate a near-constant intensity scattered background at a mean normalised intensity of the order of 10^{-5}, irrespective of the specific HMWS or APP designs used. This behaviour is due to the loss of information associated with creating a binary optic from the full complex hologram in Eq. 2.8. Although a limiting dark hole depth of 10^{-5} remains sufficient for a first prototype, it would be possible to compensate for this effect by re-optimising the APP in the presence of the scattered background.

The second effect of the multiplexing process is that, as can be seen from Eq. 2.13, the APP phase pattern introduces a set of static wavefront errors which must be disregarded by the HMWS. This can be achieved by adding static reference slope offsets to to Eq. 2.12 in a similar manner to existing NCPE correction routines (e.g. Sauvage et al. 2011), such that

$$\boldsymbol{a} = \hat{\boldsymbol{G}}^{-1}\boldsymbol{I} - \boldsymbol{a_c} \tag{2.14}$$

where $\boldsymbol{a_c}$ is the set of coefficients of ϕ_c in the sensing mode basis. This must be determined independently from $\hat{\boldsymbol{G}}$ to avoid degeneracy with static instrumental wavefront errors, either by projecting the APP onto the sensing mode basis $a_{c,i} = \phi_c(x,y)\cdot M_i(x,y)$, or by comparison with calibration data containing only the non-multiplexed HMWS component.

2.2.4 Impact of structured telescope apertures

It is important to note that throughout this paper the cMWS is evaluated for use with an un-obscured circular aperture, however it must also be applicable to more com-

plicated amplitude profiles featuring central obscurations, support spiders, and mirror segmentation. If no modifications to the cMWS design are made, any aperture modifications will degrade the orthogonality of the chosen mode basis and thus lead to increased inter-modal crosstalk. Fortunately this is not considered to be a limiting factor of the cMWS, as the effect can be effectively eliminated by performing a re-orthogonalisation of the chosen mode basis using the known aperture function, for example by using a simple Gram–Schmidt procedure (see e.g. Cheney & Kincaid 2009). This approach has now been verified during a more recent observing campaign at the WHT, the details of which will be the subject of a future work. In the case where the aperture function contains significant structures which are not azimuthally uniform, such as especially thick telescope spiders or mirror segmentation gaps, this procedure will be most effective when operated in a pupil-stabilised observation mode. This will allow the telescope aperture function to remain consistent with that of the re-orthogonalised sensing basis for the duration of each observation, however it was seen that in the case of the WHT pupil the 1.2 m circular central obscuration was in practice the only significant structure.

It is in principle also possible to develop the cMWS as a co-phasing sensor for segmented mirrors, for which the ideal sensing basis would instead consist of differential piston, tip and tilt modes which directly match the degrees of freedom of each individual mirror segment. That being said, the cMWS is not an ideal choice of sensor for co-phasing large future segmented telescopes such as the European Extremely Large Telescope (E-ELT) or the Thirty Meter Telescope (TMT), principally because the sensing basis would need to consist of an unreasonably large number of modes (2394 in the case of the E-ELT) in order to fully describe all possible phasing errors. While it may be possible to achieve this by sequentially correcting with multiple cMWS designs each containing a subset of the possible modes, such applications are much more suited to telescopes with significantly fewer mirrors where the calibration may be performed for all segments simultaneously, such as the W.M. Keck Observatory in Hawaii, or the Giant Magellan Telescope (GMT).

2.3 Idealised performance simulations

To analyse the baseline performance of the multiplexed sensor, we consider the ideal case where the aberrating wavefront consists entirely of modes to which the HMWS is sensitive. To demonstrate the interchangeability of the sensor mode basis, two distinct sensor designs are considered, which for ease of comparison both utilize six sensing modes each with bias $b_k = 0.7$ and an APP dark hole of radial extent $2.7 - 6\lambda/D$. Sensor A encodes the first six non-trivial Zernike modes (Defocus Z_3 to Trefoil Z_8) while Sensor B contains six sinusoidal 2D Fourier modes of the form $\cos((n_x X + n_y Y) + c)$, where c is equal to either 0 or π, optimised to probe three critical locations at radial separation $3.5\lambda/D$ within the APP dark hole. The diffraction-limited PSFs of these sensors can be seen in Fig 2.5b, with PSF copies showing the characteristic PSF of each sensing mode. Note that the APP of the Zernike cMWS is optimised for a $180°$ opening angle while the Fourier cMWS contains an APP optimised for $90°$, which explains the differences between the two diffraction-limited zeroth-order PSFs.

Aberrating wavefronts are generated with equal RMS wavefront error a present in each mode, giving a total RMS wavefront error $\sigma_\phi = \sum_k a_k M_k = a\sqrt{6}$ for a perfectly or-

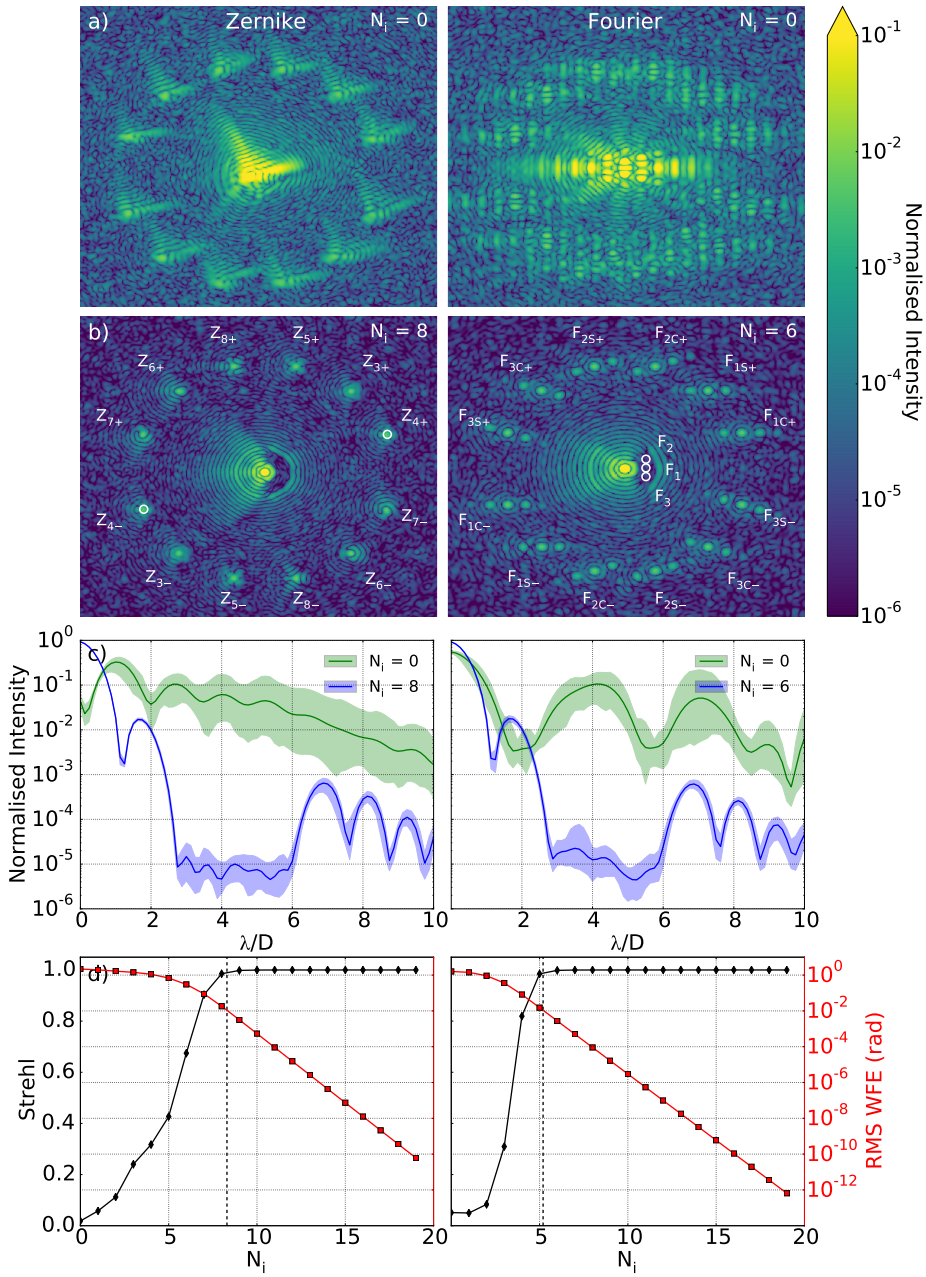


Figure 2.5: Example closed-loop performance for a 6-mode Zernike (*left column*) and Fourier (*right column*) mode cMWS for $g = 0.8$. *a*) $N_i = 0$ aberrated PSFs, with $a_k = 1.0$ radians RMS per mode. *b*) Diffraction-limited PSF after N_i closed-loop iterations required to achieve convergence. PSF copies corresponding to each mode bias are labelled $Z_{n\pm}$ for Zernike modes and $F_{nS/C\pm}$ for Fourier modes, where the index S/C denote the sine and cosine mode phases respectively, and mode number corresponds to the circled regions of influence in the APP dark hole. The white circles overlaid on the $Z_4\pm$ modes indicated the $r = 1.22\lambda/D$ region of interest used for wavefront measurement. Note also the differing angular extent of each APP, which cover 180° and 90° for the Zernike/Fourier designs respectively. *c*) Azimuthally-averaged residual intensity plots corresponding to the PSFs of panels *a*) (green) and *b*) (blue); shaded regions denote 1σ variance averaged over the APP dark hole. *d*) Science PSF Strehl ratio (black diamonds) and residual RMS wavefront error (red squares) as a function of iteration number N_i. Vertical dashed lines indicate the point of convergence.

thogonal 6-mode basis. In order to probe the upper limit of closed-loop convergence a is varied between 0.1 and 1.5 radians RMS per mode, significantly exceeding the nominal ± 0.5 radians RMS per mode linear range of the sensor. The response matrix is constructed according to Eq. 2.11 from a simulated calibration dataset, and compensation for the APP mode coefficients applied as per Eq. 2.14. Photometric apertures of radius $r_s = 1.22\lambda/D$ are applied to each PSF copy for flux measurement, which has been shown to provide optimal sensitivity for small b_k (Booth 2003). Closed-loop correction is then achieved by direct phase conjugation using a perfect simulated deformable mirror with phase $\Phi_{DM,i} = \Phi_{DM,i-1} - g\Sigma_k^N a_k M_k$, with the closed-loop gain g left as a free parameter. Convergence is taken to be achieved at iteration N_i where the total wavefront error a_r is reduced below 10^{-2} radians RMS, which is seen to correspond closely to the point at which the diffraction-limited PSF is recovered.

The panels of Fig 2.5 shows one example of closed-loop convergence for both sensors, with initial wavefront error of $a_k = 1.0$ radians RMS per mode (and thus total wavefront error $\sigma_\phi = 2.45$ radians RMS) and a closed-loop gain $g = 0.8$. It can be seen that despite this large initial wavefront error both sensors efficiently recover diffraction-limited APP performance with iteration number N_i, with residual wavefront error continuing to decline logarithmically towards the numerical noise threshold after nominal convergence is achieved. In this case the remaining intensity structure in the dark hole is limited purely by the HMWS scattered light background for each APP design. It is unclear exactly why the Fourier mode basis exhibits significantly faster convergence in this example, but a probable explanation is that the large coma aberration present as part of both APP designs pushes the Zernike mode sensor into the non-linear regime and thus lowers the initial measurement accuracy of this mode, whereas this same aberration is distributed more evenly in the Fourier mode basis.

Figure 2.6 characterises in detail the convergence efficiency of the Zernike mode sensor by considering a wide variety of closed-loop gains g and input RMS wavefront errors a_r. Both panels show that the critical failure point of this sensor lies at $a_k = 1.1$ radians RMS per mode and is independent of gain value. Below this, convergence speed is purely gain-limited for $g < 0.8$ and $g = 1$ provides the most efficient convergence for all a_k, ranging from $2 < N_i < 7$ iterations and with final $N_i = 20$ solutions consistent with the diffraction-limited wavefront at the level of numerical noise. This robust high-gain convergence behaviour stems from systematic underestimation of the wavefront outside the linear range (see Fig. 2.2), preventing oscillatory instabilities from occurring. The rapid breakdown in convergence above $a_k = 1.1$ happens when the contribution of non-linear intermodal crosstalk between 6 modes of equal a_k becomes comparable to the individual sensor response, enabling sign errors and thus irreversible divergence. The equivalent surface plots for the Fourier-type sensor was seen to be morphologically identical, confirming that the HMWS is capable of operating with any mode basis that is sufficiently complete with respect to the power spectrum of wavefront error present in the system.

It is important to note that the term “idealised” here refers to the fact that no artificial noise sources such as readout or photon noise are included in these simulations, and that the underlying light source is purely monochromatic and point-like in nature. Such factors are dealt with during the on-sky implementation of the cMWS presented in Sec. 2.4 of this paper; in this section we instead aim to demonstrate that fundamental factors such as the multiplexing process and inter-modal crosstalk do not limit the

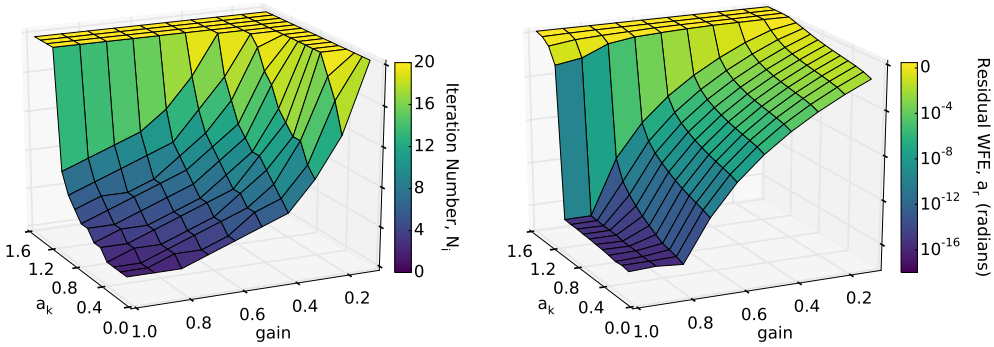


Figure 2.6: Convergence properties of the Zernike cMWS shown in Fig. 2.5 as a function of initial RMS wavefront error per mode, a_k, and closed-loop gain, g. ***Left:*** Number of iterations required until convergence $N_i(a_r < 10^{-2})$. A value of $N_i = 20$ indicates that convergence was not achieved within the allowed 20 iterations. ***Right:*** Residual wavefront error a_r after the final iteration.

final convergence of the closed-loop correction process. This explains why the residual wavefront error as presented in Fig. 2.5 reaches the numerical noise limit in both examples; this will not be the case in practice as noise sources will result in sporadic random errors in measuring the wavefront coefficients. In the absence of systematic errors this can be expected to stall the convergence process at the level of $\sim 10^{-1}$ radians RMS based on the error bars derived in Sec. 2.4, although this ultimately depends upon the signal-to-noise ratio (SNR) of individual WFS spots on a target-by-target basis. As presented in Sec. 2.4.3, use of a broadband source turns the holographic PSF copies into radially dispersed spectra, which can be useful in its own right for wavelength selection of the wavefront estimates.

2.4 On-sky demonstration

2.4.1 Instrument design

To implement the sensor on-sky at the William Herschel Telescope, we used a setup based around a BNS P512 reflective Spatial Light Modulator (SLM) as shown in Fig. 2.7, similar to that described in Korkiakoski et al. (2014). This was operated with 250 pixels across the pupil diameter, oversampling the cMWS designs by a factor of two in order to ensure the sharp boundary regions of the HMWS hologram are accurately represented. Use of an SLM allows the rapid testing of a wide variety of designs without the need to manufacture individual custom optics, but has the disadvantage of allowing only passive measurement of wavefront errors: the response rate of the SLM was seen to approach 1 Hz at times and as such is not a suitable active element for real-time phase correction. The SLM phase response was calibrated at the He-Ne 633 nm line via the differential optical transfer function (dOTF) wavefront reconstruction method of Korkiakoski et al. (2013), at which the SLM is able to produce a maximum stroke of 1.94π radians. This stroke limitation to less than 2π is unimportant as all chosen designs have peak-to-peak phase values of less than π radians. The sensor was then operated on-sky with both narrowband (650 nm, $\Delta\lambda = 10\ nm$) and broadband (Bessel-R 550-900 nm) filters, with the latter possible despite strong chromatic behaviour of SLM devices (see Sec 2.4.3 for further discussion). A high-frame

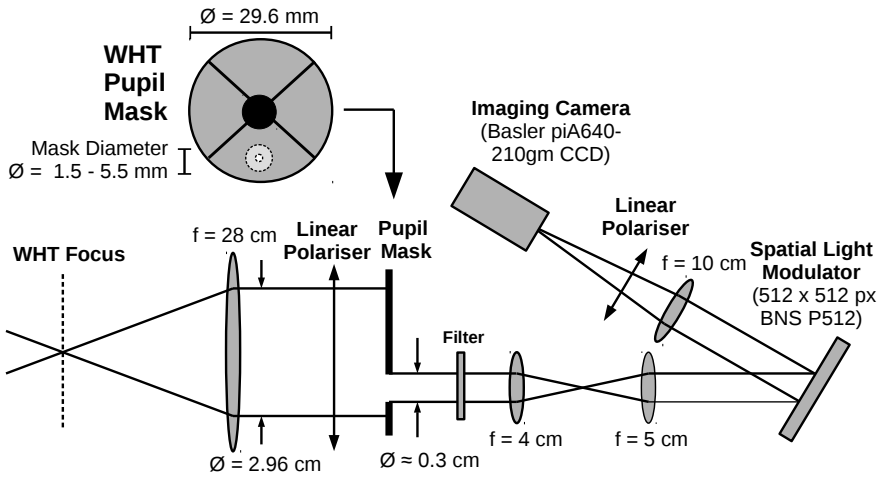


Figure 2.7: Diagram illustrating setup used at the WHT. Interchangeable pupil masks allowed the use of telescope sub-apertures with effective diameters ranging from 0.2 m to 0.8 m, to control relative aberration strengths in the absence of a classical AO system.

rate Basler piA640-210gm CCD camera was used to record the focal plane including the holographic WFS spots at a cadence of 50 Hz, comparable to atmospheric seeing timescales.

It was necessary to limit on-sky wavefront error to within the dynamic range of the sensor, which in the absence of an AO system was achieved by stopping down the WHT aperture. For this purpose an off-axis circular pupil stop was used to create an un-obscured sub-aperture of effective diameter 42.3 cm, positioned in the pupil so as to be free of telescope spiders over the elevation range 30 deg to zenith. This aperture size was chosen based on the expectation values of low-order Zernike coefficients of a pure Kolmogorov phase screen, which are constrained to $0.1 \lesssim |a_k| \lesssim 0.5$ radians RMS for the 0.7”-2.5” range of seeing conditions typical of La Palma.

Two calibration images of a 6-mode Zernike HMWS with uniform bias value $b = 1.5$ radians RMS at the calibration wavelength are shown in Fig 2.8, for a flat wavefront and for 1.5 radians RMS of defocus error introduced on the SLM. For ease of illustration, a grating amplitude of $s = 3\pi/4$ radians results here in an effective Strehl ratio of 24% compared to the un-aberrated PSF. This illustrates clearly the sensor response: since no APP is applied in this instance, the holographic copy which is biased with a focus aberration of equal amplitude but opposite sign ($b_k = -a_k$) collapses to the Airy diffraction function, while the conjugate WFS spot gains double the aberration. It should be noted that in addition to three faint filter ghosts below the zeroth order PSF, there is a significant ghost located at approximately $3\lambda/D$ which proved impossible to eliminate via optical re-alignment. This is attributed to unwanted reflection from the SLM glass cover plate which thus bypasses the active surface; a conclusion which is supported by its presence adjacent to both the central PSF and each filter ghost but not diffracted PSF copies, plus its independence of SLM-induced defocus.

An in-situ calibration of the HMWS response matrix was obtained by sequentially introducing aberrations $a_k M_k$ with the SLM. It was found however that this solution contained linear inter-modal crosstalk components (off-diagonal terms in $\hat{\boldsymbol{G}}$) on the same order as sensor linear response. This effect is not seen in simulations nor indeed in un-calibrated normalised intensity data when the defocus error is generated with an

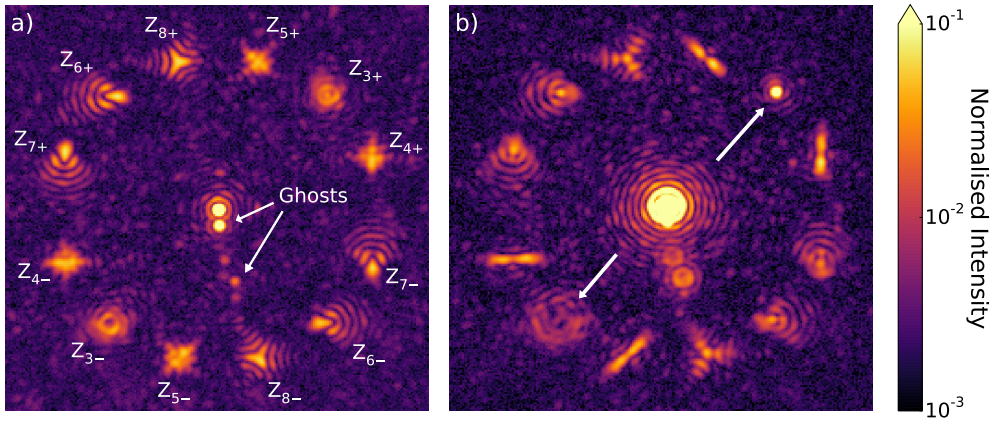


Figure 2.8: 633 nm He-Ne laser calibration data illustrating the response of a 6-mode HMWS to controlled wavefront aberrations. *a)* PSF with no induced wavefront error and *b)* with 1.5 rad of defocus error applied via the SLM, showing the asymmetric response of the $Z_3\pm$ defocus WFS spot pair, indicated *top-right/bottom-left*. Filter and SLM reflection ghosts are present below the zeroth-order PSF, though these are sufficiently separated from the PSF copies to ensure no interference with HMWS performance.

external source (see the following section). The effect is therefore attributed to errors in accurately recreating Zernike modes with the SLM due to uncontrolled spatial variations in SLM voltage-to-phase actuation response across the pupil, which degrades the mode orthogonality. We therefore rely on simulated response matrix solutions for the following analysis of on-sky sensor data.

2.4.2 Characterising HMWS on-sky response

To the authors' knowledge the HMWS has never before been implemented on-sky: the first and most important test was therefore to verify the on-sky response of the HMWS alone to known, static wavefront errors under realistic observing conditions. This is particularly important with respect to the ultimate goal of NCPE sensing, as the cMWS must be able to accurately recover coherent errors from underneath a dominant incoherent atmospheric speckle foreground. For this purpose, narrowband 650 nm observations were made of Arcturus ($m_R = -1.03$) using the HMWS design of Fig. 2.8 while scanning over a range of focus positions with the WHT secondary mirror, thereby inducing defocus error ranging between ± 2 radians RMS in a controlled manner. 60 seconds worth of 20 ms exposures were stacked for each focus position in order to increase the signal-to-noise ratio and sufficiently average out atmospheric variations. The core intensity measurements of each WFS spot were then extracted using a numerical photometric aperture mask comprising a set of circular apertures of radius $1.22\lambda/D$, aligned with the centroid location of each WFS spot. The resulting set of intensity measurements was then converted to mode coefficient estimates a_j using Eqs. 2.10 and 2.14.

Figure 2.9 shows the calibrated sensor response of all modes as a function of WHT focus position (mm), and the corresponding input defocus error a_3 in radians RMS. The mm-to-radians RMS scaling factor was obtained by least-squares fitting of the theoretical defocus response curve of Fig. 2.2 (plotted here in black) to the defocus data, which is seen to be closely consistent for input $a_3 < 1.5$ radians RMS. It is un-

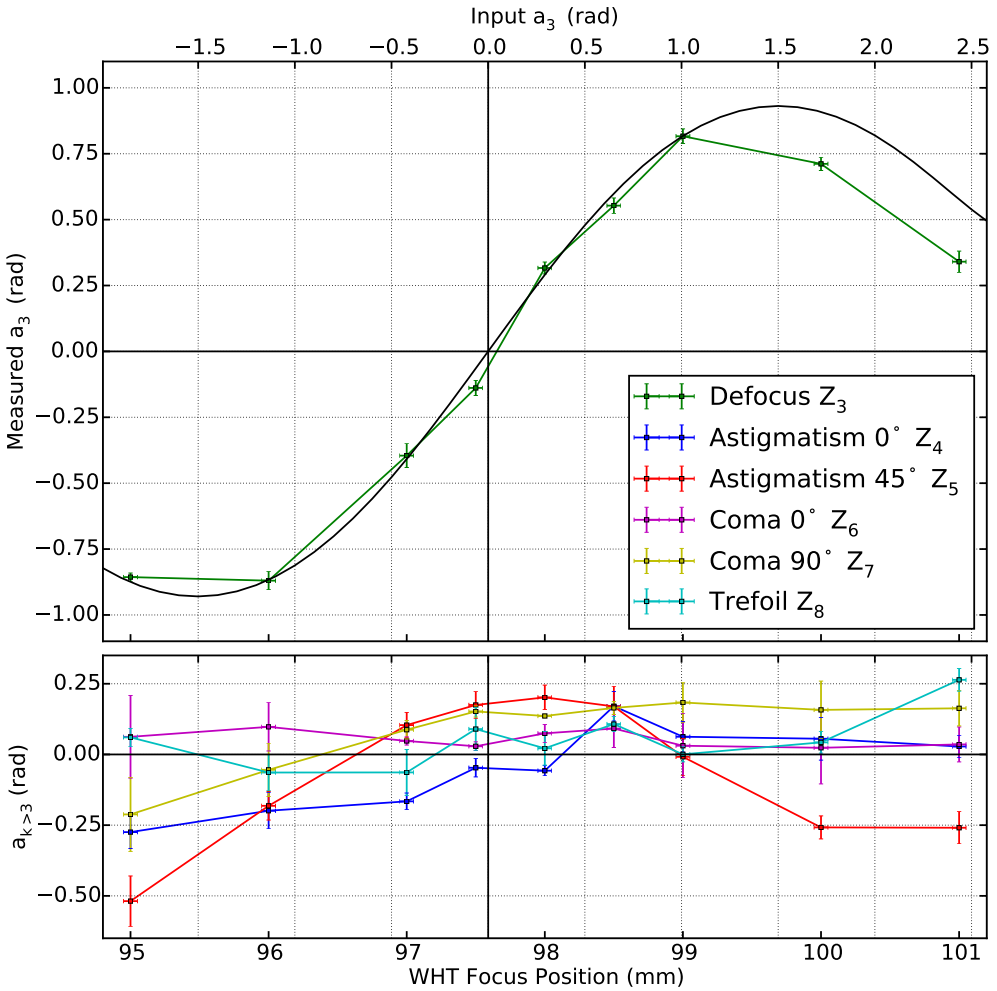


Figure 2.9: On-sky response of the Zernike HMWS sensor as a function of WHT secondary mirror focus position (mm) and corresponding induced defocus error a_k (radians RMS). ***Top:*** Response of the Z_3 defocus mode. Over-plotted in black is the theoretical a_3 response curve of Fig. 2.2, seen to correspond closely to sensor measurements over the the range $a_3 < 1.5$ radians RMS. ***Bottom:*** Response of non-defocus modes $k > 3$. Error bars correspond to uncertainties in sub-pixel centering accuracy (photon and CCD noise sources are negligible). Systematic trends in the non-defocus modes are attributed to real, instrumentally-induced wavefront aberrations upstream of the pupil mask.

clear why the final two points are underestimated with respect to the theoretical curve, but even assuming this is a real effect, the additional wavefront underestimation in this non-linear regime would have little impact on closed-loop sensor performance. The error bars on each curve represent the uncertainty in frame alignment, specifically the 1σ standard deviation of sensor measurements associated with the complete set of possible 1 pixel translational and rotational offsets of the photometric aperture mask, which was seen to be the practical limit on frame alignment accuracy. It was found that this source of uncertainty dominates over photon and readout noise when analysing seeing-averaged images; this places a limit on the precision of cMWS wavefront retrieval in the high-SNR regime of 0.04 radians per mode, or equivalently 0.1 radians RMS total wavefront error, a value obtained from the mean derived 1σ error bar of all six sensing modes where the input focus error is within the $a_3 = \pm 1$ radian RMS dynamic range of the sensor. Being azimuthally symmetric, the defocus mode is seen to be significantly more robust against small (x',y') offsets or rotations of the photometric aperture mask compared to other modes, even for large wavefront errors. Stability against positioning and/or tip-tilt errors is therefore a worthwhile consideration in choice of mode basis for future sensor designs.

It can be seen that the response curves of the other sensed modes depart significantly from the well-behaved off-diagonal terms in Fig. 2.2; in particular the astigmatism mode Z_5 displays strongly quadratic response as a function of defocus error, which cannot be corrected using a linear interaction matrix and may therefore lead to closed-loop instabilities where large focus errors are present. This behaviour indicates that either the true sensor response is not fully characterised by the simulated interaction matrix, in which case experimental calibration is necessary, or that the injected wavefront error contains variable components other than pure defocus. This second hypothesis has been further explored because of the complexities associated with using the WHT secondary mirror as the source of injected focus error when we are sampling only an off-axis sub-pupil, as illustrated in the "WHT Pupil Mask" inset of Fig. 2.7. It was found in simulations that the observed crosstalk behaviour can be recreated if there are also static higher-order wavefront errors present upstream of the pupil mask, created for instance by small mis-alignments of the upstream lens or polariser. These aberrations are not orthogonal on the Zernike basis of the sensor and can mix with the variable focus error when sampled in this way, creating a spectrum of focus-dependent higher-order Zernike aberrations which cannot be explained with the modes Z_{0-3} alone due to symmetry arguments. In this case it was found that adding an upstream wavefront error of 0.2 radians RMS of Z_4 (astigmatism of the opposite orientation) results in a Z_5 response curve which is morphologically similar to that shown in the lower panel of Fig. 2.9. Low-amplitude variations seen in the coma and trefoil response curves are more difficult to recreate using static errors of the same order and may therefore be due to other factors, but it is expected that the principle remains the same when including additional Zernike mode orders $Z_{k>8}$. By applying the correct set of upstream instrumental aberrations in this way it may be possible to account for the complete discrepancy between the non-defocus mode response of Fig. 2.2 and the lower panel of Fig. 2.9. However, due to the complexity of this effect a comprehensive treatment is beyond the scope of this paper, which is in any case specific to the non-standard pupil apodisation used in this setup and is therefore not expected to be present in subsequent observing campaigns.

It is sufficient to note that the defocus response is consistent with theoretical pre-

dictions, while the majority of the remaining sensor behaviour can be explained by unintentionally introduced additional instrumental wavefront errors and not to fundamental crosstalk effects, which appear to be accurately compensated by the theoretical response matrix calibration procedure. This confirms that the sensor is able to accurately recover (quasi-)static errors underneath a dominant fluctuating atmospheric speckle foreground, simply by integrating up to the desired timescale.

2.4.3 Broadband wavefront sensing

It is also important to characterise the broadband performance of the cMWS; a major limitation of focal-plane phase-retrieval algorithms such as phase diversity is that they only work effectively in the monochromatic case (Korkiakoski et al. 2014). By contrast, the HMWS contains no such fundamental limitations; the normalised difference metric is independent of variations in spectral transmission $T(\lambda)$, while the natural λ-dependence of the radial position of diffracted holographic PSF copies raises the intriguing possibility of performing wavelength-resolved wavefront sensing. Spatial light modulators also typically exhibit strong chromatic response variations away from the calibration wavelength (e.g. Spangenberg et al. 2014), but since all wavefront bias information is encoded into spatial variations in the binary hologram, only the effective grating amplitude s and thus $T(\lambda)$ may vary with wavelength rather than bias b_k. Altogether, the cMWS is in principle capable of delivering unbiased estimates of the wavefront coefficients a_k in radians RMS for arbitrarily wide spectral bands, at a spectral resolution set by the diffraction limit of the monochromatic telescope PSF.

Figure 2.10a shows the broadband on-sky PSF of a cMWS including the same 6-mode Zernike HMWS as in Fig. 2.8, operated with a standard 50% bandwidth Bessel-R filter. To test the full cMWS concept, this design also includes a 180 degree APP as in Fig. 2.3; in this seeing-limited image the dark hole is located in the top half of the PSF, although it is obscured by residual speckles and chromatic dispersion. In Fig. 2.10b it can be seen that the chromatic response of each mode is broadly consistent over the FWHM transmission range of 580-750 nm. A residual focus error can be clearly seen from these on-sky observations, such that this plot corresponds to the wavelength dimension of $a_3 \approx -0.3$ rad in Fig. 2.9. This mode also displays the $a_k \propto 1/\lambda$ scaling expected from physical wavefront errors.

Confirmation of a bias-free spectral response allows boosting of single-frame SNR by binning the 580-750 nm spectrum in the radial and hence wavelength dimension, making quasi-real-time wavefront sensing, with exposure times t_{exp} approximately equal to the NCPE coherent timescale τ_ϕ, a possibility. Wavelength-resolved wavefront sensing may also be achieved by using appropriately calibrated photometric subapertures along the dispersed wavefront spectra, and will be considered further in future work. Such information may be useful for optimisation of the broadband control of existing AO systems or for next-generation instrument concepts consisting of of multiple corrective elements for specific wavelength ranges.

2.4.4 Real-time atmospheric wavefront measurements

Application of the broadband sensor provided sufficient SNR for partial wavefront retrieval from individual 20 ms frames. Fig. 2.11a shows two independent estimators of RMS wavefront error σ_ϕ for a subset of frames: cMWS measurements $\sigma_{\phi,cMWS} =$

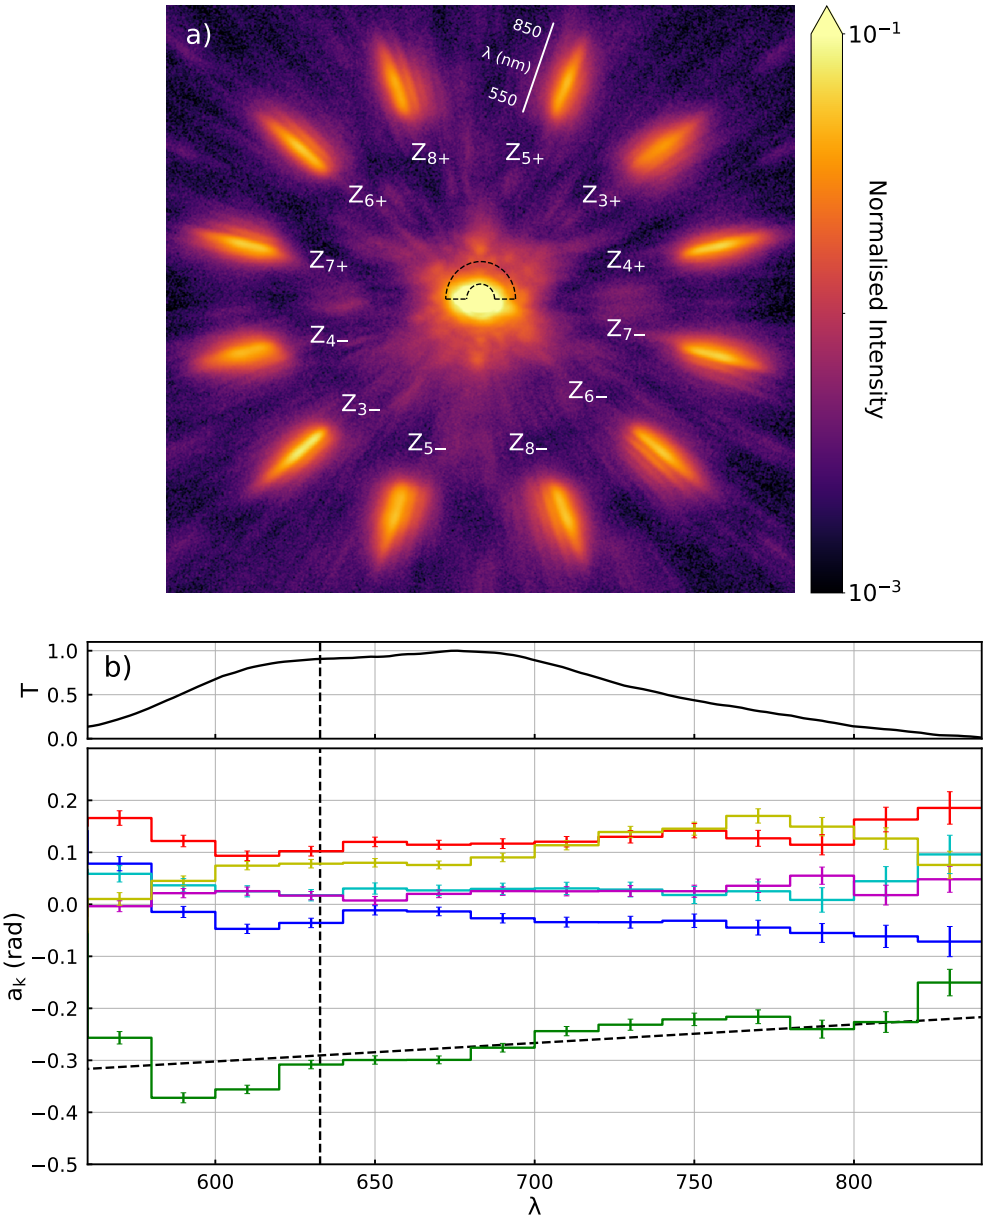


Figure 2.10: *a*) Broadband on-sky PSF of a cMWS incorporating the HMWS of Fig. 2.8, showing $1/\lambda$ radial dispersion of PSF copies. The boundary of the APP dark hole at $\lambda = 633$ nm is illustrated by the dashed line, spanning $2.7-6\lambda/D$, although this is not directly visible in the stacked image data. *b*) Chromatic response of each mode, binned to $\Delta\lambda = 20$ nm wavelength intervals; colours correspond to the modes of Fig. 2.9. The 633 nm calibration wavelength is shown by the vertical dashed line, and the mean spectral transmission $T(\lambda)$ is shown in the upper panel. Defocus Z_3 (green) shows indications of $1/\lambda$ scaling, illustrated by the diagonal dashed line.

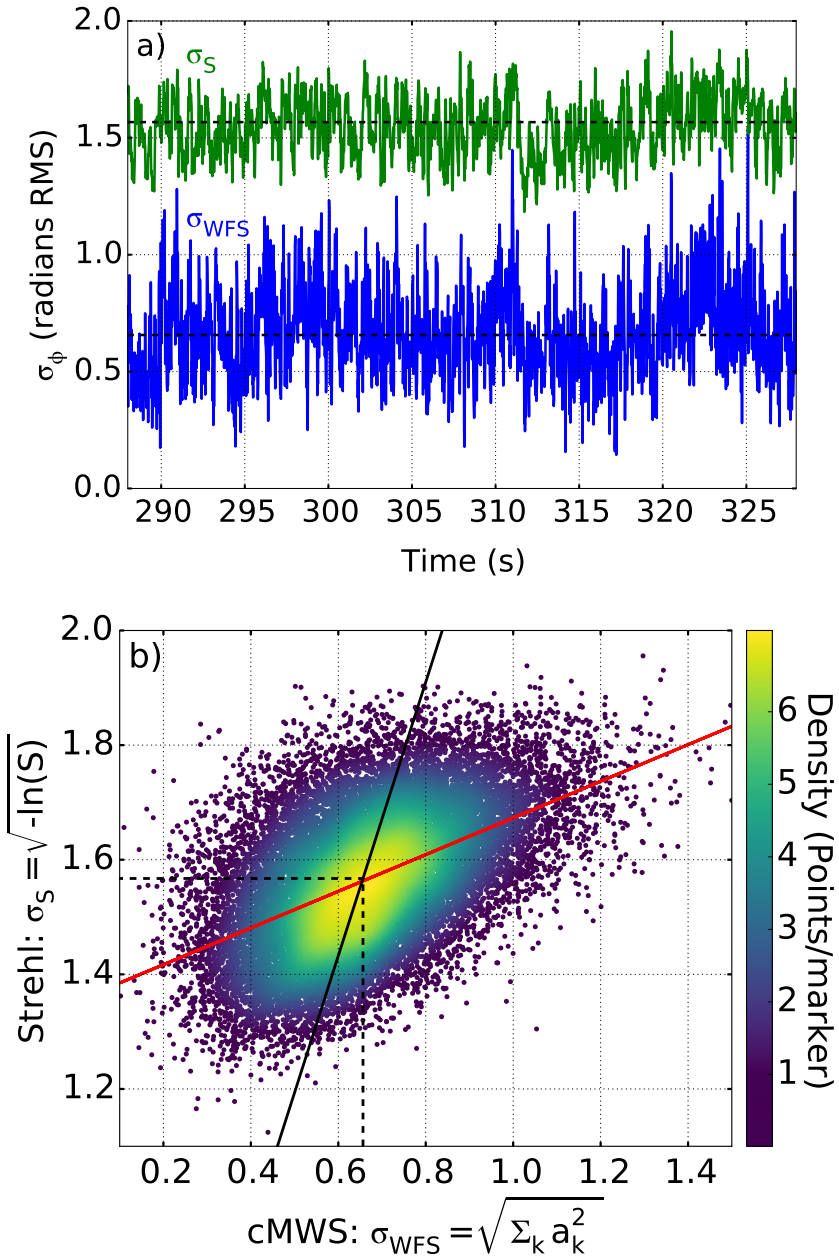


Figure 2.11: *a*) Time series of single-frame RMS wavefront error σ_ϕ as measured by i) the central PSF Strehl ratio under the Maréchal approximation (σ_S, green) and ii) calibrated cMWS measurements (σ_{WFS}, blue). Slow variations in seeing quality seen in σ_S are visibly traced by the sensor. *b*) Correlation plot between the two independent estimates of RMS wavefront error, with colour indicating point density. The solid black line denotes $\sigma_S/\sigma_{WFS} = 0.45$ as is expected theoretically from Noll (1976), which corresponds well to the core regions of the correlation ($1.45 < \sigma_S < 1.65$). The outer regions have a significantly shallower gradient (red), distorted by various systematic error sources.

$\sqrt{\Sigma_k {a_k}^2}$ (lower, blue curve) and $\sigma_{\phi,S} = \sqrt{-\ln(S)}$ from Strehl ratio measurements of the science PSF under the Maréchal approximation (upper, green curve). Despite significant noise in the measurements, the cMWS measurements trace slow trends in image Strehl ratio on timescales > 1 s, associated with changing seeing conditions. Fig. 2.11b shows the resulting correlation between these two frame quality estimators for the full 20,790 frame dataset spanning 10 minutes of observation, with a Pearson correlation index of $\rho = 0.50$. Frames with $|I_k| > 1$ were rejected as such measurements are obviously unphysical: such events are rare ($< 1\%$ of total frames affected) and attributed to cosmic ray impacts and residual hot/cold pixels. Additional confirmation that the cMWS is tracing the atmospheric wavefront is provided by the respective mean wavefront error estimates: $\langle \sigma_{\phi,cMWS} \rangle = 0.656 \pm 0.001$ and $\langle \sigma_{\phi,S} \rangle = 1.567 \pm 0.001$ radians RMS, such that on average the cMWS senses approximately 42% of the total wavefront error once known static errors have been subtracted. This is notably consistent with what is expected for pure Kolmogorov statistics as presented in Noll (1976), where 45% of the total wavefront error for $Z_{k>3}$ is contained in the first 6 non-trivial modes; this relation $\sigma_S/\sigma_{WFS} = 0.45$ is denoted by the solid black line in Fig. 2.11b. Ideally the two independent estimates should correlate along this relation for all values of wavefront error, but although there is reasonable agreement about the mean $\sigma_{WFS} = 0.65$ radians RMS, it can be seen that the correlation is significantly shallower for outlying points beyond $\sigma_{WFS} > 0.8$ radians RMS; here the best fit line, plotted in red, clearly does not intersect the origin. This may be attributed to crosstalk with high-order unsensed modes allowing the sensor to pick up some additional wavefront error to that contained purely in the 6-mode basis, or to systematic effects such as sensor saturation, making these extreme wavefront estimates unreliable. However, it is important that the majority of sensor measurement points fall close to the theoretically expected relation, where sensor performance is expected to be most reliable.

The ultimate goal of this process is to reconstruct the instantaneous wavefront in each frame. As may be anticipated from Fig. 2.11a however, such wavefronts were seen to be dominated by frame-to-frame noise. In order to assess the extent to which the independent mode coefficient measurements are degraded, we plot the modal power spectrum of the full dataset in Fig. 2.12 and contrast with that expected from Kolmogorov turbulence as rescaled to a six-mode basis. It can be seen that although there is some morphological similarity which indicates a decreasing power spectrum, the amplitude of individual modes is significantly more consistent with a flat spectrum. It is possible that the true seeing statistics are not Kolmogorov in nature, however it is difficult to justify a discrepancy of such size in this manner. Instead, it is assumed that this is due to the mixing effect of crosstalk with higher-order un-sensed modes which cannot be accounted for by the response matrix; only in this way is it possible to preserve the total wavefront variance as discussed above. The immediate solution for residual atmospheric wavefront error sensing is to increase the number of modes to encompass a larger fraction of the total power spectrum. For the application to NCPE correction of a dark hole the problem is made simpler as the power spectrum is expected to be dominated by low-order components, which may be accessed by integrating so as to sufficiently average out the unwanted high-order atmospheric errors.

Due to the dominance of frame-to-frame noise at a cadence of 50 Hz we therefore draw only limited conclusions regarding the potential of the cMWS for real-time wavefront correction in this instance, however the successful retrieval of total wavefront error σ_ϕ at this cadence is already a promising result for such a preliminary test.

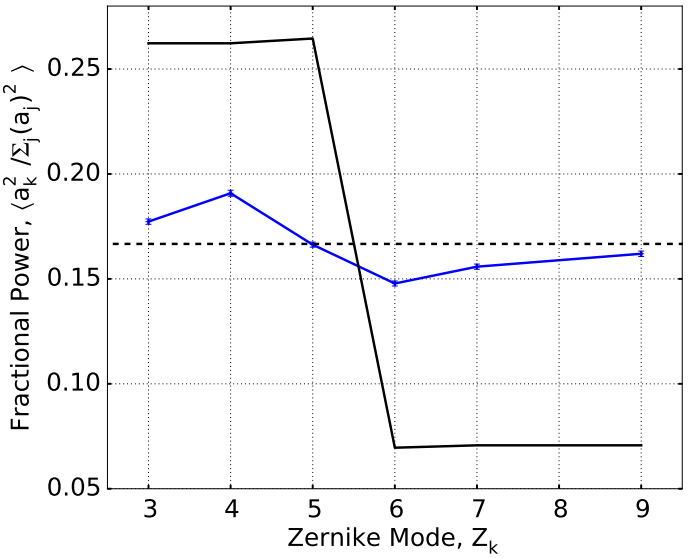


Figure 2.12: Modal power spectrum of on-sky broadband cMWS measurements as a function of Zernike mode order. The theoretical spectrum corresponding to purely Kolmogorov statistics (Noll 1976) is over-plotted in black, while the horizontal dashed line denotes a purely flat 6-mode power spectrum.

More important is that, as shown in Fig. 2.9, the cMWS is capable of recovering known static aberrations to a precision of approximately 0.04 radians RMS per mode with one-shot measurements. This is performed in the presence of a dominant and fluctuating atmospheric speckle foreground, in direct parallel with the ultimate goal of direct NCPE sensing.

2.5 Discussion & conclusions

We have demonstrated via idealised closed-loop simulations and a first on-sky implementation that the coronagraphic Modal Wavefront Sensor (cMWS) is a promising new focal-plane sensor for high-contrast imaging, highly suited to correction of non-common path errors (NCPEs) and with additional potential as a high-cadence broadband wavefront sensor. The major advantage of the cMWS over prior focal-plane reconstruction algorithms is that the measurement process requires no invasive modification of the science PSF, as is required for phase diversity approaches; this allows the correction loop to be effectively decoupled from science observations.

The performance of the cMWS is not limited by the process of multiplexing the APP coronagraph and Holographic Modal Wavefront Sensor (HMWS) components or by structures in the telescope aperture function, but by residual inter-modal crosstalk with higher-order un-sensed modes present in the wavefront. This can be addressed by using a larger sensing mode basis than the 6-mode cMWS prototype presented in this work, such that a larger fraction of the total wavefront error is encompassed by the sensor. The correction order of the cMWS is currently limited not by fundamental factors, but by the practical consideration of science PSF throughput. This may be optimised with respect to the signal-to-noise ratio of the holographic PSF copies and hence observational target brightness, however the practical limit in most cases is ex-

pected to be 20-30 modes. While expected to be sufficient for NCPE correction, this is too small to allow the removal of a classical AO sensor from the instrument design. It is however possible to avoid such limitations for applications which require only a small but extremely well-corrected field of view, such as spectroscopic characterisation of known exoplanets. We have already discovered that it is possible to manufacture APP coronagraphs which reach simulated contrasts of 10^{-10} in dark regions a few square λ/D in size (Keller 2016). These regions contain few degrees of freedom in the electric field such that they may be fully corrected with only a small basis of optimised modes.

An additional advantage of the cMWS is its computational simplicity, requiring only the relative photometry of the diffraction cores of $2N_{\mathrm{mode}}$ holographic PSF copies and a small number of linear computations for the calibration process; most importantly it does not require any Fourier transforms. This is unimportant for NCPE sensing due to the slow timescales involved, but an additional application of the sensor is then to the challenge of extremely high-cadence sensing, for the control of a limited number modes at kHz frequencies. Such an approach is expected to lead to significant improvements in wavefront quality over conventional AO update frequencies (Keller 2016). As a phase-only sensor, the on-sky performance of the cMWS will always be fundamentally limited by instrumental amplitude errors. This may be overcome by combining it with other focal-plane sensing techniques, such as electric field conjugation (Give'on et al. 2006), which are capable of reconstructing the full electric field but which lack the dynamic range to perform effectively by themselves in ground-based AO systems. The improved "Fast and Furious" algorithm of Korkiakoski et al. (2014) also lends itself to use with the cMWS, which naturally provides a large number of known phase diversities in the holographic PSF copies.

Future work will focus on implementing the optimised cMWS behind a 97-actuator AO system with a classical Shack-Hartmann WFS, previously used with the ExPo high-contrast imaging polarimeter (Rodenhuis et al. 2011). In addition to providing a significant boost in SNR, this will allow the cMWS to be tested in a realistic closed-loop environment which reflects the ultimate goal of real-time NCPE control. If successful, such a system would be ideal for inclusion into the next-generation of high-contrast imaging instruments such as EPICS for the E-ELT (Kasper et al. 2010), for the detection and characterisation of rocky exoplanets in the habitable zones of nearby stars.

Chapter 3

A 'Fast & Furious' solution to the low-wind effect for SPHERE at the VLT

Michael J. Wilby, Christoph U. Keller, Jean-François Sauvage, Thierry Fusco, David Mouillet, Jean-Luc Beuzit, and Kjetil Dohlen
Proc. SPIE 9909, Adaptive Optics Systems V, 99096C (2016)

Abstract

We present a potential non-invasive solution to sensing the so-called low-wind effect (LWE) seen in the SPHERE instrument at the VLT, based on the 'Fast and Furious' (F&F) sequential phase diversity wavefront reconstruction algorithm. This uses non-coronagraphic focal-plane images available from the near-infrared differential tip-tilt sensor (DTTS), with the closed-loop correction cycle itself providing the necessary phase diversity between frames required to reconstruct the full wavefront phase. Crucially, this means F&F does not need to apply large artificial phase probes as required by standard phase diversity algorithms, allowing it to operate in a real-time (~10 Hz) correction mode without impacting science observations. In this paper we present the results of realistic closed-loop AO simulations designed to emulate SPHERE/DTTS observations of the LWE. With this we demonstrate that the F&F algorithm is capable of effective removal of the characteristic point-spread function (PSF) aberrations of strongly LWE-affected images within a few closed-loop iterations, with the final wavefront quality limited only by the corrective order of the deformable mirror. The ultimate goal of this project is to provide an independent, real-time and focal-plane wavefront sensor for SPHERE which is capable of detecting and directly compensating the LWE as it arises, thus improving coronagraph performance under the best 15-20% of observing conditions where the effect is most pronounced.

3.1 Introduction

3.1.1 The low-wind effect

The Spectro-Polarimetric High-contrast Exoplanet REsearch instrument (SPHERE, Beuzit et al. 2008) for the Very Large Telescope (VLT) has recently finished its commissioning phase and is already providing exceptional results in the field of direct exoplanet and circumstellar disk imaging (Vigan et al. 2016a; Maire et al. 2016; Zurlo et al. 2016; Bonnefoy et al. 2016). The performance of the instrument is however limited under the best seeing conditions by the so-called low-wind effect (LWE), where the un-occulted stellar point spread function (PSF) gains two or more bright side-lobes (dubbed 'Mickey Mouse ears') on spatial scales of 2 to 4 λ/D and relative intensity of 10^{-1} or greater with respect to the central PSF core; a number of examples of the effect are shown in Fig 3.1 as observed by the SPHERE differential tip-tilt sensor (DTTS). The appearance of the LWE is strongly correlated with the most favourable observing conditions where the dome-level wind speed at the VLT drops below 3 ms^{-1}. The effect is present at some level in 15-20% of observations and therefore leads to significant degradation of the performance of the focal-plane coronagraph masks of both the visible and near-infrared (NIR) instrument arms of SPHERE.

In the current absence of direct sensing or control solutions for the LWE, the inclusion of an additional and complementary wavefront sensor is necessary in order to identify and suppress this effect without interrupting nominal science observations. The prototype Zernike sensor for Extremely Low-level Differential Aberrations (ZELDA, N'Diaye et al. 2014; N'Diaye et al. 2016) was installed into the NIR coronagraph filter wheel in 2014, converting the science camera of the IRDIS subsystem into a pupil-plane wavefront sensor. The high-resolution phase maps obtained during its testing phase provided a clear interpretation of the observed LWE-affected PSFs of Fig. 3.1, and form the basis for the wavefront models used in this paper. However, before a proposed upgrade to SPHERE is performed in which ZELDA can be fully integrated as a stand-alone wavefront sensor, it is currently not possible to operate the sensor in parallel with IRDIS or IFS science observations. An example wavefront model and corresponding PSF based on ZELDA data is shown in Fig 3.2. From this is can be seen that each 'ear' present in the PSF is a direct result of strong differential piston, tip and tilt phase error components over a single VLT pupil segment, resulting in quadrant-to-quadrant phase discontinuities which have been observed on-sky to reach amplitudes of up to 800 nm.

The current most widely accepted hypothesis regarding the origin of the LWE attributes it to changing internal dome conditions, whereby slow wind flowing across the VLT secondary mirror support spiders allows for significant thermal transfer between air and the metal surface, creating an optical path difference across each spider which gradually reduces in amplitude further downwind as the air returns to normal temperature. However, the resulting wavefront is entirely invisible to the filtered Shack-Hartmann wavefront sensor of the SPHERE extreme adaptive optics (XAO) system, SAXO (Sauvage et al. 2014; Fusco et al. 2014), making it difficult to explain the existence of quadrant-sized tip/tilt errors using only sources upstream of SPHERE. Attempts to eliminate the effect via modification of SAXO closed-loop parameters or the management of dome airflow conditions have thus far been unable to demonstrably reduce the strength of effect (J.-F. Sauvage & J.-L. Beuzit, private communication,

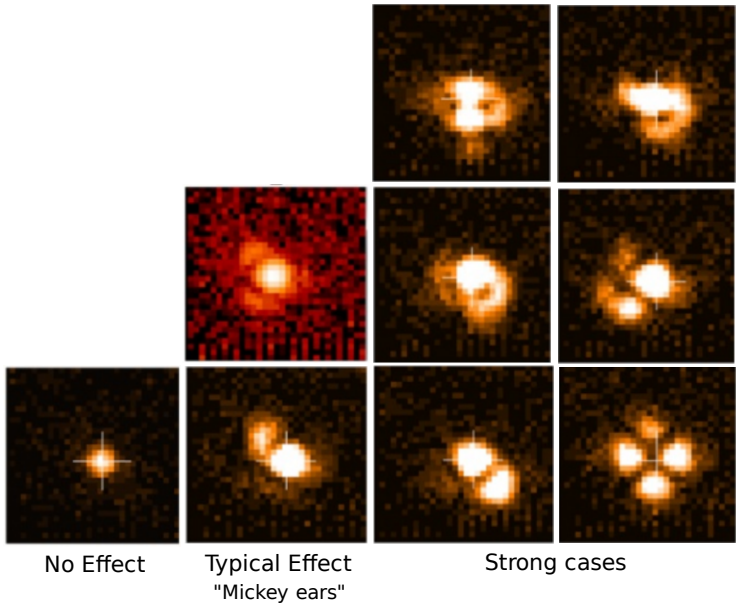


Figure 3.1: Example images of the diverse PSF morphologies which have been associated with the LWE, obtained with the SPHERE DTTS. The most extreme case (lower right) corresponds to significant tip/tilt error across all 4 VLT pupil segments and a central Strehl ratio of approximately zero. (Adapted from the SPHERE User Manual v.96, ESO 2015)

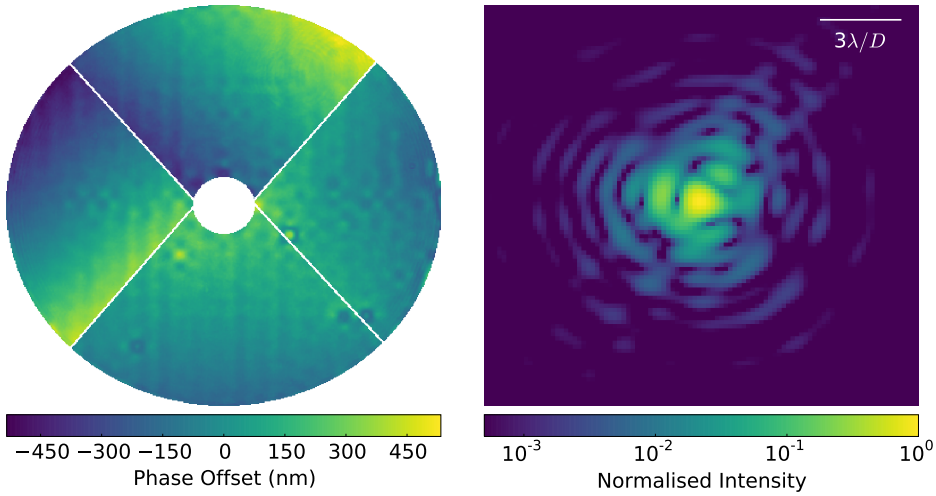


Figure 3.2: *Left:* Model VLT pupil with an example LWE wavefront based on measurements made with the ZELDA wavefront sensor, including phase error model due to faulty deformable mirror actuators. *Right:* Corresponding aberrated PSF displaying two main diffracted lobes at the location of the first airy ring (~$2.5\lambda/D$), which correspond directly with tip/tilt error across individual pupil segments in the phase error map.

2016).

3.1.2 Proposal: Focal-plane wavefront sensing using the SPHERE DTTS

In the absence of a dedicated LWE wavefront sensor within the commissioned SPHERE instrument, it is necessary in the short-term to develop a wavefront sensing approach which simultaneously meets the following criteria:

1) Provides continuous estimates of the LWE wavefront phase on second-to-minute timescales, matching the fastest observed evolution of the effect.
2) Uses only hardware currently implemented in SPHERE
3) Does not generate additional overhead or down-time for science observations
4) Compensates measured LWE without impeding the main atmospheric correction loop of SAXO.

In this paper we propose the so-called Fast and Furious (F&F) sequential phase diversity algorithm (Keller et al. 2012; Korkiakoski et al. 2014) as a focal-plane wavefront sensor which meets all of the above criteria. In order to work effectively this algorithm uses sequential images of the narrowband, non-coronagraphic stellar PSF, with frames separated by known (but not manually controlled) phase offsets. Unlike standard phase diversity, which typically requires one image in each pair to be strongly aberrated by a known phase probe such as defocus (Gonsalves 1982; Sauvage et al. 2007), F&F uses the smaller frame-to-frame phase diversity generated by actively correcting the measured aberrations to provide the additional information necessary to reconstruct the full wavefront phase.

Fortunately a subsystem already exists within the SPHERE instrument which meets these requirements, in the form of the DTTS imaging camera (Baudoz et al. 2010). This consists of a Hawaii I 1024×1024 detector operating at H-band ($\lambda_0 = 1.53\ \mu$m, 3% bandwidth) located behind a 2% beam splitter immediately upstream of the NIR coronagraph wheel, as shown in Fig. 3.3. During normal observations this runs at frame-rates between 1-1000 Hz (depending on target magnitude) to ensure the stellar image remains accurately centred on the focal-plane mask of the NIR Lyot coronagraph. These properties match extremely well with the requirements of F&F, making the DTTS a good source of real-time focal-plane images for use with the algorithm.

F&F has already been shown to be capable of performing reliable closed-loop PSF optimisation in idealised simulations (Keller et al. 2012) and optical bench tests (Korkiakoski et al. 2014) for the purpose of sensing and correcting non-common-path wavefront errors (NCPEs). The on-sky performance of the algorithm or its behaviour in the presence of complicating factors, such as residual incoherent atmospheric speckles and significant image noise, has however not yet been studied. The main goal of this paper is therefore to verify that the F&F algorithm is also able to robustly compensate the LWE under realistic SPHERE observing conditions. To this end, in this paper we present the results of closed-loop simulations which emulate real DTTS images by including factors such as the strong LWE, XAO-corrected atmospheric speckles, NCPEs and significant detector noise sources.

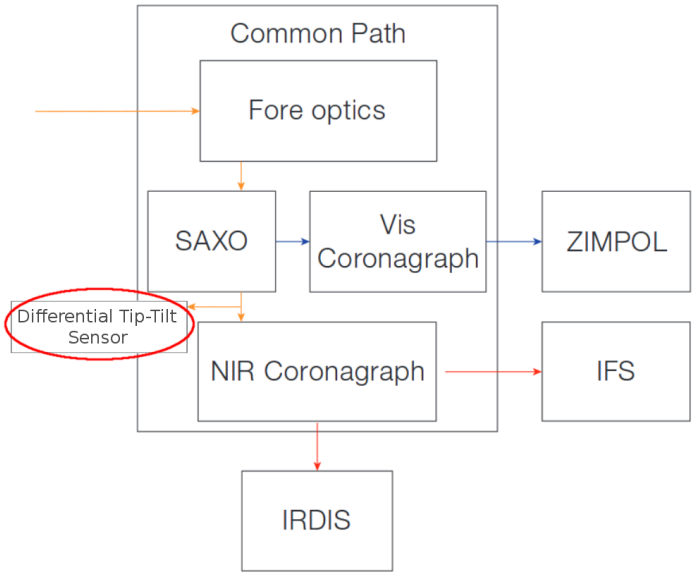


Figure 3.3: Block diagram of the SPHERE instrument, indicating the location of the DTTS directly upstream of the NIR coronagraph focal plane. (Adapted from Beuzit et al. 2008)

3.2 The Fast & Furious algorithm

The term 'Fast and Furious' here refers to a modified sequential phase diversity algorithm developed in Keller et al. (2012) and Korkiakoski et al. (2014), designed for focal-plane phase-retrieval without requiring the application of PSF-degrading phase probes. This makes it suitable to sense NCPEs in real time using the science focal plane, without interrupting science observations. F&F achieves this by using a second-order expansion of the aberrated stellar intensity image in terms of the un-aberrated telescope PSF and the Fourier transforms of the odd and even aberrating wavefront for which we need to solve. This solution is degenerate in the signs of the even focal-plane field, which must be broken with the use of a phase diversity; this is provided by the previous image in the observation sequence, for which the wavefront differs by a known phase offset.

This approach has the following advantages over conventional phase diversity and other focal-plane phase retrieval algorithms, such as the differential optical transfer function (dOTF) approach (Codona 2013):

1) No destructive phase probes or amplitude masks: closed-loop update cycle provides the necessary PSF diversity
2) Fast: requires only one 2D spatial Fourier transform and a small number of linear operations per iteration
3) Requires no physical modification to instrument components during operation (e.g. camera defocus)
4) Makes no assumptions on the spatial order or specific morphology of the wavefront phase.

Below we present a summary of the key equations required for a single closed-loop iteration of F&F, the full details of which can be found in Korkiakoski et al. (2014).

This process is also illustrated schematically in Fig. 3.4 for an example wavefront containing one even and one odd Zernike mode. For convenience, focal-plane quantities are denoted by lower case variables, while the corresponding pupil-plane quantities are represented by the equivalent upper-case variables. All quantities are two-dimensional unless otherwise specified.

The second-order expansion of the aberrated PSF can be written as

$$p = S a^2 + 2a(ia * \phi_o) + (ia * \phi_o)^2 + (a * \phi_e)^2 \tag{3.1}$$
$$= S a^2 + 2ay \quad + y^2 \quad + v^2, \tag{3.2}$$

in terms of the complex un-aberrated focal-plane electric field $a = \mathscr{F}[A]$, where A is the telescope aperture function, and the even and odd components of the aberrating focal-plane fields $\phi_{o,e} = \mathscr{F}[\Phi_{o,e}]$ respectively, with odd/even symmetry defined to be consistent with that of the Fourier transform operator $\mathscr{F}[\,]$. $*$ denotes the convolution operator. The scalar normalisation factor $S \approx (1 - \sigma_\phi^2)$ is the image Strehl ratio, applied as a second-order correction to improve energy conservation, where σ_ϕ^2 is the wavefront variance. For convenience we have defined the following focal plane quantities y, v as the (complex) focal-plane fields of the even and odd aberrating wavefronts, following the notation of Gonsalves (2001):

$$y = i\mathscr{F}[A\Phi_o] = (ia * \phi_o), \text{ and} \tag{3.3}$$
$$v = \mathscr{F}[A\Phi_e] = (a * \phi_e), \tag{3.4}$$

which for an even aperture function are real quantities due to Fourier transform symmetries, although should in general be treated as complex for real aperture functions such as the VLT pupil which include asymmetric spiders. The problem then reduces to estimating v and y by splitting Eq. 3.2 into its even and odd components and solving, yielding

$$y = a p_o / (2a^2 + \epsilon), \text{ and} \tag{3.5}$$
$$|v| = \sqrt{|p_e - (S a^2 + y^2)|}. \tag{3.6}$$

Here ϵ is a scalar noise reduction term, typically chosen to be one order of magnitude greater than the noise background in the image frames. This solution is degenerate in sign of v, where the sign ambiguity associated with the square root must be broken with the use of the second focal-plane image of known phase diversity. The map of signs of v is then computed by solving the even part of the two-equation system consisting of Eq. 3.2 and its analogue for the previous image p_1, which differs from the current image p_2 by a known phase diversity Φ_d:

$$\text{sign}(v) = \text{sign}\left(\frac{p_{2,e} - p_{1,e} - (v_d^2 + y_d^2 + 2y y_d)}{2v_d} \right), \tag{3.7}$$

where $v_d = \mathscr{F}[A\Phi_{d,e}]$, $y_d = i\mathscr{F}[A\Phi_{d,o}]$ are the Fourier transforms of the even/odd components of the known phase diversity Φ_d separating the chosen PSF images p_1 and p_2. This estimator is not used to calculate the pixel amplitudes of v because it involves the subtraction of two similar PSFs, resulting in elevated pixel-to-pixel noise with respect to Eq. 3.6. Because the first iteration of F&F must be performed without access

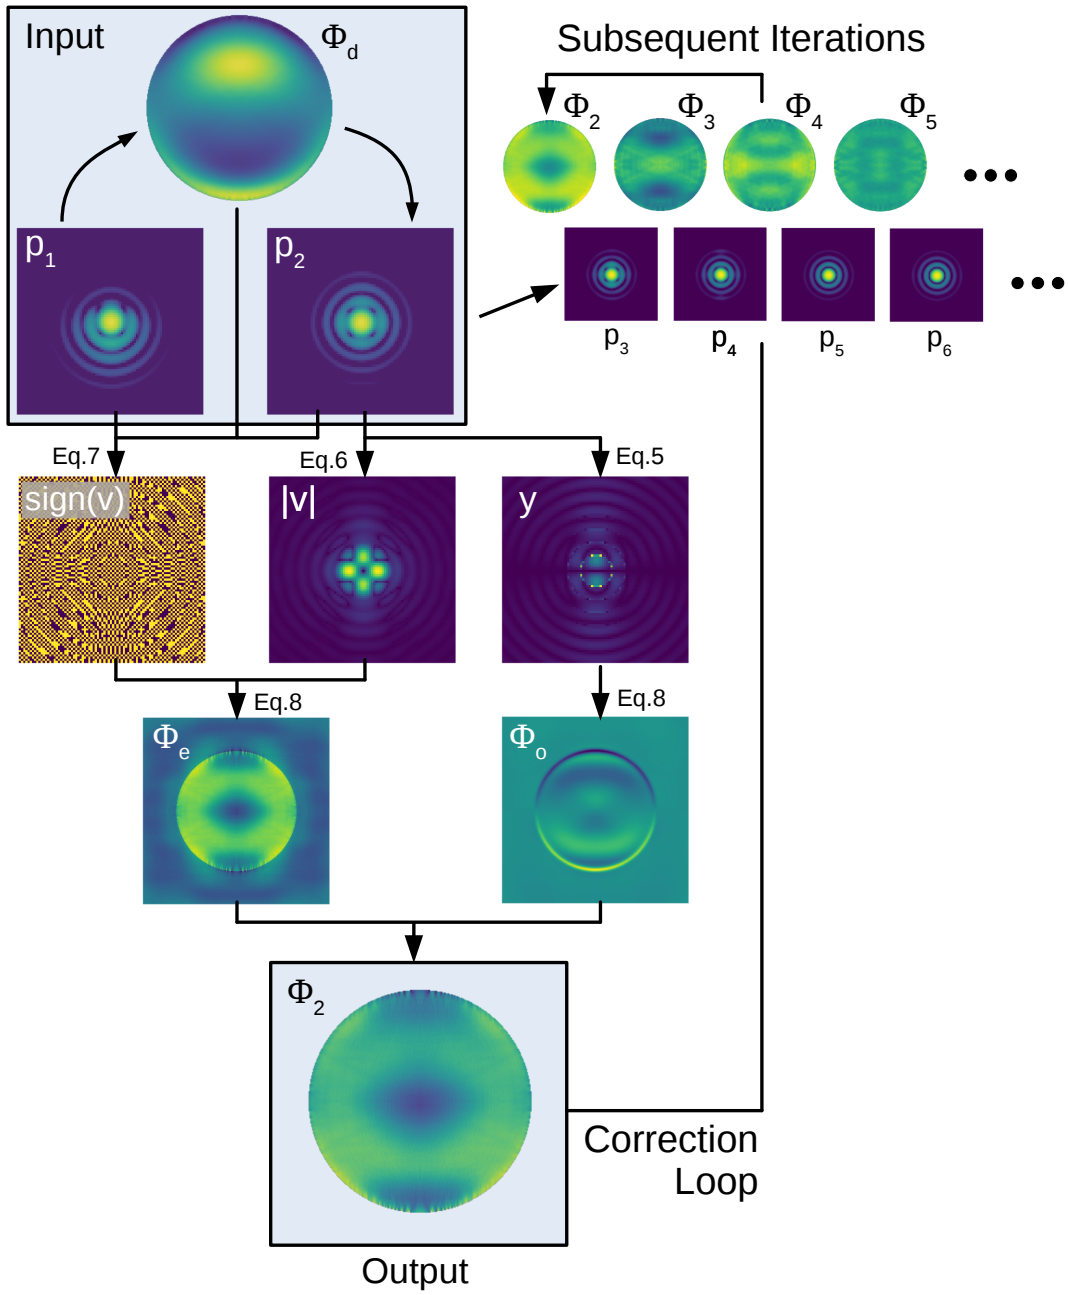


Figure 3.4: Schematic illustration of the F&F algorithm detailed in Eq. 3.5-3.8, showing the second iteration of the algorithm (i.e. the first iteration using full phase diversity information) for an example wavefront initially containing 0.25 radians root-mean-square (RMS) phase error of each of astigmatism (Z_5) and coma (Z_6) Zernike modes. The most recent PSF p_2 provides a full estimate of the odd wavefront phase Φ_o and the absolute value of the even focal-plane field $|v|$. Using p_1 and Φ_d from the previous iteration provides an estimate of the sign of v and hence completes the wavefront phase reconstruction Φ_2. This is then passed to the deformable element and becomes the phase diversity for the next iteration step.

to phase diversity information, it is convenient to take the signs of the un-aberrated field, i.e. sign(v) = sign(a), for this step only. This is found in practice to improve convergence speed over the case where all signs are assumed to be positive, and robustly results in a decrease in overall wavefront variance in the first iteration.

The final wavefront estimated is then computed in an efficient manner with the single inverse Fourier transform

$$A\Phi = \mathscr{F}^{-1}\left[\text{sign}(v)|v| - iy\right], \tag{3.8}$$

where the output may be spatially filtered or projected onto the desired correcting mode basis to reduce pixel-to-pixel noise in the final wavefront estimate.

3.3 Simulating F&F performance in SPHERE

3.3.1 Simulation parameters

The goal of this paper is to demonstrate that the F&F algorithm is capable of reliably eliminating LWE-like wavefronts given a noisy image sequence supplied by the SPHERE DTTS. The following parameters are therefore included in these simulations:

1) **Low-wind effect:** Uses the LWE phase screen shown in Fig. 3.2, which is allowed to grow linearly in amplitude from zero to a maximum discontinuity of 800 nm across the VLT spiders (equal to 1200 nm peak-valley error), over a timescale of 10 seconds.

2) **XAO-corrected atmospheric speckles:** Generated using a Kolmogorov spatial power spectrum with three wind layers of effective $r_0 = 30$ cm, at an observing wavelength of 1.5 μm. A spatial filter is applied to low frequencies to approximate a SAXO-like control radius of approximately 13 λ/D (see top panel of Fig. 3.5). The filtered phase screens are passed over the aperture with a coherence time-scale of $\tau_0 \approx 1$ ms, producing a mean Strehl ratio of S $\approx$ 0.75.

3) **Non-common-path errors:** Aberrations are generated by randomly drawing coefficients from a 100-mode Zernike basis with a simple flat spatial power spectral density. The coefficients are individually convolved with Gaussian temporal evolution profiles with a mean full width at half maximum of 1 s and the resulting time-series normalised to provide a mean RMS wavefront error of 75 nm. An example NCPE coefficient time series and wavefront are shown in the lower panel of Fig. 3.5.

4) **Detector Noise:** The simulated 512×512 pixel detector has a resolution of 3 pixels per λ/D, producing simulated DTTS-like images with an effective exposure time of 0.1 s. This includes a constant intensity background at 5×10^{-3} contrast and including photon noise, 50 e$^-$/s dark current and 20 e$^-$ readout noise sources with a final 16-bit analog-to-digital unit (ADU) conversion matched to an effective well depth of $100,000$ e$^-$ and bias offset of 500 e$^-$.

5) **Corrective element:** The final wavefront estimates are spatially filtered by projection onto a basis of the first 50 Zernike modes, as a conservative estimate of the corrective ability of the 41×41 actuator SAXO deformable mirror (DM). This filtered correction map is then assumed to be implemented accurately on

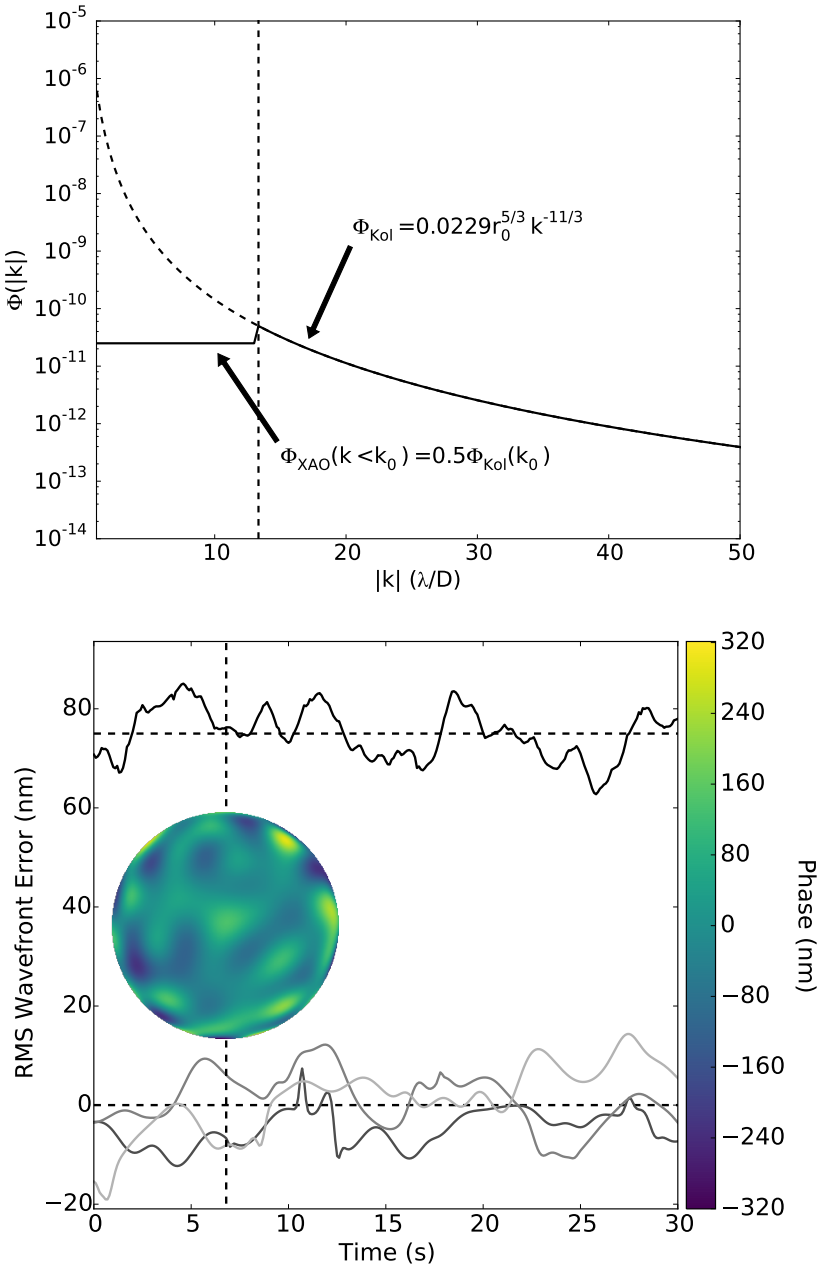


Figure 3.5: *Top:* XAO-filtered spatial power spectrum used to generate atmospheric speckle phase screens, with the x-scale converted to units of λ/D to correspond directly with focal-plane coordinates. The dashed line indicates the $k^{-11/3}$ power law of Kolmogorov turbulence. *Bottom:* Total non-common-path RMS wavefront error (top) and three examples of the temporal evolution of single Zernike mode coefficients (bottom), with a coherence time of ~1 s. Inset: Representative NCPE phase map comprising all 100 Zernike modes, evaluated at $t = 7$ s.

the DM at a cadence of 10 Hz to match the simulated DTTS exposures, with zero frames of lag.

In order to test the limiting performance of the algorithm, these parameters have been deliberately adjusted to provide a conservative estimate of SPHERE performance: the correction radius, H-band Strehl ratio and residual speckle noise inside the control radius have all been deliberately degraded with respect to reported SAXO performance (Fusco et al. 2014; Sauvage et al. 2015). The aim of this is to maximise the potential impact of noise sources (in particular the presence of an incoherent speckle background) on F&F, such that a successful verification in simulation provides greater confidence in the on-sky performance of the algorithm under less strenuous conditions. Examples of the simulated DTTS images generated using the above parameters are shown in Fig. 3.6. Here the control radius due to the XAO power spectrum can be seen at $\sim 13\lambda/D$, although this and the majority of diffracted intensity structure beyond the first Airy ring are hidden beneath the detector noise floor in final images. This is comparable to the observed image quality presented in Fig. 3.1 whilst still containing all the major factors against which the F&F algorithm must be tested in order to evaluate its potential for on-sky operation.

3.3.2 Simulation results

Figure 3.7 shows the results of a demonstrative closed-loop F&F simulation with the parameters specified in the previous section. The total duration of 30 seconds spans $3\times10^4\tau_0$ and 30 NCPE coherence time-scales, and as such fully samples the behaviour of the secondary aberration sources. The first 5 seconds involve only these two components, after which a LWE is allowed to slowly evolve over the course of 10 seconds, reaching the maximum 800 nm phase discontinuity amplitude at the 15 s half-way point in the simulation. In order to simulate a scenario in which F&F is activated once the LWE reaches a critical aberration threshold, correction begins at the 10 second mark (blue curve), when the LWE amplitude reaches 400 nm. The final 15 seconds after the LWE has finished evolving are then used to produce long integrations of the simulated post-correction DTTS images shown in the left panels of Fig. 3.7, which are used to evaluate the performance of the algorithm with respect to the un-corrected case.

The final LWE discontinuity amplitude of 800 nm represents an extreme case exceeding the maximum effect observed on-sky, resulting in a final Strehl of less than 15% once residual atmospheric speckles and NCPEs are included. Nonetheless, it can be seen that the F&F algorithm is able to effectively remove the characteristic 'Mickey Mouse ears' by compensating only the first 50 lowest-order Zernike modes, recovering a diffraction limited core and first Airy ring along with a significant gain in image Strehl ratio. The convergence of the algorithm is also rapid once initialised; for a closed-loop gain of 0.3 the Strehl ratio increases by 20% within the first five iterations. It can be seen however that image quality then continues to degrade slowly up to 15 s simulation time as the strength of the LWE continues to increase, with the final converged wavefront Strehl ratio stabilising at approximately 60%, 16% below its initial pre-LWE value. This loss of PSF quality can be explained entirely by the high-order components of the injected LWE wavefront which are not addressed when correcting only the first 50 Zernike modes. This is shown in Fig. 3.8, where the majority of the residual aberration is seen in the third panel to be localised around the VLT spiders,

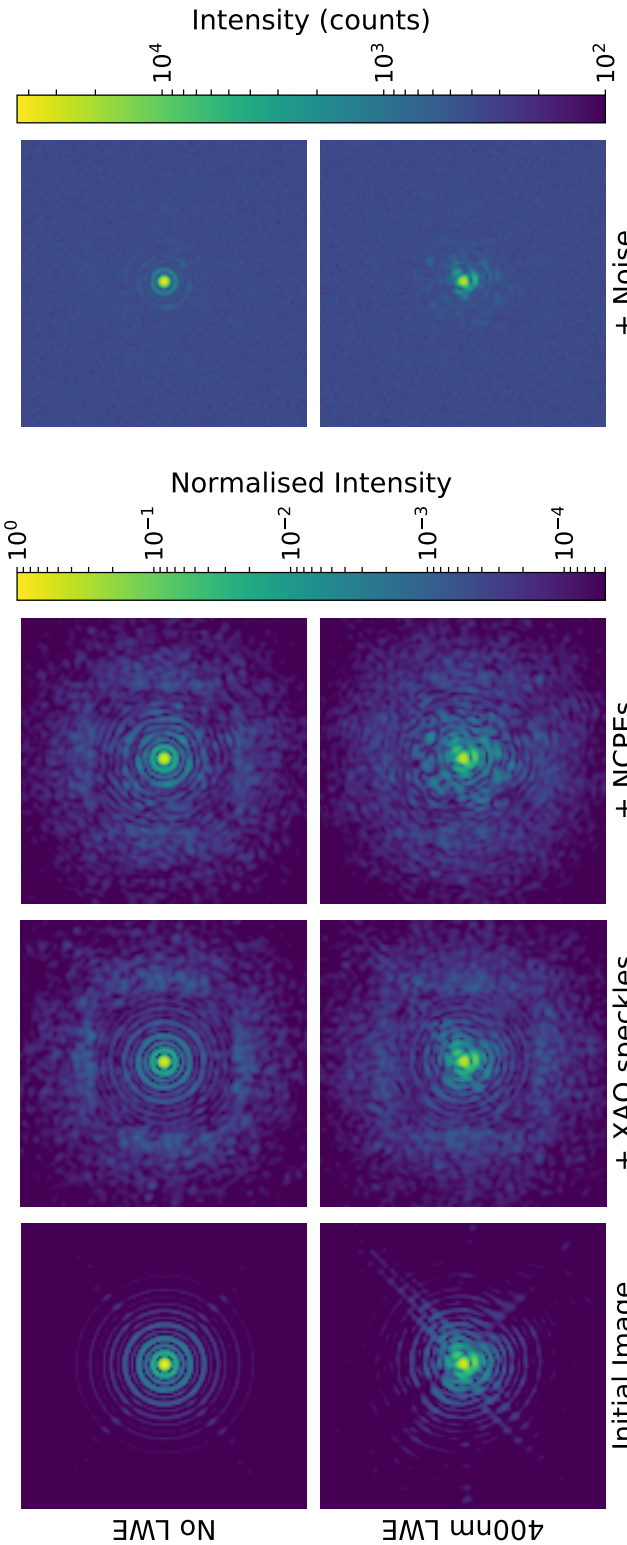


Figure 3.6: Construction of simulated DTTS images with (lower panels) and without (upper panels) a 400 nm amplitude LWE. The first column shows the ideal stellar images with only the LWE, the second column includes XAO-corrected atmospheric residuals, the third column adds NCPEs and the fourth column shows the final noisy DTTS images, re-scaled to simulated ADU counts.

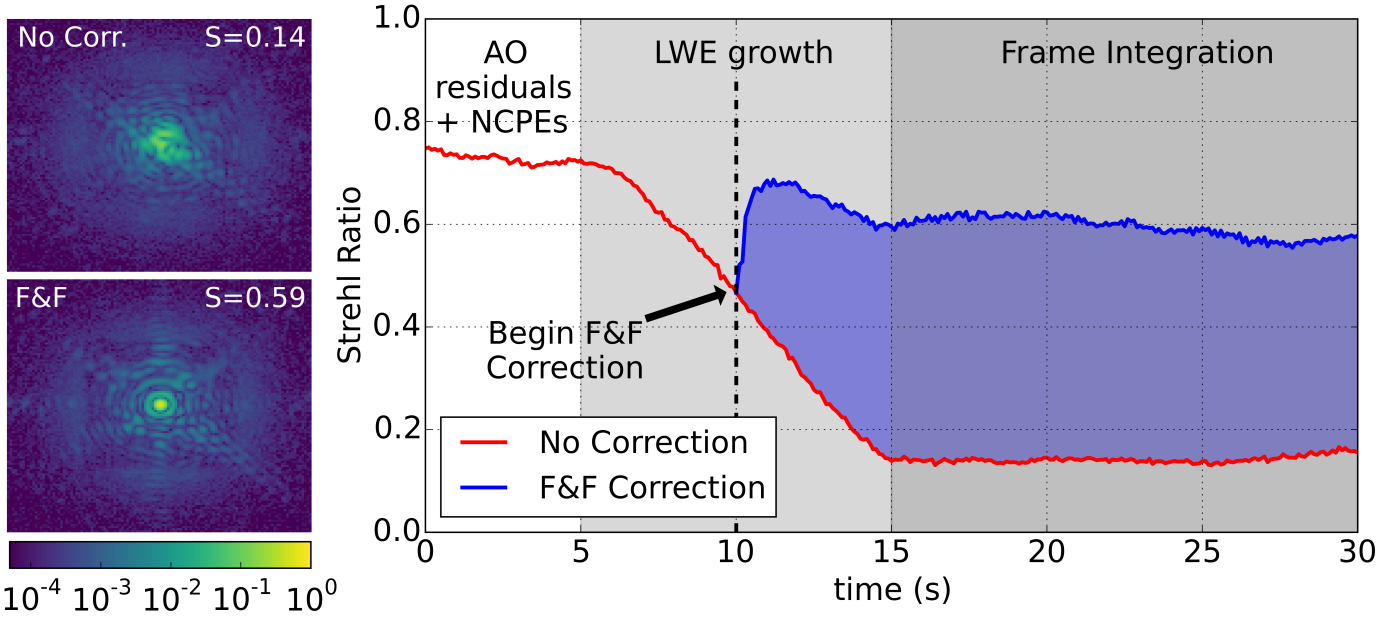


Figure 3.7: Closed-loop performance of F&F for the LWE wavefront shown in Fig. 3.2 with simulated SPHERE/DTTS observing parameters described in Sec. 3.3.1. *Left:* 15s integrations of the final pre-coronagraphic DTTS PSFs with and without F&F correction (bottom and top respectively). *Right:* Strehl ratio of individual 0.1s DTTS exposures as a function of simulation time, showing significant improvement made by F&F (blue) over the un-corrected wavefront (red).

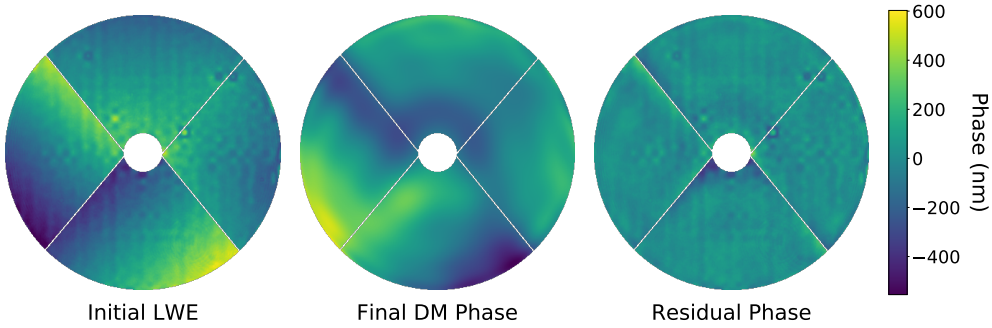


Figure 3.8: Wavefront phase maps illustrating the corrective ability of the first 50 Zernike modes, in a closed-loop simulation involving only LWE aberration. ***Left:*** Input LWE wavefront. ***Middle:*** Final corrective phase on deformable mirror which matches well the tip/tilt components of the wavefront. ***Right:*** Residual wavefront error after correction, seen to be dominated by high-frequency errors around the spiders, and additional high-order waffle and faulty actuator patterns present in the initial wavefront.

in addition to a weak waffle pattern and the influence functions of faulty actuators present in the initial wavefront.

3.4 Discussion

In these simulations F&F was found to have an effective closed-loop dynamic range capable of correcting an existing LWE aberration with discontinuity amplitude of 600 nm across the spiders, corresponding to a wavefront phase variance of $\sigma_\phi^2 = 1$ radian. This is twice the typical weak-phase approximation criterion $\sigma_\phi^2 \leqslant 0.5$ radians (Polo et al. 2013; Gonsalves 2001) due mostly to the inclusion of the second-order Strehl correction factor in Eq. 3.2. This increased dynamic range is in itself sufficient to cope with the majority of the fully-evolved cases of the LWE seen on-sky at SPHERE. Any further convergence issues may be simply avoided by ensuring F&F correction is activated as soon the LWE passes a critical detection threshold, in the same manner as Fig. 3.7. Alternatively, it is hoped that F&F will prove sufficiently stable to run continuously in the nominal SPHERE observing mode without manual intervention, such that it is possible to prevent the LWE from arising in the first place.

Spatial filtering of the wavefront estimates produced by F&F is necessary in order to remove high-frequency noise and artefacts of pupil asymmetry generated by the algorithm, in addition to providing an approximation of the corrective order of the chosen deformable mirror. The majority of the algorithm works analytically at the pixel level without any assumption of wavefront geometry or spatial order; it is therefore prone to propagating pixel-to-pixel noise in the input DTTS images through as high-frequency noise the final wavefront estimates. The re-formulating of Eq. 3.5 to include ϵ following the approach of Gonsalves (2001) provides a level of noise reduction for the odd wavefront, however the solution of v is significantly more prone to pixel-to-pixel noise due to the independent estimation of signs of the focal-plane field using Eq. 3.7. Filtering has been found to be effective at removing these noise sources, although other approaches such as amplitude apodisation of the input images has also been seen to reduce the severity of noise propagation (Korkiakoski et al. 2014). F&F operates under the implicit assumption of a completely even aperture function

A, hence any asymmetric structures such as VLT spiders result in high-frequency, low-amplitude phase 'ghosts' located at the conjugate positions of the amplitude masking element in the raw F&F output wavefront estimates. While such effects were not found to grow significantly in amplitude in closed-loop in cases where no spatial filtering techniques were used, they nonetheless degrade wavefront quality by a small amount and are best removed from final estimates.

It was also an initial concern that SPHERE's main 41×41 actuator deformable mirror would be unable to cope with the sharp wavefront discontinuities present in all LWE wavefronts, hence the use of a 50-mode low-order basis throughout this paper to remove the high-frequency components of F&F estimates. As presented by Sauvage et al. 2016 (Sauvage et al. 2016) however, this may in reality be highly conservative as SAXO is expected be able to accurately implement a custom projection basis describing independent piston, tip and tilt errors across each VLT pupil segment and, if desired, also include a secondary smooth basis such as Zernike modes for independent control of NCPEs. A limited correction basis tailored to the LWE is also expected to help maintain algorithm stability by minimising the number of degrees of freedom available for the secondary loop to control. It is still worth noting however that the majority of the LWE may be eliminated by controlling only the first 50 Zernike modes as shown in Fig. 3.8, should a solution implementing only low spatial frequencies prove to be necessary for practical implementation in SPHERE.

Additional outstanding practical considerations not addressed by these simulations include:

1) **XAO loop conflicts:** The potential for conflict of F&F commands with the existing adaptive optics operations loop may be avoided by applying these corrections as reference offsets to the Shack-Hartmann sensor in the same manner as is proposed for NCPEs in Sauvage et al. 2011 (Sauvage et al. 2011). Further investigation is required to verify that the additional stroke used in compensating the LWE does not significantly degrade the performance of the main atmospheric correction loop.

2) **DTTS windowing modes:** Although the Hawaii I detector of the DTTS consists of an array of 1024×1024 pixels, it is usually run in a highly windowed mode to increase readout speed which cuts off the majority of the high-order PSF structure. However, it is expected that the F&F algorithm is able to cope with this reduction of field of view through appropriate normalisation, given that the majority of information beyond $\sim 3\lambda/D$ is dominated by pixel-to-pixel noise in the simulations presented in this paper.

3) **Dead Actuators:** SPHERE currently has a number of known dead or faulty actuators, which modify the phase response of the mirror and are typically masked out in amplitude during science operation. In order to provide the correct phase diversity to F&F, it is important to ensure an up-to-date DM model is used for estimating Φ_d.

3.5 Conclusions

The 'Fast and Furious' (F&F) algorithm is a promising solution for the real-time detection and correction of the low-wind effect (LWE), which is also in principle capa-

ble of compensating quasi-static non-common-path errors in part of the near-infrared optical path for the IRDIS and IFS instruments. This is achieved using no additional hardware and functions entirely within the nominal observation mode of the instrument by using non-coronagraphic images supplied by the existing DTTS imager present within the NIR instrument arm. Realistic closed-loop simulations show that F&F can robustly eliminate the characteristic 'Mickey Mouse ears' associated with a LWE-affected PSF within five iterations at 10 Hz cadence, providing significant recovery of Strehl and hence improvement in coronagraph performance under the best observing conditions. Future work will focus on validating F&F on a representative experimental optical testbench and, if successful, its implementation within the SAXO extreme adaptive optics system of SPHERE. Such algorithms may also prove useful in the long term for overcoming similar (or as-yet unforeseen) issues with XAO systems on ELT-class telescopes, where additional complications from mirror segmentation and associated phasing errors may lead to similar effects.

Chapter 4

Laboratory verification of 'Fast & Furious' phase diversity

Towards controlling the low-wind effect in the SPHERE instrument

Michael J. Wilby, Christoph U. Keller, Jean-François Sauvage, Kjetil Dohlen, Thierry Fusco, David Mouillet, and Jean-Luc Beuzit
A&A 615, A34 (2018)

Abstract

The low-wind effect (LWE) refers to a characteristic set of quasi-static wavefront aberrations seen consistently by the SPHERE instrument when dome-level wind speeds drop below 3 ms^{-1}. The LWE produces bright low-order speckles in the stellar PSF, which severely limit the contrast performance of SPHERE under otherwise optimal observing conditions. In this paper we propose the Fast & Furious (F&F) phase diversity algorithm as a viable software-only solution for real-time LWE compensation, which would utilise image sequences from the SPHERE differential tip-tilt sensor (DTTS) and apply corrections via reference slope offsets on the AO system's Shack-Hartmann wavefront sensor. We evaluated the closed-loop performance of F&F on the MITHIC high-contrast test-bench, under conditions emulating LWE-affected DTTS images. These results were contrasted with predictive simulations for a variety of convergence tests, in order to assess the expected performance of an on-sky implementation of F&F in SPHERE. The algorithm was found to be capable of returning LWE-affected images to Strehl ratios of greater than 90% within five iterations, for all appropriate laboratory test cases. These results are highly representative of predictive simulations, and demonstrate stability of the algorithm against a wide range of factors including low image signal-to-noise ratio (S/N), small image field of view, and amplitude errors. It was also found in simulation that closed-loop stability can be preserved down to image S/N as low as five while still improving overall wavefront quality, allowing for reliable operation even on faint targets. The Fast & Furious algorithm is an extremely promising solution for real-time compensation of the LWE, which can operate simultaneously with science observations and may be implemented in SPHERE without requiring additional hardware. The robustness and relatively large effective dynamic range of F&F also make it suitable for general wavefront optimisation applications, including the co-phasing of segmented ELT-class telescopes.

4.1 Introduction

The Spectro-Polarimetric High-Contrast Exoplanet REsearch instrument (SPHERE, Beuzit et al. 2008), is a second-generation high-contrast imaging instrument for the Very Large Telescope (VLT), which finished its commissioning phase in 2014. Since then it has been routinely delivering unprecedented science results in the fields of dual-band imaging, differential polarimetry, and integral field spectroscopy of directly-imaged protoplanetary disks and young exoplanets (e.g. Vigan et al. 2016a; Maire et al. 2016; Zurlo et al. 2016; Bonnefoy et al. 2016; de Boer et al. 2016; Ginski et al. 2016). The extreme adaptive optics (XAO) system of SPHERE, SAXO (Petit et al. 2016; Fusco et al. 2014, 2016), is capable of routinely achieving Strehl ratios of 90% in the H-band. When this performance is combined with coronagraphic observation modes and optimised reduction pipelines, it is possible to achieve 5σ planet-star companion detectability ratios of better than 10^{-6} beyond angular separations of 375 mas (Zurlo et al. 2016).

However, the instrument performance and ultimately the science yield of SPHERE is currently limited under the best observing conditions by the so-called low-wind effect (LWE). This effect refers to a systematic degradation of the image quality of all three SPHERE detector arms (IRDIS, IFS, and ZIMPOL), which occurs when the wind speed at the altitude of the VLT dome drops below approximately $3\ \mathrm{ms}^{-1}$ (Sauvage et al. 2016). The characteristic LWE wavefronts consist of independent piston-tip-tilt (PTT) phase errors across one or more of the VLT pupil segments, and have been observed to reach up to 800 nm peak-to-valley error (PVE) on-sky as measured by a prototype of the Zernike wavefront sensor ZELDA (N'Diaye et al. 2014; N'Diaye et al. 2016). As shown in Fig. 4.1, this leads to a significant degradation of the imaging point-spread function (PSF) by creating multiple bright side-lobes at the location of the first Airy ring and increasing the amount of diffraction structure around the secondary mirror (M2) support spiders. This is an issue for both the coronagraphic and non-coronagraphic high-contrast observing modes of SPHERE, due firstly to increased photon noise and a lower Strehl ratio of off-axis companion sources. In reality the LWE is also a quasi-static phenomenon, and generates significant additional speckle noise on timescales and angular separations particularly detrimental to reference PSF subtraction and other high-contrast data reduction techniques, such as the angular differential imaging (ADI) and principal component analysis (PCA) classes of algorithm (e.g. Marois et al. 2006; Lafrenière et al. 2007; Soummer et al. 2012; Amara & Quanz 2012).

This high-amplitude, quasi-static LWE may be considered a specific example of a more general 'island effect' (N'Diaye et al. 2018), which encompasses all differential PTT aberrations associated with pupil segmentation irrespective of underlying cause or temporal behaviour. Since examples of island effect behaviour are now also being reported intermittently at the SCExAO (Jovanovic et al. 2015) and GPI instruments (Macintosh et al. 2008, V. Bailey, private communication, 2016), solutions developed for the LWE in SPHERE may well be applicable to similar issues faced by other instruments. This is expected to be especially important for the upcoming extremely large telescopes (ELTs), which will feature significantly more complex pupil geometries and may be correspondingly prone to these effects.

Figure 4.2 illustrates the degradation of raw contrast performance for the SPHERE apodised pupil lyot coronagraph (APLC) for two example LWE cases drawn

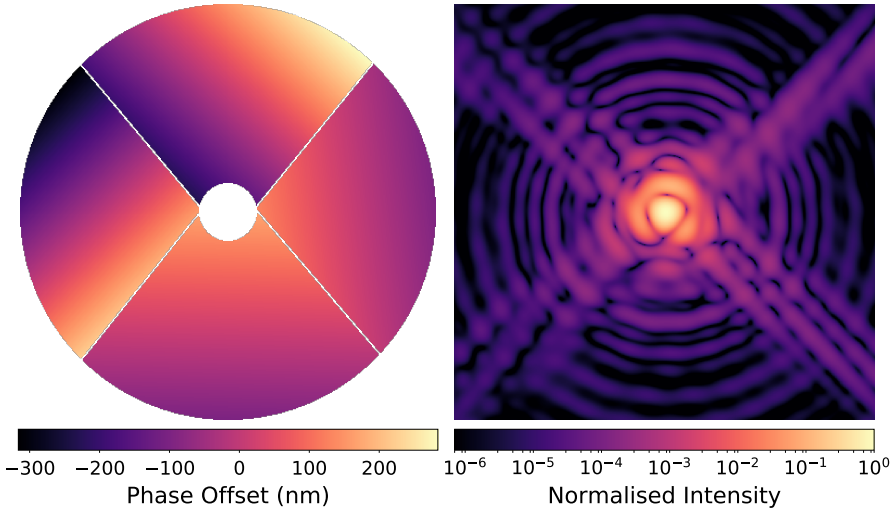


Figure 4.1: Typical example of the LWE phenomenon, based on on-sky measurements made with the ZELDA wavefront sensor. *Left:* Parametrised PTT wavefront model based on a single ZELDA phase measurement. *Right:* The aberrated PSF corresponding to this LWE wavefront map, simulated at a wavelength of 1.536 μm. The PSF displays three notable side-lobes at the location of the first Airy ring (2.5 λ/D), which correspond to the differential tip-tilt components seen across individual pupil segments in the aberrating phase map.

from the SPHERE user manual (ESO 2016), using a Fourier propagation model of the three-plane coronagraph system described in Guerri et al. (2011). It can be seen that diffraction-limited simulations (black curve, main panel) predict a raw contrast ratio of significantly better than 10^{-4} between 2-4 λ/D, however this is not representative of real systems containing sources of non-common path error (NCPE). In order to provide a more realistic performance estimate, each PSF in Fig. 4.2 (and the corresponding green, blue, and red curves in the lower panel) includes the incoherent average of 100 random realisations of low-order, low-amplitude wavefront aberrations. These low-order wavefronts are created by drawing random Zernike mode coefficients, with the resulting phase maps then spatially filtered in the Fourier domain to have a $1/f^2$ decreasing spatial power spectrum often used to model NCPEs (Sauvage et al. 2007; Lamb et al. 2016). These are then scaled to have a 30 nm root-mean-square (RMS) error, representative of the calibration accuracy achieved in SPHERE after baseline NCPE calibration routines (Fusco et al. 2014). From this it can be seen that for a typical LWE amplitude of 600 nm PVE there is an increase in off-axis transmission of the central source of an order of magnitude with respect to the NCPE-limited case between 2-4 λ/D. This alone would result in a factor of three increase in photon noise in the final reduced image at these angular separations, notwithstanding the inevitable impact of speckle variability due to a quasi-static LWE on the ultimate achievable contrast.

The current working hypothesis is that the LWE is caused by slow laminar airflow across the deep but narrow VLT M2 support spiders, which allows time for significant thermal exchange to occur (Sauvage et al. 2015). This results in sharp temperature changes and hence variations in the optical path depth of the air column across the width of each spider, thus generating discontinuities in phase such as that illustrated in the left-hand panel of Fig. 4.1. This optical path depth hypothesis is sup-

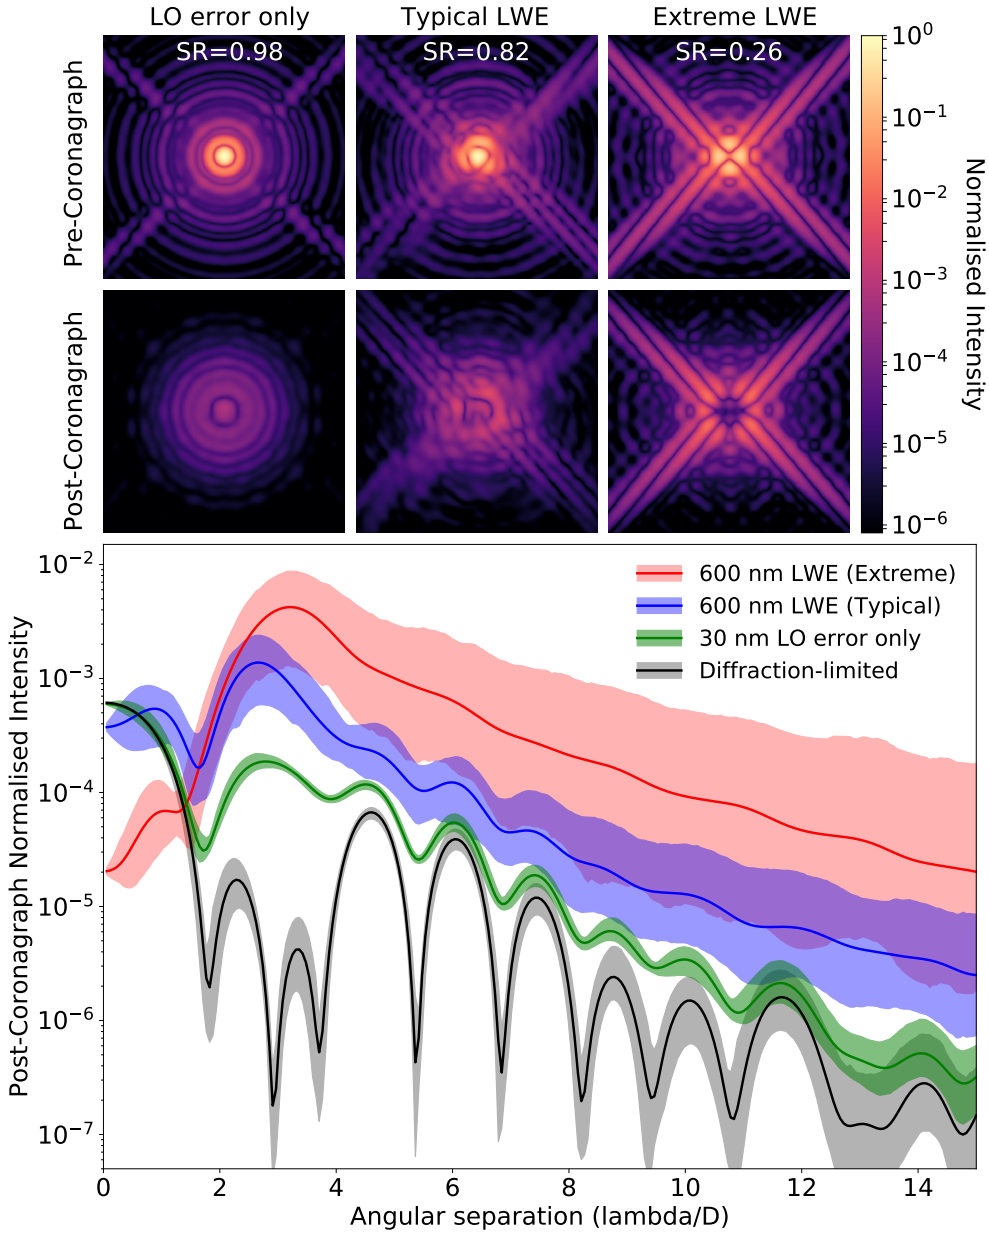


Figure 4.2: Simulated performance of the SPHERE APLC in the presence of various LWE models and low-order wavefront error, for an observing wavelength of 1.536 μm. *Left image column:* LWE-free PSFs containing only 30 nm RMS of low-order aberrations with a $1/f^2$ spatial power spectrum, incoherently averaged over 100 random realisations. *Centre column:* a typical three-lobed LWE model (600 nm PVE) identical to Fig. 4.1. *Right column:* an extreme LWE with wavefront similar to a four-quadrant phase mask pattern (600 nm piston on two opposite VLT pupil quadrants). *Top image row:* non-coronagraphic PSFs, including image Strehl ratio with respect to the diffraction-limited case. *Bottom image row:* Corresponding on-axis coronagraphic PSFs. *Main figure panel:* Radial average contrast curves of each post-APLC PSF, with shading denoting the 1σ upper and lower bounds on azimuthal variation. Diffraction-limited performance in the absence of all aberration is shown by the black curve for comparison purposes.

ported by ESO simulations (Sauvage et al. 2016), which can reproduce the strength and overall morphology of the various wavefronts associated with the effect under realistic dome conditions. These characteristic wavefronts are however not seen in either SAXO Shack-Hartmann wavefront sensor (SH-WFS) data or in the deformable mirror (DM) actuator voltages during on-sky operation when the effect is present. This implies that the AO system is at best blind to this class of wavefront error, and at worst may be partially responsible for creating the effect due to unreliable sensing of phase discontinuities near the spiders. For this reason a complete understanding of the LWE, and other instances of the island effect, remains an active area of investigation with potentially significant implications for the design of future high-contrast imaging instruments.

Attempts to eliminate the LWE phenomenon via dynamic control of telescope dome conditions (including increased ventilation, temperature control, and telescope pointing with respect to the prevailing wind direction) or manipulation of the AO closed-loop parameters have so far proved unsuccessful in reducing the strength or occurrence rates of the effect. Current efforts are ongoing to improve the thermal properties of the spiders by directly applying coatings with improved near-infrared (NIR) emissivity (M. Kasper, private communication, 2016), which if successful would allow the structure to better remain in thermal equilibrium with the surrounding air and thereby prevent phase discontinuities from arising. Another approach is to directly sense and compensate the LWE wavefront in real-time by introducing an additional wavefront correction loop into the SPHERE instrument: the ZELDA wavefront sensor has been shown to be an accurate truth sensor for the LWE during trials at the VLT in 2016, however until an upgrade to the SPHERE instrument can be performed ZELDA must convert the IRDIS focal plane into a pupil-plane sensor, preventing it from being used simultaneously with NIR science observations.

This paper proposes an immediately implementable solution to directly sense and compensate the LWE, by using phase diversity techniques to turn the existing differential tip-tilt sensor (DTTS) camera, used for centring the stellar PSF on the NIR coronagraph (Baudoz et al. 2010), into a focal-plane wavefront sensor. This is an attractive solution as it requires no additional hardware or modification to the operation of existing SPHERE subsystems, as the correction commands may be applied as reference slope offsets to the main SAXO SH-WFS and therefore should be able to operate in parallel with the atmospheric XAO loop without conflict. For this task we propose the Fast & Furious (F&F) modified sequential phase diversity algorithm (Keller et al. 2012; Korkiakoski et al. 2014), so named because it uses a simplified model of the imaging system to obtain an analytical, computationally efficient phase reconstruction procedure (see Sect. 4.2.1 for details). F&F is also capable of using its own phase correction update cycle to provide the necessary phase diversity for complete focal-plane wavefront retrieval. This is in contrast to the majority of phase diversity approaches (Gonsalves 2001; Sauvage et al. 2007; Lamb et al. 2016), which require the periodic application of large controlled probe phases in order to reconstruct the aberrating wavefront. By eliminating this requirement, F&F has the major advantage of being able to run continuously in closed-loop without degrading or interrupting science observations, enabling continuous real-time wavefront control.

Fast & Furious has been successfully tested in proof-of-concept simulations under imaging conditions emulating the SPHERE DTTS (Wilby et al. 2016b) and in a general laboratory environment not specific to the LWE (Korkiakoski et al. 2012; Korkiakoski

et al. 2014). While these preliminary simulations indicated that the algorithm should be robust against operating with the DTTS camera, before this solution may be implemented on-sky it is essential to verify that this performance is reflected in an appropriate laboratory environment. In this paper we therefore present the results of LWE-specific lab testing using the Marseille Imaging Testbed for HIgh Contrast (MITHIC, Vigan et al. 2016b), located at the Laboratoire d'Astrophysique de Marseille (LAM), the results of which are combined with improved closed-loop simulations to evaluate the potential performance of F&F on-sky in SPHERE at the VLT.

This paper is divided into the following sections: Sect. 4.2 outlines the principle of the F&F algorithm and presents details of both the MITHIC test-bench environment and supporting simulation tools. Sect. 4.3 presents the main laboratory closed-loop results and compares these to simulated performance predictions, and investigates the stability of F&F at extremely low image signal-to-noise ratio (S/N). The limiting factors and lessons learned from this investigation are discussed in Sect. 4.4, and final conclusions are drawn in Sect. 4.5.

4.2 Methodology

4.2.1 The Fast & Furious algorithm

The F&F algorithm refers to a sequential phase diversity technique based on Gonsalves (2002), which has been extended to improve dynamic range and stability. It is capable of performing real-time wavefront phase retrieval when provided with a time-series of non-coronagraphic, narrowband focal-plane images and knowledge of the frame-to-frame phase commands applied by deformable elements in the system. Wavefront reconstruction is achieved by solving an analytical approximation to the stellar PSF in terms of the even and odd focal-plane intensity distributions, corresponding to the Fourier symmetries of the wavefront to be sensed. Phase diversity information is used only to break a sign ambiguity associated with calculating the even wavefront, which is most effectively provided by the phase correction command from the preceding iteration of F&F. This has the major advantage that the sequential phase diversity process continually improves wavefront quality, allowing it to operate in parallel with continuous science observations. The algorithm is also highly computationally efficient, requiring only a single complex Fourier transform per iteration plus a small number of linear operations on image data. The correction cadence of F&F will therefore be limited by the imaging camera readout frequency in most practical applications, whereas other phase diversity approaches are limited by the (significantly lower) frequency of phase probe injection. Unlike classical phase diversity however, F&F is not capable of performing one-shot phase retrieval and must be operated in closed-loop for full wavefront compensation. Despite this it is always possible to reconstruct the odd wavefront component from any single PSF image, and with an appropriate choice of initial conditions it is almost always possible to achieve a systematic improvement in wavefront quality even on the first iteration.

A full derivation and analysis of the numerical properties of the F&F algorithm can be found in Keller et al. (2012) and Korkiakoski et al. (2014); a summary of the key equations necessary to implement the algorithm is presented here for reference. In the following description capitalised variables are used to denote two-dimensional pupil-plane quantities and lower-case variables denote two-dimensional focal-plane

quantities, unless otherwise noted. The PSF $p = |\mathscr{F}[Ae^{i\Phi}]|^2$ of an aberrated stellar image may be Taylor expanded to second order as a function of the even and odd focal-plane electric fields as

$$p \approx S a^2 + 2a(ia * \phi_o) + (ia * \phi_o)^2 + (a * \phi_e)^2, \tag{4.1}$$

where $a = \mathscr{F}[A]$ is the complex Fourier transform of the telescope aperture function A, which is assumed to be real and symmetric, while $\phi_o = \mathscr{F}[\Phi_o]$ and $\phi_e = \mathscr{F}[\Phi_e]$ are the complex Fourier transforms of the odd and even components of the wavefront phase map Φ, with $\Phi_o = -\Phi_o^{\mathrm{T}}$ and $\Phi_e = \Phi_e^{\mathrm{T}}$, such that $\Phi = \Phi_o + \Phi_e$. The scalar normalisation factor $S = (1 - \sigma_\phi^2)$ is approximately equal to the Strehl ratio of the most recent image, where σ_ϕ^2 is the total wavefront variance: this is effectively the first-order Taylor expansion of the Maréchal approximation (Roberts et al. 2004). Here $\mathscr{F}$ is the Fourier transform operator, and $*$ is the convolution operator. This formulation takes advantage of the enforced symmetry properties of all pupil-plane quantities to simplify the expressions, and to remain consistent with the symmetry properties of the complex Fourier transform.

The expression above may be more conveniently expressed by defining the odd and even focal-plane fields as the two real quantities

$$y = i\mathscr{F}[A\Phi_o] = (ia * \phi_o), \text{ and} \tag{4.2}$$
$$v = \mathscr{F}[A\Phi_e] = (a * \phi_e), \tag{4.3}$$

leading to analytical solutions for the complete odd field and the absolute value of the even field, by separating Eq. 4.1 according to symmetry and solving. This yields

$$y = a p_o / (2a^2 + \epsilon), \text{ and} \tag{4.4}$$
$$|v| = \sqrt{|p_e - (S a^2 + y^2)|}, \tag{4.5}$$

where p_o and p_e are the odd and even components of p by direct analogy with Φ_o and Φ_e. Here the scalar ϵ parameter is introduced as a method of regularisation for pixels where a tends towards zero, which is typically set to a factor of ten above the noise threshold in the image sequence.

In order to estimate the signs of the even focal-plane field, it is necessary to introduce a second PSF image which differs from the first by a known phase offset Φ_d, such that

$$p_1 = S a^2 + \quad 2ay \quad + \quad y^2 \quad + \quad v^2, \text{ and} \tag{4.6}$$
$$p_2 = S a^2 + 2a(y + y_d) + (y + y_d)^2 + (v + v_d)^2, \tag{4.7}$$

where v_d and y_d correspond to this additional phase diversity between frames by analogy with Eqs. 4.2 & 4.3. Solving for v yields an independent expression in terms of the even components of each PSF, p_{1e} and p_{2e}, however this solution is extremely prone to noise due to the subtraction of two similar PSFs. It is therefore significantly more robust to take only the signs of this estimate,

$$\mathrm{sign}(v) = \mathrm{sign}\left(\frac{p_{2e} - p_{1e} - (v_d^2 + y_d^2 + 2y y_d)}{2v_d} \right), \tag{4.8}$$

and combine these with the magnitude $|v|$ computed from Eq. 4.5.

The final estimate of the wavefront phase may then be reconstructed via a single complex inverse Fourier transform, by taking into account the symmetries of the odd and even focal-plane fields to give

$$A\Phi = \mathscr{F}^{-1}[\text{sign}(v)|v| - iy]. \tag{4.9}$$

This estimate may then be applied using a wavefront correcting element, which in turn becomes the new phase diversity Φ_d for the next iteration. Since $\text{sign}(v)$ is undefined on the first iteration where phase diversity information is not available, it is optimal in the case of small wavefront aberrations to choose $\text{sign}(v_0) = \text{sign}(\mathscr{F}[A])$.

Key assumptions made implicitly by F&F include the use of a symmetric telescope pupil function, the presence of phase-only aberrations and a monochromatic light source. An improved version of the algorithm dubbed FF-GS, which includes Gerchberg-Saxton style iterative steps, has been developed (Korkiakoski et al. 2014) which can overcome these first two limitations by also enabling amplitude retrieval. It was not however found to be necessary to implement FF-GS in this work, due to the near-symmetric SPHERE pupil generating only small systematic reconstruction errors which may be removed by spatial filtering, therefore not warranting the increased complexity and lower stability of FF-GS compared to F&F alone.

4.2.2 Implementing F&F in SPHERE-like simulations

The first step towards using F&F as a LWE-compensator for SPHERE is to simulate as closely as possible the observing conditions of the DTTS camera, and verify that the algorithm is both efficient at eliminating LWE wavefronts and stable during continuous operation in the absence of the effect. The DTTS Hawaii I camera is capable of operating at high frame-rates (1 Hz - 1 kHz) with only a 32×32 pixel field of view at 3.5 pixels per λ/D sampling, with images stacked to provide tip-tilt correction at a cadence of 1 Hz. This is operated at a wavelength of 1.536 μm (H-band) with a 3% bandwidth and is situated behind a 2% beam-splitter to avoid unnecessary science throughput losses, hence DTTS images normally have low S/N. F&F must therefore first and foremost be robust against dominant detector noises sources and a limited image size. Other considerations include the presence of amplitude variations, most notably the presence of the SPHERE NIR coronagraph amplitude apodiser (located upstream of the DTTS beam-splitter) and pupil rotation, but also errors in representing wavefront control commands due to fitting errors associated with a finite DM resolution, as well as systematic errors in applying reference offsets on an existing AO loop.

A dedicated python simulation package has been developed for the purpose of validating F&F under SPHERE-like conditions, which is also used throughout this work. This code is capable of generating realistic DTTS-like image sequences using an XAO-corrected turbulence phase-screen simulator and quasi-static NCPE model with appropriate coherence timescale and spatial power spectrum, with photon and detector noise sources added to achieve the desired S/N. A comprehensive overview of the main code features is provided in Wilby et al. (2016b), which also reported initial simulation results demonstrating that F&F should be capable of providing significant and robust improvements in wavefront quality for the case of the LWE. The key steps used in this implementation F&F are described below, which also apply to the code used for MITHIC laboratory testing (see Sect. 4.2.3).

As part of the preliminary data reduction step, a windowing function was applied to images in order to suppress pixel-to-pixel noise by removing the high spatial frequencies containing no retrievable wavefront information. This took the form of a radial sigmoid low-pass spatial filter, with a radial cut-off determined by image S/N: this step is described further in Sect. 4.3.2. Two consecutive images p_1 and p_2 were then used along with the appropriate phase diversity command Φ_d to calculate the F&F phase reconstruction $A\Phi$, as described by Eqs. 4.4, 4.5, 4.8 & 4.9. The maximum spatial frequency which can be controlled by F&F is ultimately set by the field-of-view (FOV) of the input images, and the Fourier transform operation must be constructed so as to ensure that the reference field a has both the same focal-plane sampling as the input images and at least the same FOV. This may be achieved by using a discrete fast Fourier transform (FFT) operation with an appropriately zero-padded aperture function A, but to fully optimise the algorithm a non-FFT based method such as the semi-analytical approach described in Soummer et al. (2007b) should be used. This allows for an arbitrary focal-plane sampling and FOV whilst computing $A\Phi$ at an appropriate resolution for the deformable element being used, thereby maximising computational speed.

In order to further suppress pixel-to-pixel noise and systematic artefacts associated with asymmetric pupil features (such as spiders), the raw F&F wavefront estimates were projected onto a low-order modal correction basis, which can also be customised to constrain the degrees of freedom that F&F is able to control. The most appropriate modal correction basis to use given the current model of LWE wavefront morphology is a segmented basis consisting of independent PTT components for each VLT pupil quadrant, although Zernike or disk harmonic modes are less sensitive to specific pupil orientation and would also allow F&F to correct additional errors including NCPEs. Finally, the resulting phase maps were spatially filtered with a Gaussian kernel to approximate the limited spatial frequency response of the 41×41 actuator SPHERE DM. For simplicity it was assumed that these final filtered wavefront estimates could be applied accurately on the DM via reference offsets to the zero points of the main SAXO AO loop. Wavefront corrections were therefore directly applied on the (simulated) deformable element throughout this work. This assumption is encouraged by successful preliminary testing of the SPHERE DM response to piston offset commands (Sauvage et al. 2016), but an end-to-end simulation involving the SPHERE filtered SH-WFS and DM interaction matrix is beyond the scope of this paper.

More detailed and diagrammatic explanations of F&F wavefront correction loop architectures can be found in Korkiakoski et al. (2014) and Wilby et al. (2016b). The code implementation described in the latter paper has since been enhanced to explore additional considerations for SPHERE performance, with the capability to apply both systematic and random temporal variations in the gain of individual DM actuators (and hence imperfect reproduction of F&F phase estimates), the capacity to add additional amplitude aberrations and pupil rotation, a more realistic spatial filtering window function for atmospheric phase-screen generation, and a $1/f^2$ filtered spatial power spectrum of injected NCPEs. This allows an extension of the already published simulation results, especially for addressing the impact of factors such as amplitude aberrations or reconstruction errors when applying F&F update commands through a real AO system.

4.2.3 Laboratory verification of F&F on the MITHIC bench

In order to verify whether the simulation package discussed above is accurate in its implementation of F&F and in its treatment of the most important factors for on-sky observation, we performed closed-loop tests with the MITHIC high-contrast testbench at LAM (Vigan et al. 2016b). The F&F code package was implemented on the bench as described in the previous sub-section and executed for a number of controlled convergence tests, with variable parameters including the amplitude and type of wavefront error, image FOV, pupil apodisation, and image S/N. Each set of laboratory conditions was also run through an equivalent closed-loop simulation to provide a direct means of comparison between the two approaches. The various parameters and results of these tests are listed in Table 4.1, and are discussed in detail in Sect. 4.3.1.

The MITHIC bench operates at visible wavelengths (677 nm) and is optimised for high-contrast coronagraphic imaging using the Roddier and Roddier phase mask coronagraph (Roddier & Roddier 1997; N'Diaye et al. 2011), with the capability for pupil-plane wavefront measurement provided by Shack-Hartmann and ZELDA wavefront sensors, in addition to the COFFEE coronagraphic phase diversity estimator (Paul et al. 2014a) for focal-plane NCPE control. Figure 4.3 illustrates the optical layout of the bench as was used for this investigation, which is similar to that described in Paul et al. (2014a). A newly-installed phase-screen turbulence simulator (Vigan et al. 2016b) can be used to inject either single layer dynamic turbulence or a variety of static wavefront error patterns into the beam, including a LWE aberration with the same morphology as in Fig. 4.1 and an estimated amplitude of 98 nm PVE. Wavefront control was achieved using a Hamamatsu liquid-crystal-on-silicon spatial light modulator (LCOS-SLM), which is situated behind a beam-splitter in a face-on reflective configuration and samples the re-imaged pupil with 273 pixels across the diameter. The CoolSnap HQ2 1392 × 1040 pixel interline CCD camera was used for final focal-plane imaging, sampling the final PSF at 9.6 pixels per λ/D. For this work the camera was used in non-coronagraphic imaging mode, with the images numerically binned to a resolution of 3.3 pixels per λ/D and cropped to a 32×32 pixel FOV, matching the default imaging parameters of the DTTS. The PSF S/N was adjusted by varying the laser input power for a fixed exposure time of 1 ms, and was calculated based on the central image pixel with respect to the background noise floor.

In the absence of artificially injected sources of wavefront error, the MITHIC bench was measured by the HASO-3 SH-WFS to contain 96 nm RMS of astigmatism-dominated low-order static error, due to optical mis-alignments. After wavefront flattening using the SLM it was estimated that residual aberrations were reduced to the level of 10 nm RMS, with the imaging camera PSF displaying four consecutive unbroken Airy rings. This phase correction was then manually applied as a flat wavefront command using the SLM during this investigation, ensuring that injected aberrations comprised the vast majority of total wavefront error in the system for closed-loop tests.

A weak amplitude apodiser was included in the setup for a small subset of tests, inserted in the pupil shortly after the SLM in both an on-axis and off-axis position as shown in Fig. 4.4 in order to test the response of F&F to amplitude effects. During most SPHERE coronagraphic observations the DTTS imaging path also includes the strong APLC apodiser shown in black in Fig. 4.4, but a similar apodiser was not available in MITHIC for these tests. This is not in itself a concern for F&F performance: the radial amplitude function of the APLC apodiser is completely symmetric, allowing it to be included when calculating the reference focal plane field a which defines the zero

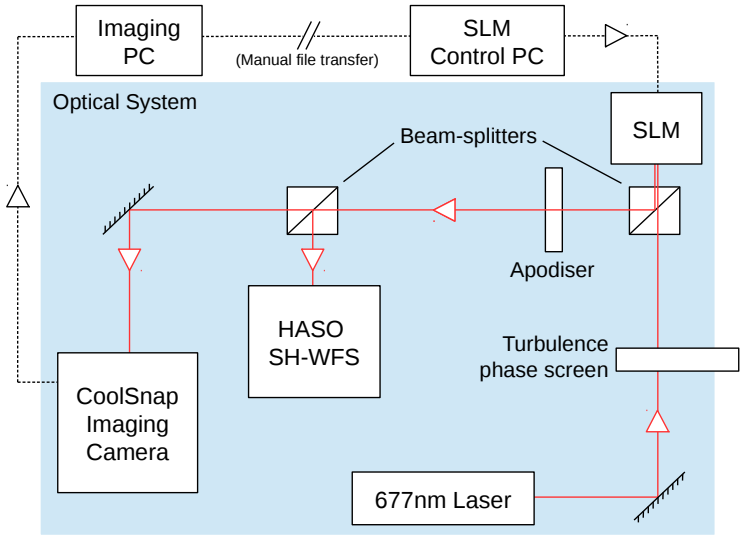


Figure 4.3: Schematic of the MITHIC bench configuration used for this investigation. The spatial light modulator (SLM) is placed in a face-on reflective configuration via a double pass through a beam-splitter. Images taken via the CoolSnap camera GUI were manually transferred from the imaging PC to the SLM control PC via a separate LAM server, before performing F&F wavefront reconstruction as described in Sect. 4.2.2 and applying the new wavefront commands on the SLM via a custom GUI. Each closed-loop iteration of F&F typically took between 30 s and a minute to complete, of which the F&F code runtime was a negligible fraction. The HASO SH-WFS was used only for the pre-compensation of MITHIC bench alignment errors and was not operated during F&F closed-loop tests. The turbulence phase screen was aligned either on the static LWE pattern or a clear aperture, and was not used to simulate dynamical turbulence.

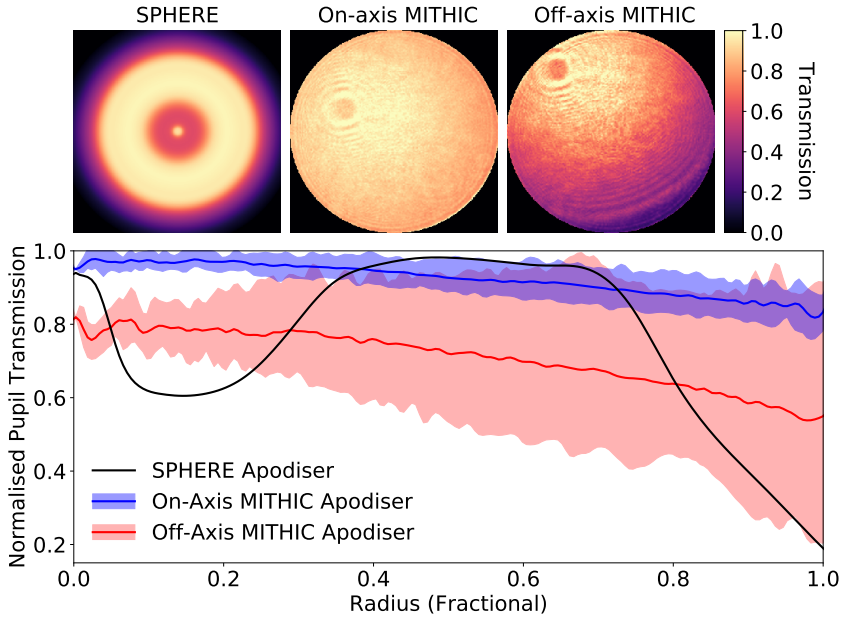


Figure 4.4: Throughput images (top panels) and radial transmission profiles (lower panel) of the amplitude apodisers relevant to this investigation. *Top row, from left:* SPHERE APLC apodiser model of Guerri et al. (2011), and pupil-imaging measurements of the on-axis and off-axis transmission (respectively) of the weak MITHIC bench apodiser. *Lower panel:* Dark lines denote the mean radial throughput curve of each transmission profile, with shading illustrating the maximum and minimum bounds in the azimuthal direction.

point for F&F wavefront reconstruction, without violating any of the assumptions on pupil geometry made in Sect. 4.2.1. The important consideration which is investigated here is whether the stability of F&F is limited by unknown amplitude errors, for example due to optical mis-alignments or errors in the pupil model used for F&F. The MITHIC apodiser pupil functions were therefore not provided to F&F when generating the reference field a for these tests, which instead used a uniform circular aperture matched to the MITHIC beam.

Due to the practicalities of using separate GUI interfaces for imaging and SLM control, coupled with the non-networked computing architecture currently implemented in MITHIC, wavefront correction could only be achieved by manually closing the loop. This limited the number of iterations which could reasonably be performed to a maximum of 25, due to time restrictions and risk of human error. As simulations and previous laboratory tests of F&F indicated that stable convergence is typically achieved in 5-6 iterations, this approach was deemed sufficient to characterise the initial convergence behaviour and place limits on the short-term post-convergence stability.

The limited number of closed-loop iterations, combined with the slow and irregular wavefront update frequency of manual control, also made it impractical to properly apply wavefront dynamics representative of SPHERE on-sky observations. All injected LWE and low-order aberrations in this investigation were therefore static, whether applied on the SLM or via static patterns on the turbulence phase screen. Most importantly, this means that MITHIC images did not include a simulation of atmospheric turbulence, quasi-static NCPEs or variability of the LWE itself. However, none of these dynamical factors were found to limit F&F performance when studied in the prior closed-loop simulations of Wilby et al. (2016b), provided that the wavefront correction loop can run sufficiently fast compared to the LWE coherence timescale: this point is discussed further in Sect. 4.3.2. The lack of an incoherent atmospheric speckle background does mean that the Strehl ratios quoted in this paper are close to the diffraction limit, and must be scaled down by the typical H-band performance of SAXO in order to make a comparison with expected on-sky performance.

In future it would be advantageous to fully automate the MITHIC wavefront correction loop, which would eliminate user error and allow more thorough investigation into the long-term closed-loop stability of focal-plane sensors, and their correction cadence requirements under variable conditions more closely resembling on-sky observations.

4.3 Results

4.3.1 Comparison of MITHIC bench and simulation results

Figure 4.5 shows the results of the most important convergence test performed on MITHIC, the details of which are listed in the first row of Table 4.1. This test applied a static 319 nm PVE LWE wavefront phase pattern on the SLM, which was chosen to provide half a wave of error at the 677 nm MITHIC laser wavelength. For the purposes of F&F this is directly equivalent to solving a 724 nm PVE LWE in the H-band, which is close to the strongest LWE amplitudes regularly seen on-sky with SPHERE. The setup did not contain a SPHERE aperture mask but used a clear circular pupil: regions of LWE wavefront phase corresponding to the VLT M2 obscuration and support spiders

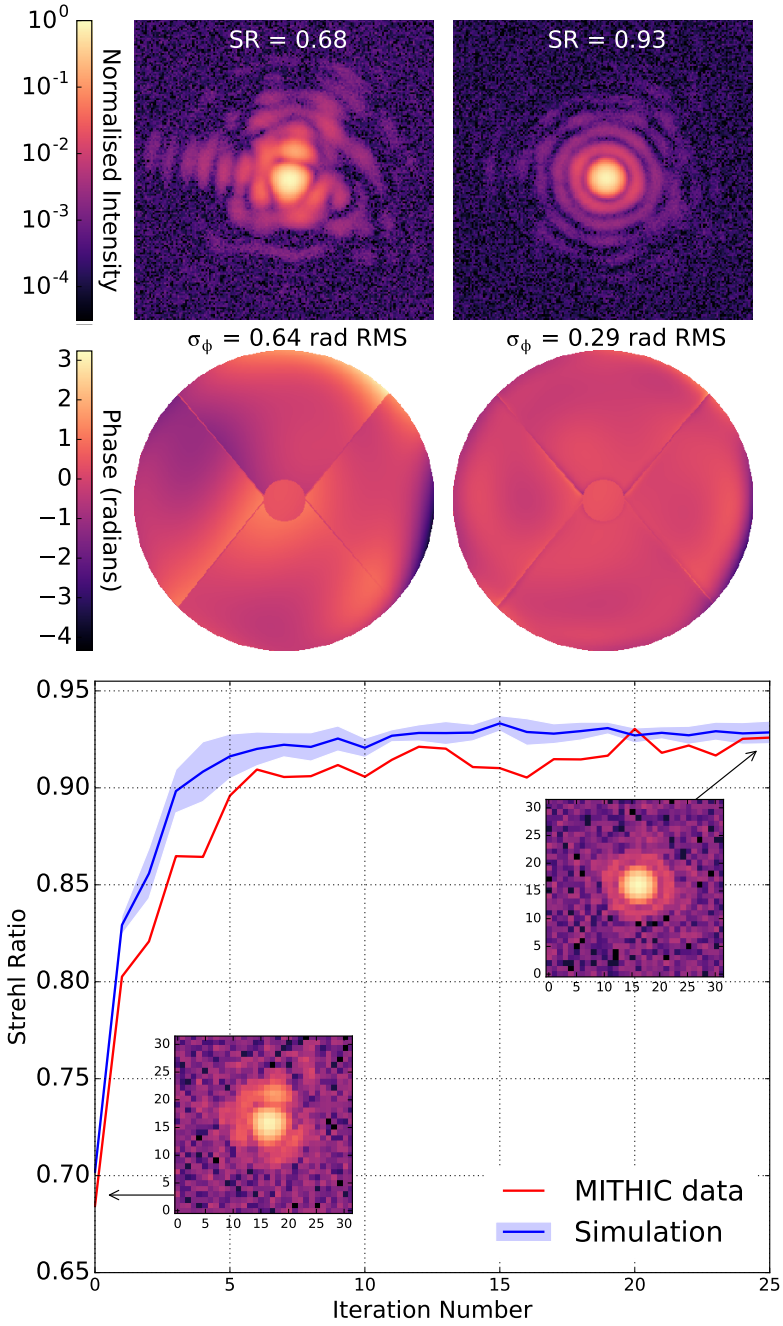


Figure 4.5: Comparison of F&F convergence behaviour between MITHIC laboratory data and equivalent simulations (see Table 4.1, row 1 for details). ***Top panels:*** High-resolution, high-S/N MITHIC focal-plane images (upper row) and corresponding residual wavefront error maps (lower row), before and after F&F correction (first and second columns respectively). High-frequency residuals visible in the final wavefront error map are dominated by fitting error enforced by the DM-like filtering of wavefront corrections. ***Bottom panel:*** Plot of estimated image Strehl ratio as a function of closed-loop iteration number, showing close agreement between a single convergence of MITHIC data (red) and an average of ten simulations (blue), with shading denoting the 1σ limit. The two inset images show the $N = 0$ and $N = 25$ DTTS-like input focal-plane images provided for F&F for wavefront reconstruction.

were therefore set to zero phase. In addition to the LWE, 39 nm RMS of low-order aberrations were also applied on the SLM, with coefficients randomly drawn from the first 19 Zernike modes. The image S/N and 32 x 32 pixel FOV were chosen to be representative of DTTS images, and a weak on-axis apodiser was included in the beam. Corrections were made on the SLM by first projecting the output of F&F onto a combination of the segmented PTT mode basis and a 50-mode low-order Zernike basis, and then spatially filtering to mimic the finite actuator influence function of the 41×41 actuator SAXO DM. This was achieved by convolving the resulting phase map with a Gaussian of full-width half-maximum equal to two thirds of the SAXO actuator spacing, which is qualitatively representative of the resolution seen in reference slope control tests by Sauvage et al. (2016). By not applying phase corrections at the full 273×273 pixel resolution of the Hamamatsu SLM with which the aberrating phase was initially implemented, we mimicked the interplay between a finite resolution DM and non-discrete upstream aberrating wavefront, thereby providing a better estimate of the correcting power of the SPHERE DM.

The topmost image row shows initial and final PSFs taken at high S/N, showing the clear improvement in Strehl ratio from 68% to 93% associated with the elimination of LWE and the majority of the low-order wavefront within five closed-loop iterations. The residual wavefront error maps show that the uncorrected wavefront error was dominated by high-frequency components along the edges of the spiders and along one edge of the pupil. This was a result of the DM-like spatial filtering of the correction wavefront, and is representative of the capabilities of such an implementation on-sky.

In the bottom panel, the image Strehl ratio is directly compared for each iteration between MITHIC data and an average of ten simulations matching the laboratory conditions. Multiple simulations were run in order to place a limit on the reproducibility of convergence under noisy conditions, with only the individual pixel-to-pixel noise allowed to vary between realisations. This plot shows that there is an extremely close match between predicted and obtained performance in this scenario, with very similar final Strehl ratios achieved and laboratory convergence only very slightly lagging behind the simulated curve. It was also estimated from SH-WFS measurements during this convergence process that there existed approximately 10 nm RMS of residual NCPE in the system, for which F&F is in principle also capable of correcting. This means that the laboratory implementation was dealing with additional wavefront error when compared to the simulated case, which may explain this small decrease in convergence efficiency.

All Strehl ratios for MITHIC bench images quoted in this paper were estimated according to a modified encircled energy metric described by

$$S_{\mathrm{MITHIC}} \approx \mathrm{EE}(p_{\mathrm{MITHIC}}) \frac{S_{\mathrm{sim}}}{\mathrm{EE}(p_{\mathrm{sim}})} \text{ , where} \tag{4.10}$$

$$\mathrm{EE}(p) = \frac{\sum p(r < 1.22\lambda/D)}{\sum p(r < 3.5\lambda/D)} \tag{4.11}$$

is the ratio of encircled energies between the Airy core and first Airy ring of the given PSF p, which provides an estimate of image quality using the regions visible above the noise floor. As the encircled energy metric systematically over-estimates the true Strehl ratio of MITHIC images p_{MITHIC} by ignoring high-frequency aberrations, the second term in Eq. 4.10 used to account for this bias using the known true Strehl ratio

S_{sim} and encircled energy of the simulated PSF p_{sim} of matched simulations. Such a correction was possible due to the close consistency between laboratory and simulation results, such that the correction factors could be expected to be very similar for the two cases. This metric was chosen over more readily available metrics such as using the SLM residual wavefront map or a reference image PSF from the MITHIC bench, as these can be biased by factors such as residual NCPEs or intensity variations in the system, and in practice were seen to occasionally predict Strehl ratios of greater than one.

Table 4.1 also summarises the laboratory performance obtained with F&F in MITHIC for a range of additional key tests, which illustrate the behaviour of the algorithm under variable conditions. It can be seen that for all but the final three test cases F&F returned the image Strehl ratio to over 90%, a gain of typically greater than 20% on the starting value. The right-most column quotes the post-convergence RMS wavefront error, which was in most cases estimated directly from the last phase command applied on the SLM after closed-loop convergence. These RMS residuals were found to be typically on the same order as the 10 nm RMS MITHIC bench alignment residual after HASO pre-compensation, and agree well with the quoted Strehl ratios through the use of the Maréchal approximation. In cases where the target wavefront error was not injected with the SLM, such as when using the turbulence screen LWE aberration or removing the MITHIC bench wavefront flat command, this approach is inherently biased. For these tests the estimate of the final RMS residual was then simply made using the Maréchal approximation with the final image Strehl ratio. For the final three test cases, the aim was not to correct all wavefront error present in the system but to leave specific aberrations uncorrected, in a manner which would allow F&F to solve only LWE-like wavefronts without impacting the system NCPE budget. These cases are discussed in more detail in bullet point 8 below, and in Sect. 4.4.

The following specific observations can be made about the performance of F&F by comparing the various scenarios presented in Table 4.1:

1) **Field-of-view:** Shrinking the square FOV available to F&F from 100 to 32 pixels had only a few percent impact on final PSF quality (rows 9. & 10.); this was likely due to the removal of high-order wavefront information from the PSF supplied to F&F.

2) **Signal-to-noise:** Lowering the image S/N by a factor of 16, from 6380 to 382 (rows 2 & 7), had a negligible impact on final performance in the high-S/N regime (see Sect. 4.3.2 for a simulated treatment of low-S/N performance).

3) **Choice of mode basis:** Comparing the use of segmented PTT (row 9.) and 50-mode Zernike (row 8.) bases shows that, as expected, Zernikes were less able to replicate the high-frequency LWE wavefront. There is negligible difference between a PTT-only basis and the full PTT + Zernike basis (row 7.), although the latter is expected to perform better in the presence of additional non-LWE aberrations, or for erroneous pupil rotation angles where the PTT basis is no longer a good description of the LWE wavefront.

4) **Using no mode basis:** Comparing the 100×100 pixel FOV PTT basis test (row 10.) with a comparable test applying only Gaussian-filtered (i.e. DM-resolution) F&F outputs without mode basis projection (row 13.) shows that using the PTT basis actually resulted in a three percent increase in final Strehl ratio. This can

Table 4.1: Summary of the key parameters of all laboratory F&F convergence tests performed on the MITHIC bench. The first row corresponds to the most challenging scenario, which is presented in Fig. 4.5. Strehl ratio (SR) estimates are calculated from normalised encircled energy measurements of the PSF core and first Airy ring, as described by Eq. 4.10. Final RMS error is estimated from the residual wavefront map implemented on the SLM, or the Maréchal approximation when significant non-SLM induced aberrations are involved. The quoted 1σ errors on S/N, Strehl ratio, and residual RMS are calculated from the final five frames in each test sequence, and as such are a reflection of post-convergence stability. The final three rows denote test cases attempting to force F&F to avoid correcting specific wavefront errors, by manipulating the correction basis or subtracting reference offsets from F&F output wavefronts.

Row	LWE (PVE / RMS)	Low-Order (RMS)	Corr. Basis	Apodiser	FOV (px)	Final S/N	Initial SR	Final SR	Final RMS
1.	319 / 59 nm	39 nm (Zernike[1])	PTT + Zern	On-axis	32	445 ± 2	68 %	93 ± 1 %	32.2 ± 0.5 nm
2.	319 / 59 nm	0 nm	PTT + Zern	None	32	382 ± 1	72 %	94 ± 1 %	20.5 ± 0.8 nm
3.	98 / 16 nm[2]	96 nm (Bench[3])	PTT + Zern	None	32	6247 ± 6	56 %	91.5 ± 0.5 %	28 ± 1 nm[4]
4.	319 / 59 nm	0 nm	PTT + Zern	Off-axis	32	4400 ± 30	72 %	94 ± 1 %	24 ± 2 nm
5.	319 / 59 nm	0 nm	PTT + Zern	On-axis	32	6410 ± 20	74 %	97.1 ± 0.5 %	21 ± 2 nm
6.	98 / 16 nm[2]	0 nm	PTT + Zern	None	32	6450 ± 30	94 %	96.2 ± 0.1 %	20.7 ± 0.3 nm[4]
7.	319 / 59 nm	0 nm	PTT + Zern	None	32	6380 ± 40	73 %	94.9 ± 0.1 %	18 ± 1 nm
8.	319 / 59 nm	0 nm	Zern	None	32	5980 ± 20	73 %	91.1 ± 0.1 %	24.8 ± 0.5 nm
9.	319 / 59 nm	0 nm	PTT	None	32	5970 ± 40	73 %	94.5 ± 0.2 %	14.4 ± 0.6 nm
10.	319 / 59 nm	0 nm	PTT	None	100	6070 ± 30	73 %	97.9 ± 0.1 %	14.9 ± 0.5 nm
11.	319 / 59 nm	0 nm	PTT[5]	None	100	6020 ± 30	73 %	98.0 ± 0.1 %	16.4 ± 0.7 nm
12.	98 / 16 nm[2]	0 nm	None	None	100	5850 ± 20	93 %	97.1 ± 0.2 %	18.5 ± 0.6 nm[4]
13.	319 / 59 nm	0 nm	None	None	100	5820 ± 30	72 %	94.8 ± 0.2 %	22.3 ± 0.7 nm
14.	319 / 59 nm	0 nm	PTT[6]	None	32	5570 ± 70	74 %	86 ± 1 %	32 ± 2 nm[7]
15.	319 / 59 nm	43 nm (Focus)	PTT[8]	None	32	5430 ± 40	66 %	82 ± 1 %	42 ± 1 nm[7]
16.	98 / 16 nm[2]	77 nm (Zernike[1])	PTT	None	32	3790 ± 20	66 %	71 ± 1 %	64 ± 1 nm[4]

(1) Low-order static wavefront, consisting of the first 19 non-trivial Zernike modes (excluding PTT) with randomly drawn coefficients.
(2) LWE aberration is applied via the weak static phase pattern included in the MITHIC turbulence module (Vigan et al. 2016b), instead of via the SLM.
(3) Residual astigmatism-dominated MITHIC bench alignment errors, which were measured by the HASO SH-WFS and pre-compensated for all other tests.
(4) RMS residuals are estimated via Maréchal approximation, as the majority of wavefront error is not applied via the SLM and thus residual phase is poorly known.
(5) No DM-like spatial filter is applied to F&F outputs, hence wavefront corrections are applied at the full SLM resolution.
(6) Global tip-tilt is subtracted from F&F estimates before projection onto the PTT basis. The PSF is initially shifted by 1×1 pixels from the F&F reference centroid.
(7) F&F exhibits divergent behaviour within ten closed-loop iterations.
(8) 43 nm of focus is subtracted from each F&F output wavefront estimate before projection onto the PTT basis, in an attempt to apply a zero-point offset.

be attributed to the effective removal of unwanted high-frequency noise propagation and pupil asymmetry systematics from the F&F output estimates by the tailored mode basis.

5) **Wavefront corrector resolution:** Disabling the DM-like Gaussian spatial filtering and applying corrections at the full resolution of the SLM had negligible impact on the final correction performance (rows 10 & 11). In this case the maximum correctable spatial frequency was effectively set by the FOV of DTTS images. Residual wavefront error maps show high-frequency residuals around the locations of the spiders for filtered wavefront tests as in Fig. 4.5, however these result in only a small amount of additional diffraction along the spiders and hence have a low impact on Strehl ratio.

6) **Pupil apodisation:** A weak, on-axis pupil apodiser (row 5, and top-centre panel of Fig. 4.4) in fact resulted in marginally better performance than the equivalent un-apodised case (row 7.) despite F&F calculations still assuming a uniform pupil; this is attributed to the suppression of high-order aberrations. F&F was still stable in the presence of a stronger, asymmetric apodisation of the pupil which even exhibited some vignetting (row 4, and top-right of Fig. 4.4) when the algorithm was still not provided with the modified pupil function, showing only a few percent loss in Strehl compared to the above cases. This indicates that F&F is extremely stable against unknown amplitude aberrations, and even severe pupil mis-alignments which violate the even-pupil assumption implicit in F&F. For this particular application it is therefore unnecessary to implement additional amplitude retrieval steps (such as the FF-GS extension to F&F mentioned in Sect. 4.2.1) to ensure robust performance in SPHERE.

7) **Source of phase aberrations:** F&F was equally capable of correcting strong aberrations from external (non-SLM) sources (row 3.) as it was for SLM-induced aberrations (row 1.), returning both to a final estimated Strehl ratio above 90%. This indicates a sufficiently accurate orientation, alignment, and phase-to-voltage calibration of SLM commands was achieved for closed-loop correction.

8) **Applying reference offsets:** Attempting to make F&F insensitive to PSF centre by removing global tip-tilt components in the output wavefront (row 14.) or attempting to induce specific wavefront reference offsets by manually subtracting them from the F&F output on each iteration (row 15.) resulted in unstable convergence and were not viable methods for this implementation of F&F (see below and Sect. 4.4 for details). Using the natural lack of sensitivity of a PTT-only basis to low-order Zernike modes to try and correct only the LWE component of an aberrated wavefront (row 16.) resulted in a stable convergence, however the final correction may have included some unwanted partial compensation of non-LWE errors, as the final residual RMS error of 64 nm was smaller than the 77 nm of low-order (i.e. non-LWE) aberrations initially applied on the SLM.

As with the main results presented in Fig. 4.5, the various test cases presented in Table 4.1 were found to be highly representative of closed-loop simulations directly emulating the conditions of each test case, indicating that F&F was performing very close to the expected level in this MITHIC implementation. The predictive power of these

simulations extends to the identification of the two key limitations so far identified for F&F; a sensitivity to the centroid zero-point of the PSF in the image (row 14.), and difficulty effectively converging to a non-flat wavefront via direct reference phase map subtraction (row 15.). In both cases divergent behaviour was seen to set in within ten iterations after an initial improvement in wavefront quality, and therefore it is important to evaluate the underlying causes and potential solutions; this analysis is presented in Sect. 4.4.

4.3.2 Simulated low-S/N performance of F&F

One of the most important concerns with using the SPHERE DTTS as a focal-plane wavefront sensor is that the low throughput to the DTTS camera results in low S/N images, especially for faint targets. The DTTS control loop is typically operated at a cadence of 1 Hz and is designed to function down to a S/N of approximately ten, and so F&F should also be stable under these conditions. However, any attempt to sense LWE-like wavefronts from focal-plane images will ultimately be limited by the S/N of the first Airy ring, which corresponds to the dominant spatial frequencies present in this type of wavefront. This forms a significantly stronger S/N constraint than that required for simple tip-tilt correction using the PSF core, and will by necessity limit the efficiency of LWE correction at low S/N.

Figure 4.6 illustrates the simulated performance of F&F as a function of input image S/N, at the laboratory wavelength of 677 nm for comparison with MITHIC results. This was performed in the absence of atmospheric residuals or NCPEs, with only the 319 nm PVE LWE phase aberration (identical to that used for MITHIC bench tests) present. Each data point shows the average final Strehl ratio and wavefront error of ten independent simulations each of 25 iterations, such that the error bars provide an estimate of the post-convergence frame-to-frame stability of F&F. Here all S/N values are quoted for the central pixel of the PSF, and in all simulations a constant value of $\epsilon = 10^{-3}$ was used for wavefront reconstruction. It can be seen from the red (upper) data points that the algorithm is stable over the entire range of S/N values, and still makes some statistical improvement to the wavefront quality even at the minimum S/N of five. Above this the residual wavefront error declines logarithmically, and Strehl ratios of greater than 95% are achieved for S/N greater than 100. The blue (lower) residual RMS wavefront error data points also provide a useful estimate of the S/N-limited sensitivity of this implementation of F&F to low-amplitude aberrations. This curve is consistent with the equivalent MITHIC test result (row 2 of Table 4.1), which for an image S/N of 382 achieved 20 nm of residual RMS error. This is equivalent to a sensitivity limit of 45 nm RMS in the H-band, since F&F operates in radians and so performance can be expected to scale linearly with wavelength.

Such robust performance at low S/N was only possible with the use of an adaptive focal-plane spatial filter to attenuate pixel-to-pixel noise. This filter modifies the input DTTS PSF p by smoothly replacing noise-dominated pixels with the reference (diffraction-limited) PSF $|a|^2$, such that F&F sees the higher spatial frequencies as perfect and so provides zero correction on these scales. The resulting filtered PSF p_{filt} is

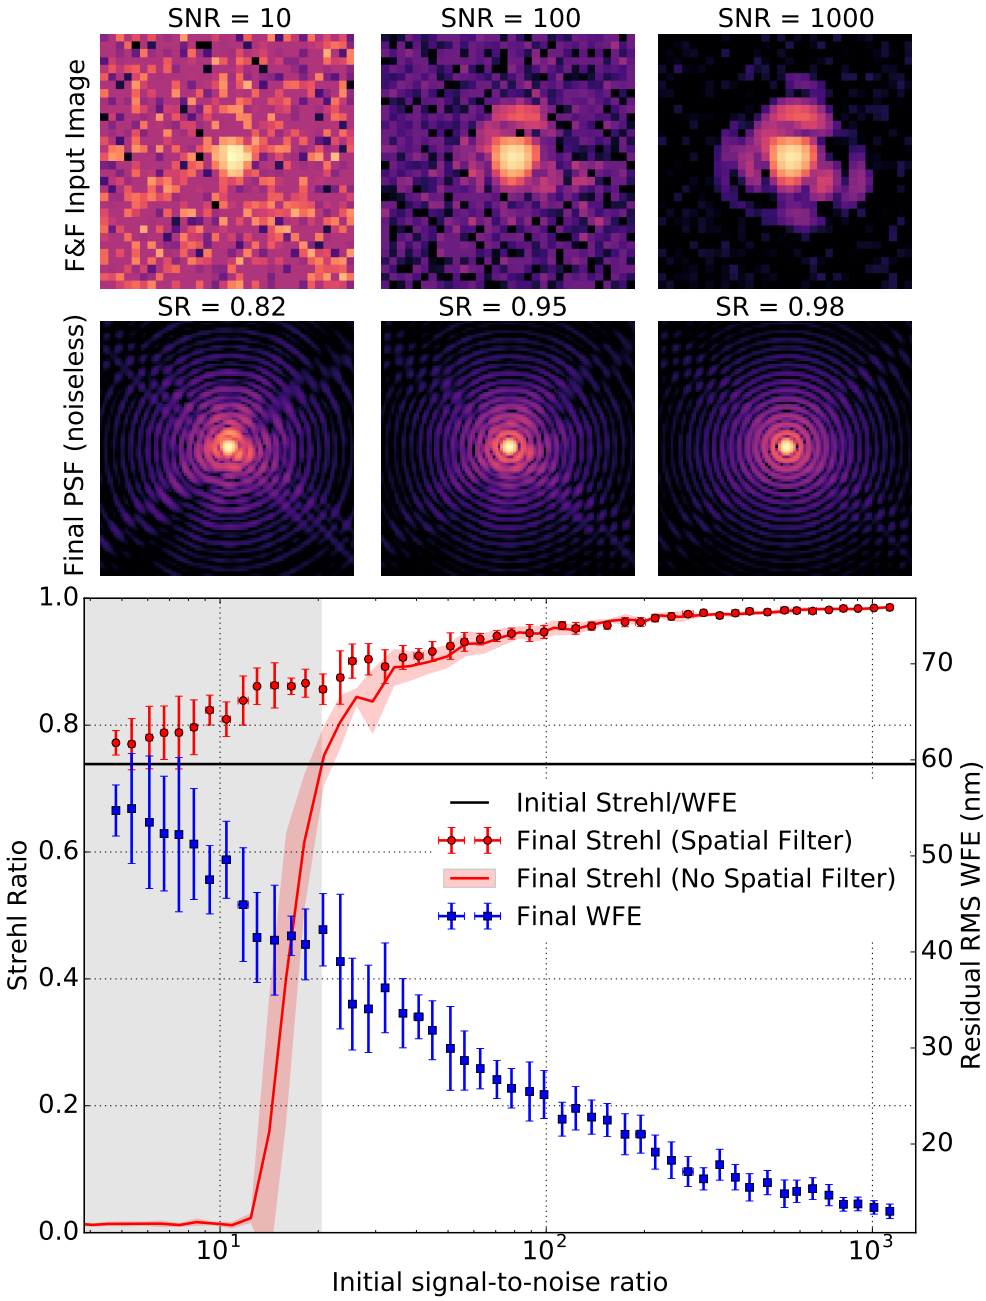


Figure 4.6: Simulated convergence quality of F&F as a function of image S/N, for half-wave (319 nm) PVE LWE at a wavelength of 677 nm. *Upper image row:* Initial DTTS images at specific S/N, showing increasing visibility of the aberrated first Airy ring. *Lower image row:* Final noiseless images after 25 F&F iterations, showing a corresponding improvement in final PSF quality. *Main panel:* Plot of final Strehl ratio (red) and residual RMS wavefront error (blue) as a function of initial S/N. Each point is the average of ten simulations. The black horizontal line denotes both the starting Strehl ratio and wavefront error RMS, of 74% and 59 nm respectively. The shaded red line shows the Strehl ratio behaviour for F&F using the full 32 × 32 pixel DTTS image as input (i.e. without an adaptive spatial filter), with the shaded region below S/N = 20 denoting the region where this implementation diverges.

defined by

$$p_{\text{filt}} = w(r)p + (1 - w(r))\,|a|^2 \ \text{, where} \tag{4.12}$$

$$w(r) = \frac{1}{1 + e^{r - r_c}} \tag{4.13}$$

is a radial sigmoid windowing function with a critical cut-off radius of r_c. This critical radius is defined to be the point at which the local maxima of the diffraction-limited PSF first drop below a critical S/N of 2.5, which in this instance was found to be an acceptable balance between rejecting noise and preserving wavefront information in the image. It may be possible to further improve the performance of this filter by optimising the functional form for $w(r)$, for instance by using the generalised logistic function. A detailed investigation is however beyond the scope of this paper.

For comparison, the red shaded curve shows the natural behaviour of F&F in the absence of an appropriate spatial filter, using the full noisy 32×32 pixel FOV as the input. In this case the pixel-to-pixel noise at high spatial frequencies is directly propagated into strong modal noise in the final correction, resulting in a rapid divergence in wavefront quality below a S/N of 20. With spatial filtering applied it can be seen that the algorithm instead 'fails gracefully', simply correcting less of the aberrating wavefront as the corresponding spatial frequencies fall below the image noise threshold. This improved approach also achieves a final post-convergence wavefront quality of better than 40 nm RMS and stability of better than 20 nm RMS 1σ jitter for all but the lowest S/N values, which can most likely still be improved by more careful optimisation of the adaptive spatial filter profile and cut-off radii as a function of S/N.

For any given observing conditions it should also be possible to further improve F&F performance by stacking individual 1 s DTTS frames for a longer effective exposure and hence higher S/N, provided that the correction cadence remains significantly shorter than the variability timescale of the LWE. A recent study of the morphology and temporal evolution of the LWE as seen by the SPHERE-IRDIS subsystem (J.F. Sauvage, ESO, private communication, 2017) concluded that under typical conditions (1 m/s wind speed, average LWE of 600 nm PVE, 10 Hz imaging cadence) the majority of the LWE-related structures were coherent on timescales of longer than 10 s, although some small amount of short-term variability was also observed. Frame stacking up to this 10 s threshold would facilitate an additional S/N boost of a factor of three, which even for a faint target with a 1 s exposure image S/N of ten would return image S/N to the regime in Fig. 4.6 where we can expect final corrected Strehl ratios of 90% or greater, corresponding to the removal of the majority of low-frequency LWE error and a significant reduction of spider diffraction effects. In general however, it is advantageous to sense and correct at least ten times faster than the shortest wavefront coherence timescale of interest. This means that F&F will most likely perform better when operated at the fastest cadence (i.e. with the lowest image S/N) which can be expected from Fig. 4.6 to provide sufficiently good correction for any given application.

4.4 Discussion

It is important to discuss the two key limiting factors identified in F&F, which are found both in simulations and MITHIC lab tests as mentioned in Sect. 4.3.1. The first of these is sensitivity to the centroid location of the image PSF: the tip-tilt zero-point

to which F&F will try to converge is set by the centroid location of the reference PSF $|a|^2 = \left|\mathscr{F}[A]\right|^2$, and so the algorithm naturally attempts to apply global tip-tilt corrections if the image centroid differs from this reference. While this tip-tilt correction was observed to be robust up to a dynamic range of approximately 1 radian, it raises the potential for loop conflict if the F&F zero-point differs from that of the main tip-tilt sensing loop for which the DTTS is primarily used, or if frame-to-frame DTTS tip-tilt correction residuals approach 1 radian. It is therefore desirable to make F&F completely insensitive to tip-tilt error, for which the naïve approach is to directly subtract any measured global tip-tilt components from the output wavefront estimate. However, this approach only results in the build-up of differential tip-tilt between individual VLT pupil segments and a slow divergence of wavefront quality over time. This is because F&F is not a perfect one-shot phase reconstructor: the exact tilts measured across each pupil segment tend to differ slightly from the global gradient, with the residual between the two still included in the wavefront correction command. As F&F is still sensitive to the subtracted global tip-tilt error on each subsequent iteration, this residual differential tip-tilt map is re-applied on each iteration and thus builds up steadily over over time. The DTTS image can always be re-centred to pixel precision by shifting the image array (in this case corresponding to a precision of 0.3 λ/D) to somewhat limit the extent of this centroiding issue, however the best approach is instead to ensure that the reference PSF p_0 is constructed to exactly match the zero point of the DTTS to sub-pixel accuracy, thereby ensuring that there are no conflicts between the two correction loops.

The other outstanding issue associated with tip-tilt control is whether a stable, converged F&F control loop adds any additional positional jitter to the PSF, and whether this remains within the specifications of the SPHERE design requirements for coronagraph centring. This was investigated on the MITHIC bench for the main 25-iteration convergence test, previously presented in Fig. 4.5 and row 1 in Table 4.1. The absolute deviation of the PSF centroid from its mean location at 3.3 pixels per λ/D focal-plane sampling is shown for all frames in Fig. 4.7, comparing F&F control with a short reference image sequence at the same cadence, containing only the natural bench image jitter without an active control loop. This shows that F&F achieved sub-DTTS pixel stability with an RMS of one tenth of a pixel, equivalent to $0.03\ \lambda/D$ and hence 1.2 mas on-sky for SPHERE, only slightly higher the natural bench jitter of $0.02\ \lambda/D$ (0.8 mas). It is likely that this result was dominated by normal thermal and mechanical fluctuations in the MITHIC optical path, given the low temporal bandwidth of the manual F&F control loop and the lack of any other form of active PSF centring control. The on-sky positional stability of SPHERE on the other hand is dominated by the telescope high-frequency vibration environment, with a target of 3 mas (0.07 λ/D) RMS required for baseline coronagraph operation (Fusco et al. 2016). If it may be assumed that F&F would add the same amount of additional jitter as was seen in MITHIC, this would constitute only a few percent of the total error budget. Due however to the fact that F&F would contribute at frequencies between 0.1-10 Hz depending on correction cadence (which is significantly slower than the 10-100 Hz vibrations limiting SPHERE stability), a more detailed investigation would be needed to determine its exact impact on the SPHERE vibrational error budget.

The second known limitation is in attempting to force F&F to converge to a specific non-flat wavefront, which is in general a useful feature enabling the application of controlled reference wavefront offsets. In this case, the differential optical path be-

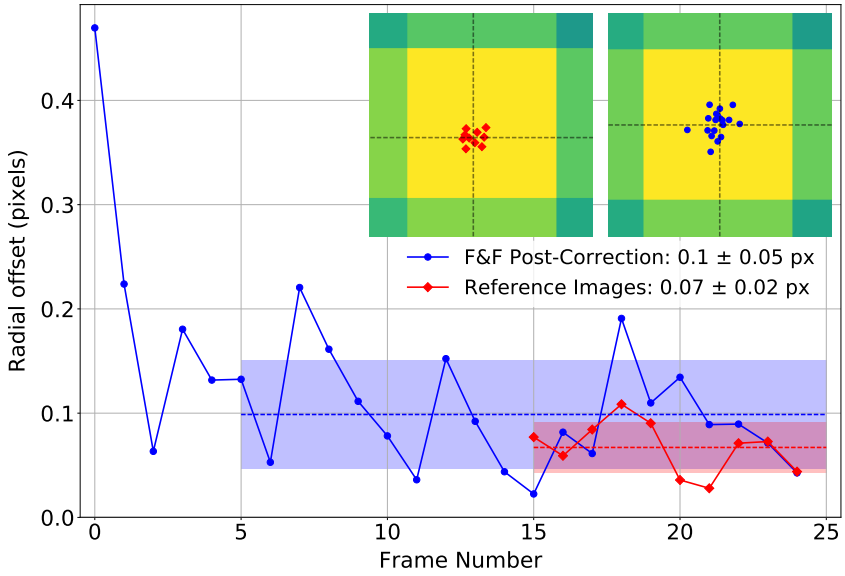


Figure 4.7: Impact of closed-loop F&F correction on the tip-tilt stability of the MITHIC bench. *Main panel:* Measured radial offsets from the mean PSF centroid location as a function of iteration number, for the headline MITHIC F&F convergence test of Fig. 4.5 (blue) and a set of reference images containing only natural MITHIC bench image jitter (red). F&F convergence is taken to be achieved after five iterations. *Inset panels:* The central 0.3 λ/D pixels of the mean reference image (left) and the F&F-corrected image (right) respectively, each over-plotted with the centroid data used to compute jitter estimates. Dashed cross-hairs show the mean PSF centroid for each dataset, showing that F&F applies a systematic sub-pixel offset to align the PSF with its internal reference zero-point.

tween the DTTS and IRDIS focal planes is known to contain 20 nm of static focus error plus some additional higher order NCPEs, which ideally should not be introduced into the science beam by a DTTS-based WFS. While it was found in simulations that the most straightforward approach of subtracting the fixed reference offset from the F&F output wavefront on each iteration is stable in the noiseless case, under realistic conditions the ϵ regularisation parameter in Eq. 4.5 (which is typically chosen to be comparable to the image noise floor) results in a systematic underestimation of the even wavefront. As for the case of the global tip-tilt drift phenomenon described above, this results in the injection of the residual between the sensed and true offset phase map on each iteration, resulting in the divergent behaviour seen in row 14 of Table 4.1. While it is possible to systematically ignore specific modes entirely (for example focus or astigmatism) by subtracting the measured coefficient of that mode from the raw F&F output wavefront before final mode basis projection, it is unclear how this would affect the efficiency of LWE correction: the segmented PTT basis is not fully orthogonal to such Zernike modes, and astigmatism in particular is typically present in the LWE wavefronts we wish to correct.

Attempts to modify the F&F algorithm itself to properly treat reference offsets are ongoing, however this is made challenging by the fact that the simplifying assumption of a real, even pupil function necessary for an analytical solution also prevents direct modification of the target PSF to include reference phase aberrations, i.e. $a = \mathscr{F}[A] \rightarrow \mathscr{F}[Ae^{i\phi_{ref}}]$. This is because a number of terms in the original Taylor expansion of the PSF which correspond to non-even and complex aperture terms are deliberately neglected from the derivation before arriving at Eq. 4.1 in order to obtain

an analytical solution for the focal plane fields. Because these neglected terms now become significant, such a substitution for a is no longer valid. A generalised version of F&F capable of arbitrary wavefront reference offsets would also be of great interest for operation with coronagraphic images, especially the apodising phase plate (APP) coronagraph (Kenworthy et al. 2010c; Snik et al. 2012; Otten et al. 2017). However, given the low level of NCPEs in the SPHERE case of interest there would be a minimal impact on science image Strehl ratio associated with allowing free convergence to the DTTS focal-plane: the small degradation in image quality is expected to be vastly outweighed by the gain in raw contrast performance from controlling the LWE. In high-S/N environments, it would also be expected that using F&F to stabilise NCPEs up to the NIR coronagraph with a general Zernike mode basis would also provide an improvement in final high-contrast imaging performance despite inducing 20 nm of static focus error into the science beam.

The results presented in this paper are somewhat idealistic in that they assume the XAO system is composed of the sole DM component to which the output phase commands from F&F can be accurately implemented via phase conjugation. In reality, correction must be achieved by modifying reference slope offsets on the SAXO SH-WFS during operation of the main XAO sensing and correction loop. In addition to the potential for control loop conflicts, it is currently unknown how this approach will filter the high spatial frequencies present in the LWE wavefront, and at what point the finite dynamic range of the SH-WFS will limit the correction of high-amplitude LWE cases. The remaining stroke on the SAXO DM during closed-loop operation will also determine how effectively the highest amplitude LWE cases can be corrected. For example, an 800 nm PVE LWE would constitute 11% of the total ± 3.5 μm SPHERE DM stroke and 26% of the ± 1.2 μm inter-actuator stroke (Fusco et al. 2006). Since the LWE occurs under good seeing conditions where there is less strain on the AO system, it is likely that almost all LWE cases would see significant improvement before being limited by DM saturation. Due to the complicating factors listed here, applying any form of focal-plane wavefront control using the SPHERE DM and SH-WFS is clearly still an important area to be addressed, and may require the development of a dedicated control scheme. However, initial tests have been encouraging in showing that the DM can accurately reconstruct a strong differential piston via the reference offset approach in both open and closed loop, with the width of the phase discontinuity boundary at pupil segment edges consistent with the influence function of the DM actuators (Sauvage et al. 2016).

Altogether it is expected that with appropriate calibration the current implementation of F&F is capable of providing at least an order of magnitude of raw contrast improvement in SPHERE coronagraphic imaging performance at 2-4λ/D in typical LWE-affected cases for S/N greater than 20, by returning the distorted first Airy ring to near-diffraction-limited performance (see Fig. 4.2). Additional gains may also be made in post-processing if F&F has sufficient image S/N to stabilise quasi-static speckle structure, allowing for more effective reference PSF subtraction or removal via ADI or PCA-based PSF subtraction techniques. By running in continuous closed-loop mode, on-sky performance can also be expected to be superior to that presented in Fig. 4.6 and Wilby et al. (2016b), since both this work and the previous study are concerned with compensating an established LWE wavefront error. Provided that F&F is operated above the critical cadence of ten times the variability timescale of PTT wavefront errors as they arise. In low-S/N cases a piston-only correction loop can still be

expected to reduce the impact of many LWE wavefronts on the coronagraphic PSF, and is also less likely to conflict with the main AO loop as the SH-WFS is in principle insensitive to differential piston aberrations.

4.5 Conclusions

We have demonstrated that the Fast & Furious sequential phase diversity algorithm is capable of robustly eliminating strong LWE wavefronts in the MITHIC high-contrast laboratory testbench environment, where it reliably returned image Strehl ratios to better than 90% within five closed-loop iterations. This was achieved in the presence of strong static low-order aberrations, low S/N, and small FOV images representative of the SPHERE DTTS, but in the absence of incoherent atmospheric speckle residuals or an active primary XAO loop, and assumed an idealised SPHERE DM response for correction. We find no significant discrepancies between these MITHIC laboratory results and the predictions of dedicated LWE simulation code (Wilby et al. 2016b) designed to emulate focal-plane wavefront sensing with the DTTS sensor. Therefore, further work carried out using this code is expected to be representative of performance achievable with F&F on-sky with the SPHERE instrument.

Supporting simulations showed that this implementation of F&F is also stable over the full working S/N range of the DTTS sensor down to at least S/N = 5, and is capable of efficiently removing the dominant structures of the LWE for $S/N \geqslant 20$. For targets where this condition can be satisfied for correction cadences faster than the dominant LWE variability timescale (estimated from IRDIS observations to be longer than 10 s) an on-sky implementation of F&F should be capable of effectively maintaining a near-diffraction-limited PSF core under the strongest LWE conditions routinely seen by SPHERE. Such an improvement is expected to provide at least an order of magnitude gain in raw contrast over typical LWE-affected PSFs close to the coronagraphic inner-working angle, greatly improving the ultimate contrast performance of the SPHERE instrument under the best seeing conditions.

Further efforts will focus on understanding the interplay between F&F and a realistic AO environment, including the spatial filtering properties of WFS reference slope offset based control, and the potential for conflicts in a multi-control loop system. It is also of great interest to develop a generalised version of the algorithm which is capable of converging to arbitrary non-flat reference wavefronts: in addition to providing greater flexibility for closed-loop control, this would also allow F&F to operate directly with many types of coronagraphic science image.

In addition to this specific application for controlling the LWE in SPHERE, the stability and versatility of F&F makes the algorithm highly suitable for other real-time focal-plane wavefront control tasks, such as NCPE correction or mirror co-phasing of segmented telescopes, provided that narrowband image data is available at a sufficiently fast cadence. With instances of the more general 'island effect' now being repeatedly seen in high-contrast instruments beyond the VLT, it is clear that focal-plane wavefront control methods such as F&F could become increasingly essential for the field. This will be especially important for the upcoming ELT-class telescopes, which will feature highly segmented pupils and a large amount of obscuring support structure. It can be expected that these telescopes will be more prone to island effect and LWE phenomena than simple four-quadrant pupil geometries, and without appro-

priate mitigation strategies these effects may severely limit the performance of their XAO-fed high-contrast instruments. The computational simplicity of F&F allows it to scale efficiently to work with high-resolution deformable elements, and in principle makes it sufficiently fast for high-speed (kHz) wavefront control applications: this makes it an attractive focal-plane phase control solution for both current and future instruments.

Chapter 5

These are not the voids you're looking for
Using Ceres to calibrate SPHERE-IRDIS coronagraphic DPI observations of TW Hydrae

Michael J. Wilby, Jos de Boer, Rob G. van Holstein, Christian Ginski, Julien Girard, Anthony Boccaletti, and Christoph U. Keller
In preparation

Abstract

The use of coronagraphic optics and dual-band polarimetric imaging (DPI) techniques to remove unwanted stellar contamination from high-contrast images leads to significant transmission losses at small angular separations. This makes it challenging to determine whether features observed in these regions are due to intrinsic disk features such as inner cavities, or are simply artefacts of the imaging system. This work proposes two calibration approaches of differing levels of complexity, which can be used to correct the relative photometry of circumstellar disk observations. The first uses measurements of the radial coronagraphic extinction profile to directly normalise observations, while the second uses a full model of the imaging system to account for non-linear effects when forward-modelling the object in question. e analyse observations of the minor planet Ceres in order to measure the extinction profile of the SPHERE-IRDIS apodised Lyot coronagraph. From this an effective model of the coronagraphic imaging system is developed, which is used to calibrate polarimetric observations of the TW Hydrae protoplanetary disk. We find that a four-component power law model provides an accurate reconstruction of the observed TW Hydrae signal below 600 mas, and confirms that the innermost gap feature below 110 mas observed by van Boekel et al. (2017) cannot result from instrumental effects alone. The results of both calibration approaches are found to be consistent down to 75 mas, below which point the normalisation calibration is dominated by diffracted instrumental polarisation artefacts. Observing solar system bodies such as Ceres provides a useful approach for calibrating coronagraphic imaging systems, especially when combined with optical instrument modelling. We also find that for DPI observations, signal losses due to convolutional depolarisation effects are of equal importance to coronagraphic extinction, especially between angular separations of 100-300 mas. This factor should therefore always be accounted for before the detection of an inner disk feature is claimed.

5.1 Introduction

High-contrast, high-resolution scattered light imaging at near-infrared (NIR) wavelengths has in recent years developed into an incredibly versatile tool for observing protoplanetary disks, providing constraints on the disk morphology by allowing us to distinguish spiral arms (e.g., Benisty et al. 2015) and rings and gaps that yield the surface-height profile (de Boer et al. 2016; Ginski et al. 2016) and micron-scale dust composition (Muro-Arena et al. 2018) of the scattering surface of these objects. With ALMA now providing comparable spatial resolution in the sub-mm, it is now also possible to use multi-wavelength studies to place strong constraints on the vertical structure and decoupling between small and large grains due to pressure traps within of these disks (Pinilla et al. 2015), and hence build a more complete picture of the planet-forming environment (e.g. Marino et al. 2015; de Juan Ovelar et al. 2016). Dual-band polarimetric imaging (DPI, see e.g. Kuhn et al. 2001) has in particular become an extremely popular observing mode for scattered-light disk science. This is due to its capacity for effective starlight rejection, whilst also introducing minimal image artefacts when compared with the angular differential imaging (ADI, Marois et al. 2006) family of reduction algorithms, which tend to erase or distort extended disk signal.

Transitional disks are of particular interest for understanding the processes of disk evolution and planet formation. These objects are classified based on the morphology of their spectral energy distribution (SED), which displays a characteristic suppression of the disk's infra-red excess at NIR wavelengths, indicating the absence of the hottest dust components (e.g. Strom et al. 1989; Calvet et al. 2002; Merín et al. 2010). This is interpreted as being caused by dust depletions, or cavities, in the inner regions where disk material is being cleared in an inside-out fashion by radiation pressure from the central star (Clarke et al. 2001; Alexander & Armitage 2007). This interpretation has been confirmed for multiple transition disks (e.g. HD 142527 by Canovas et al. 2013, HD 169142 by Quanz et al. 2013b, and HD 135344B by Garufi et al. 2013).

Observations of such disks have revealed a wide variety of morphological features including multiple rings and gaps (e.g. RX J1615: de Boer et al. 2016, HD 97048: Doucet et al. 2007; Ginski et al. 2016, and HD 141569: Biller et al. 2015; Perrot et al. 2016); spiral structures (e.g. MWC 758: Grady et al. 2013; Benisty et al. 2015, HD 100453: Wagner et al. 2015; Benisty et al. 2017 and HD 135344B: Muto et al. 2012; Garufi et al. 2013); and potential exoplanet accretion sites (e.g. HD 100546: Quanz et al. 2013a, 2015; Currie et al. 2015, LkCa 15: Kraus & Ireland 2012; Sallum et al. 2015, and most recently PDS 70: Keppler et al. 2018). It is important to understand whether such features are indeed indicators of ongoing planet formation, or whether they can be explained by unrelated processes operating within the disk.

Since the existence of a central cavity is the hallmark of the transition process, there is a strong incentive to probe the innermost regions of protoplanetary disks in order to identify such features. However, a nearly all high-contrast imaging observations make use of attenuating coronagraphs to suppress on-axis starlight, which produce distinctive but often poorly characterised signal losses in the vicinity of the inner-working-angle (IWA[1]) of the instrument. Such instrumental features may easily be mistaken for a gap or cavity if not correctly accounted for, especially in low-inclination disks where the two can appear functionally identical. Even when

[1]For observations using a focal-plane coronagraph, this is commonly defined as the angular separation at which the instrument throughput drops to 50% of the wide-angle baseline.

analysing non-coronagraphic polarimetry data, the influence of an extended instrument point-spread function (PSF) can lead to significant, non-linear erasure of disk signal at angular separations interior to a few hundred milli-arcseconds. This "PSF smearing" effect was originally noted in Avenhaus et al. 2014a: here we refer to it as "convolutional depolarisation", and discuss the effect extensively in Sec. 5.4 onwards.

It is therefore desirable to develop a general calibration approach which can accurately recover the relative surface brightness of the disk as close to the IWA limit of the imaging system as possible. In this work we investigate the various instrumental factors which can result in artificial flux attenuation at small angular separations, with a view to calibrating observations made with the Spectro-Polarimetric High-contrast Exoplanet REsearch instrument (SPHERE, Beuzit et al. 2008). Particular attention is paid to the most popular coronagraphic and polarimetric imaging modes of the InfraRed Dual-beam Imager and Spectrograph (IRDIS) subsystem (Dohlen et al. 2008).

To this end we obtain H-band IRDIS coronagraphic observations of Ceres, in order to measure the two-dimensional extinction profile of the SPHERE-IRDIS Apodised-Pupil Lyot Coronagraph (ALC) for a bright, extended target of known flux distribution without the complication of stellar residuals (Sec. 5.2). The data is then used to validate a model of the coronagraphic imaging system, which allows analysis of the various contributions to measured flux as a function of angular separation (Sec. 5.3). A recipe for calibrating coronagraphic disk observations is then developed and validated using the Ceres data and disk models (Sec. 5.4), which includes a discussion of additional effects particular to DPI observations. This calibration procedure is then demonstrated for the case of TW Hydrae (Sec. 5.5), using the coronagraphic H-band DPI observations of van Boekel et al. (2017).

5.2 Ceres observations and data reduction

5.2.1 Observations

Calibration observations of Ceres were taken on the night of the 13th December 2016 as part of a wider effort to characterise the SPHERE NIR coronagraphic system. These observations were performed with IRDIS in H2/H3 dual-band imaging mode ($\lambda_c = 1588.8/1667.1$ nm, $\Delta\lambda = 53.1/55.6$ nm), and using the N_ALC_YJH_S apodised-pupil Lyot coronagraph (ALC) configuration. This mode is particularly useful for a first calibration attempt as it is currently the most commonly used SPHERE coronagraph for disk science observations.

At the time of observation, Ceres displayed an angular diameter of 0.6", a V-band magnitude of +8.3 and a phase angle of 17.6° (JPL horizons database, Giorgini et al. 2015). This makes it a perfect target for coronagraphic calibration as it provides an extended, near-uniform target without the complication of a bright central point source. With a diameter of 0.6", the disk of Ceres extends beyond the 0.2"-radius sphere of influence of the ALC coronagraph, whilst still being significantly smaller than the 1.0" angular resolution of the 40×40-lenslet SPHERE-SAXO Shack-Hartmann wavefront sensor (Sauvage et al. 2014), required for the XAO system to effectively use it as a natural guide star.

Sequences of 8 s exposures were taken in two positional configurations: firstly with the disk of Ceres centred on the coronagraphic axis, and secondly with the target

offset from the ALC mask to provide a non-coronagraphic reference image. The on-axis dataset consisted of five exposure sequences totalling 9.6 minutes of integration time, during which the IRDIS detector was dithered in a 1×1 pixel square grid pattern. The off-axis configuration comprised two sequences totalling 2.1 minutes, at a pointing offset of 0.4"-0.5" with respect to the on-axis dataset. All data were taken in pupil-tracking mode in order to stabilise the coronagraphic optics, resulting in 3° of field rotation during the observing sequence.

Such a method of obtaining the coronagraphic extinction profile is complementary to another approach carried out in parallel, where a disk-less star (i.e. a simple point source) is observed and slowly shifted out from under the coronagraph mask by applying tip/tilt reference offsets via the AO system's Shack-Hartmann wavefront sensor. The extinction profile is then reconstructed from the peak throughput of each image at the corresponding angular separation from the coronagraph axis. The results of this calibration can be found in the SPHERE instrument manual (ESO 2018) for IRDIS H-band and IFS H- to J-band measurements. The two methods are highly complementary, as this "point-source" method provides an estimate of the direct throughput of the coronagraphic system as a function of angular separation from the optical axis. On the other hand, the "extended-source" method using Ceres inherently includes non-linear diffraction and scattering contributions from all resolution elements of the target. The comparison of both approaches therefore highlights those regions where the extinction curve is no longer dominated by direct (useful) throughput for a particular extended target. The point-source method can also be sensitive to temporal variations in image Strehl ratio and errors in positioning accuracy, and provides only a one-dimensional estimate of the coronagraphic extinction profile. By providing a full image of the profile using only two images, the extended-source method is significantly more robust against these effects. Such an approach is also necessary to properly characterise asymmetric coronagraphic optics, such as the wedge coronagraphs of HST-STIS (Heap et al. 2000) and JWST-NIRCam (Krist et al. 2009), which cannot be calibrated using a simple 1-D radial profile.

5.2.2 Data reduction

All Ceres H2/H3-band images were first appropriately dark- and flat-corrected. Each frame was then independently corrected for sky background signal and frame-to-frame readout striping artefacts via column-wise median subtraction, using areas containing no significant Ceres signal. Sky frames were available but were not used in this case, as it was seen that they retained latent image signal from the previous object exposures. While negligible when compared to the total flux of Ceres (less than 0.2%), this latent signal was seen to create minor but noticeable artefacts beyond the edge of the Ceres disk when used in the image reduction process. A mask of the many hot and rogue pixels of the IRDIS detector was however created using the dark/flat-calibrated sky frame data, based on a manually optimised threshold chosen to capture all bad pixels whilst avoiding overly masking additional pixels.

On-axis frames were then shifted to correct for the 1×1 px dither pattern used during the observation sequence, before extracting the H2 and H3 channels from the image. The stability of the ALC focal-plane mask with respect to the IRDIS detector plane is extremely high, making it unnecessary to perform any additional sub-pixel correction when co-aligning the coronagraphic extinction profile of each image. How-

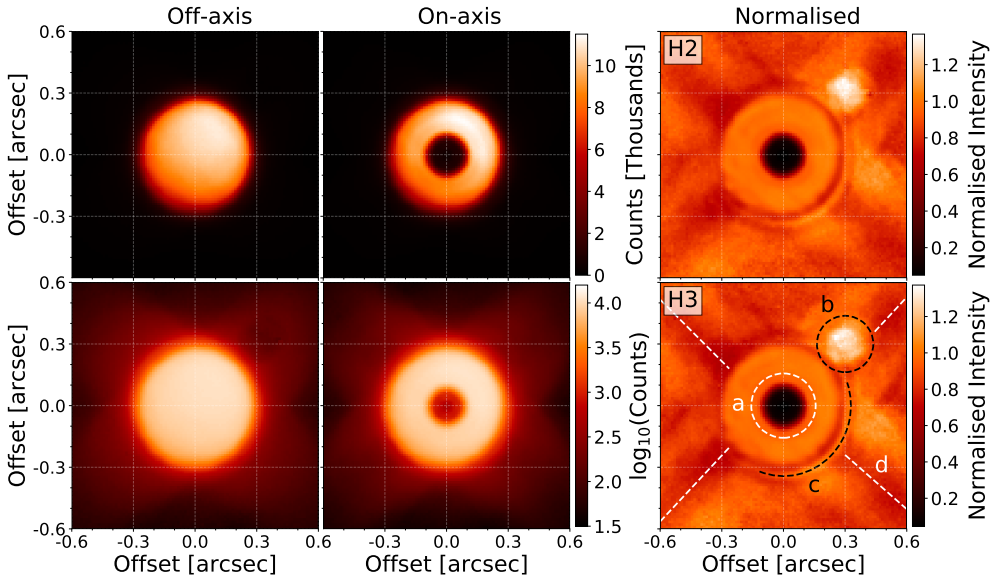


Figure 5.1: Reduced H2-band images of Ceres calibration observations, and the normalised ALC throughput profiles in both the H2 and H3 bands. ***Left column:*** non-coronagraphic images, with the ALC focal-plane mask located an average of 0.42" to the top-right of the Ceres disk edge. ***Middle column:*** coronagraphic observations with the focal-plane mask centered on the Ceres disk. The bottom row shows the log-scaled equivalent of linearly scaled top row, showing the inner regions of the coronagraphic mask and modifications to the diffraction pattern beyond the disk edge. ***Right column:*** linearly-scaled ratios of on-axis and off-axis intensity images, which correct for spatial variations in the Ceres surface brightness profile and are used to measure coronagraphic throughput curves of Fig. 5.2.

ever, due to relatively poor observing conditions there were significant frame-to-frame variations in image quality, which mostly manifested in changes in the shape of the first Airy ring and hence the centroid location of the target with respect to the ALC mask.

To account for these relative positional offsets, the central coordinates of each frame in both the on-axis and off-axis image sets was registered by fitting a model Ceres profile using the Levenberg-Marquardt least-squares fitting routine of the astropy python package. The observed Ceres surface brightness profile was found to be accurately modelled as a spherical Lambertian reflector with the appropriate illumination angle, convolved with a simulated un-occulted ALC instrument PSF. This model was used as the least-squares fitting function, with central coordinates and peak flux allowed to vary as free parameters. The central 175 mas of on-axis images was first masked out, in order to avoid biases due to the coronagraphic extinction profile.

The off-axis image frames were then co-aligned to sub-pixel accuracy and median-stacked, to form a master non-coronagraphic image reference for the Ceres disk. This profile is shown in the left panels of Fig. 5.1. The master off-axis reference image was then separately co-aligned with each on-axis image after correcting for the average field rotation angle of each exposure sequence, then used to individually normalise each on-axis frame. De-rotation of each frame within an exposure sequence, or when co-aligning the off-axis reference images, was found not to be necessary: the low amount of field rotation in this dataset meant that pixels at the edge of the Ceres disk shifted by less than 0.2 pixels during individual exposure sequences.

The final coronagraphic extinction profile image was then obtained via median-

stacking of the normalised dataset: this profile is shown in the right-hand side of Fig. 5.1 for both the H2 and H3 bands. This frame-wise reduction approach provides the optimum normalisation of the edges of the Ceres disk under non-ideal observing conditions, thereby maximising the useful field over which the coronagraphic extinction profile can be analysed without the influence of disk-edge effects.

5.2.3 Analysis

5.2.3.1 Image features

The circular extinction profile of the coronagraph is clearly visible in the center of the Ceres normalised intensity images (feature *a* in the bottom-right panel of Fig. 5.1). Both H2 and H3-band images were seen to be nearly identical, save for some small increase in scale of the diffraction pattern beyond the edge of the Ceres disk, as would be expected for the small change in wavelength between bands. The disk of Ceres is seen to be well corrected for surface brightness variations out to an angular separation of approximately 250 mas, which allows for a straightforward extraction of the full ALC extinction profile.

Numerous secondary features are also visible in the outer regions of these normalised images, which can all be explained with reference to the un-normalised data. Firstly, feature *b* is an artefact of the ALC focal-plane mask blocking part of the spider diffraction structure in the off-axis images (visible in the bottom-left panel of Fig. 5.1). The asymmetric arc-like feature, labelled feature *c*, is attributed to a coma-like quasi-static aberration in the system which varied between the on-axis and off-axis datasets. By changing the morphology of the first Airy ring, this aberration correspondingly alters the brightness distribution of the outer regions of the Ceres disk in the same way as discussed in Sec. 5.2.2. In this case, the effect results in an over-normalisation in the top-left half of the image and a corresponding under-normalisation in the lower-right region. The four dark channels in the spider diffraction structure, labelled feature *d*, is a result of occultation of the central regions of the Ceres disk in the on-axis images, as is visible in the lower-middle panel of Fig. 5.1. This is because the majority of this diffraction structure is produced by the spider mask component of the Lyot stop, which sits downstream of the focal-plane mask and therefore diffracts less flux along the diagonals corresponding to the occulted regions of the disk.

5.2.3.2 Radial extinction profile

The radial coronagraph extinction profile for these observations was extracted after precisely determining the location of the optical axis in the image, since this is not known a priori. Due to the high level of azimuthal uniformity seen in the ALC extinction region normalised images, the centering was refined via an iterative grid search which minimised the variance of the extracted radial profile for angular separations below 200 mas. The resulting extinction curves are shown for both H2 and H3 bands as the solid coloured lines in Fig. 5.2. It can be seen that the coronagraph profile has extremely low variance in the regions of significant extinction (<125 mas), while the standard deviation of the outer regions of the profile which is not influenced by disk-edge effects (<250 mas) remains below the 2% level. We can therefore conclude that the extinction profile of this SPHERE-IRDIS ALC coronagraphic system is az-

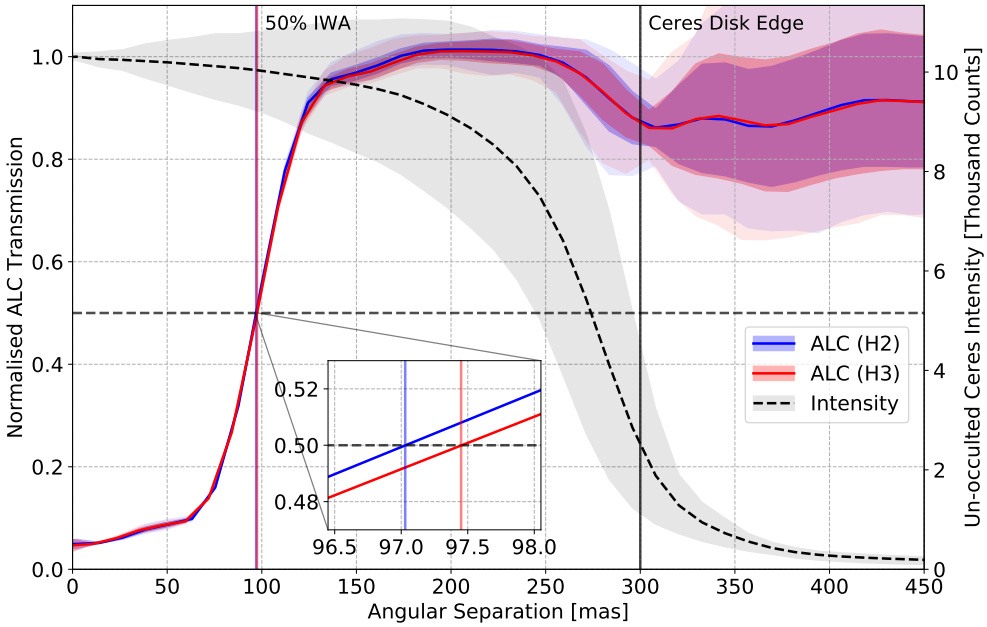


Figure 5.2: Radial ALC coronagraph extinction profile, as measured from Ceres data (right column of Fig. 5.1). Blue and red curves denote the median azimuthally averaged normalised intensity profile for the H2- and H3-bands respectively, with dark shading denoting the 1σ azimuthal variability and light shading denoting the extrema of the data. The grey dashed curve denotes the un-occulted Ceres surface brightness profile (left column of Fig. 5.1), with shading denoting the extrema. The edge of the Ceres disk is located at 300 mas, and the 50% IWA contour is located at approximately 97 mas, denoted by red and blue vertical lines in the inset panel for the H2- and H3-bands, respectively.

imuthally uniform to a high degree, facilitating an accurate calibration using a 1-D radial profile in this case.

The 50% IWA is measured to be 97.0 ± 0.3 mas and 97.5 ± 0.3 mas for H2 and H3 respectively, which is within 2 mas of the value reported by the results of the point-source offset method (ESO 2018). Note that this value is higher than the physical 92.5 mas radius of the coronagraph mask which is often incorrectly taken as the masking radius for analysis of coronagraphic data: this figure clearly shows that the coronagraph is responsible for significant extinction out to approximately 200 mas, which if not correctly normalised may be erroneously claimed as the outermost edge of a central cavity in an equivalent protoplanetary disk observation.

It can be seen that in Fig. 5.2 the profile levels off in the inner regions behind the coronagraph mask (<75 mas), tending to a value much greater than the $< 5 \times 10^{-3}$ on-axis H-band extinction ratio which was measured for this coronagraph during commissioning laboratory testing (Guerri et al. 2011) and on-sky (Fusco et al. 2016). This is behaviour is not seen in direct throughput curve measurements made via the point-source calibration approach, since the additional signal is produced by diffraction from all resolution elements of the target outside the ALC mask. Accurately reproducing this non-linear contribution to the total signal for an extended target is one of the major motivators of the coronagraphic system modelling presented Sec. 5.3, since it ultimately determines the smallest angular separation at which useful astrophysical signal can be recovered.

Additional modelling is also required to determine why the region between 175-

250 mas exceeds a normalised intensity of unity: this should not be possible if the coronagraphic mask is the only source of extinction. It was found however that this can be effectively explained by variations in Strehl ratio between the on- and off-axis images: this consideration is presented in detail in Sec. 5.3.3 with the aid of a coronagraphic system model.

5.3 Modelling the SPHERE ALC Coronagraph System

5.3.1 Description of the coronagraphic system model

It is clear from Sec. 5.2.3 that a model of the ALC system is required in order to fully explain the observed extinction curve, and to facilitate an optimal correction scheme for science images. For this purpose we use a three-plane Fourier propagation model which incorporates the amplitude apodiser, focal-plane mask and Lyot stop designs shown in Fig. 5.3. These planes are based on the comprehensive description of the ALC design and laboratory testing provided by Guerri et al. (2011). The additional structures visible in the Lyot stop pupil plane were added during commissioning of SPHERE, and mask dead actuators on the SAXO deformable mirror (Dohlen et al. 2016).

A coronagraphic image $I(x, y)$ is produced by convolution of the surface brightness distribution of the astrophysical object $O(x, y)$ with the spatially variable PSF $p_{ij}(x, y)$ which describes the attenuation of the coronagraphic system. This can be written in the form of a discrete convolution, as

$$I(x, y) = \sum_{ij} O(i, j) p_{ij}(x, y), \tag{5.1}$$

where the indices i and j refer to the discrete sampling points in the focal plane along the x and y axes respectively, which functionally correspond to pixels in the detector array. In this work, pupil plane coordinates are denoted in (u, v) space while focal-plane coordinates are denoted in (x, y) space.

The individual point-spread functions required for this computation are calculated by re-imaging the respective pupil and focal-planes of the coronagraph, such that

$$p_{ij}(x, y) = \left| \mathscr{F} \left[L(u, v) \, \mathscr{F} \left[m(x, y) \, \mathscr{F} \left[A(u, v) \mathrm{e}^{i\phi_{ij}(u,v)} \right] \right] \right] \right|^2 . \tag{5.2}$$

In this equation $A(u, v)$ denotes the pupil amplitude apodisation function (panel 1 of Fig. 5.3), $m(x, y)$ the focal-plane mask (panel 2), and $L(u, v)$ the downstream Lyot stop (panel 3). The location of the PSF in the focal plane is controlled by adding a tip-tilt wavefront error $\phi_{ij}(u, v)$ to the first pupil plane, which is calculated to exactly match the location of the desired sampling point ij. The Fourier transform operation is denoted by $\mathscr{F}$ [], implemented here using the array-based Fourier transform routine described in Soummer et al. (2007b) which is suited to the efficient computation of Lyot coronagraphic PSFs.

The intermediate focal-plane containing the coronagraphic mask is sampled at a fixed resolution of 5 pixels per λ/D, as a trade-off between speed and resolution when computing a large grid of PSFs, each with a significant final field-of-view. The

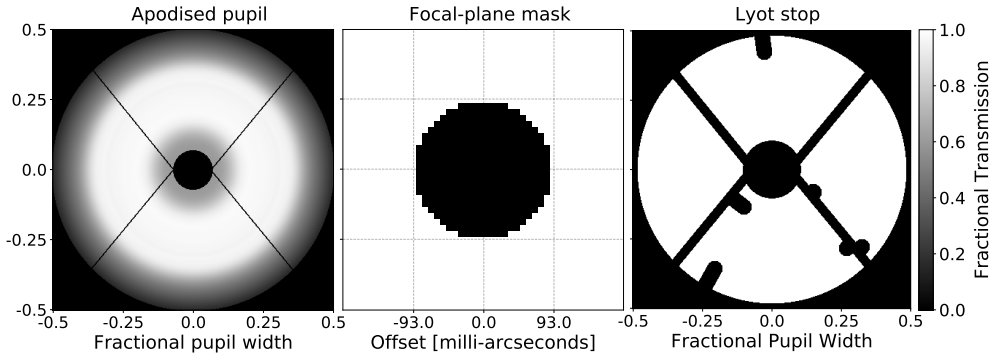


Figure 5.3: Three-plane ALC model used for optical system modelling, matching the SPHERE-IRDIS N_ALC_YJH_S ALC coronagraph configuration (ESO 2018). The VLT pupil geometry and ALC_2 Lyot stop design, including dead actuator masks, is based on Fusco et al. (2014). The APO1 pupil apodisation function is based on the H-band laboratory measurements of Guerri et al. (2011). The focal-plane mask is computed at a resolution of 5 pixels per λ/D.

simulation wavelength can be freely adjusted by varying the effective radius of the focal-plane mask and appropriately scaling the sampling grid of the final image plane. Phase aberrations from sources such as XAO-filtered atmospheric residuals and non-common path errors can also be added to either the upstream or downstream pupil planes, such that $A(u,v)$ and $L(u,v)$ may become complex in Eq. 5.2.

5.3.2 Diffraction-limited model

In order to directly compare with observations, we evaluate our ALC model with an object model $O(x,y)$ emulating the Ceres surface brightness profile at the time of observation. This consists of the spherical, perfectly Lambertian reflector model that was previously used for image registration in Sec. 5.2.2, with a phase angle of $17.6°$. A datacube of diffraction-limited ALC imaging PSFs $p_{ij}(x,y)$ was then computed according to Eq. 5.2 for each resolution element of the IRDIS focal plane where $O(x,y)$ is non-zero, at the H2-band central wavelength of 1588.8 nm. An equivalent un-occulted datacube was also computed by omitting the focal-plane mask and directly computing the Fourier transform squared of $L(u,v)A(u,v)e^{i\phi_{ij}(u,v)}$, to act as a reference for normalisation. The Ceres observations presented in Sec. 5.2 were then emulated by computing the spatially-variable convolution of Eq. 5.1 with both the occulted and un-occulted $p_{ij}(x,y)$ datacubes, and normalising the first resulting image with the second.

Figure 5.4 shows the normalised intensity radial profile of this simulated image model (blue curve), which is directly comparable to the H2-band Ceres observations of Fig. 5.2, shown here as black data points with grey shading. Here the Ceres curve has been normalised to unity between 200-250 mas to properly account for the effects of variable seeing conditions during the observation window: this approach is justified in Sec. 5.3.3 where such factors are explored in detail.

The top panel of Fig. 5.4 shows that this diffraction-limited ALC model provides an excellent fit to the measured Ceres extinction profile, which is accurate to better than 2% for all angular separations above the 50% IWA and better than 5% up to the 92.5 mas focal-plane mask edge. The IWA itself is predicted to be 98.1 ± 0.6 mas, only 1 mas larger than that measured on-sky. This close agreement is especially impressive given that we use here the simplest possible forward model of the optical system,

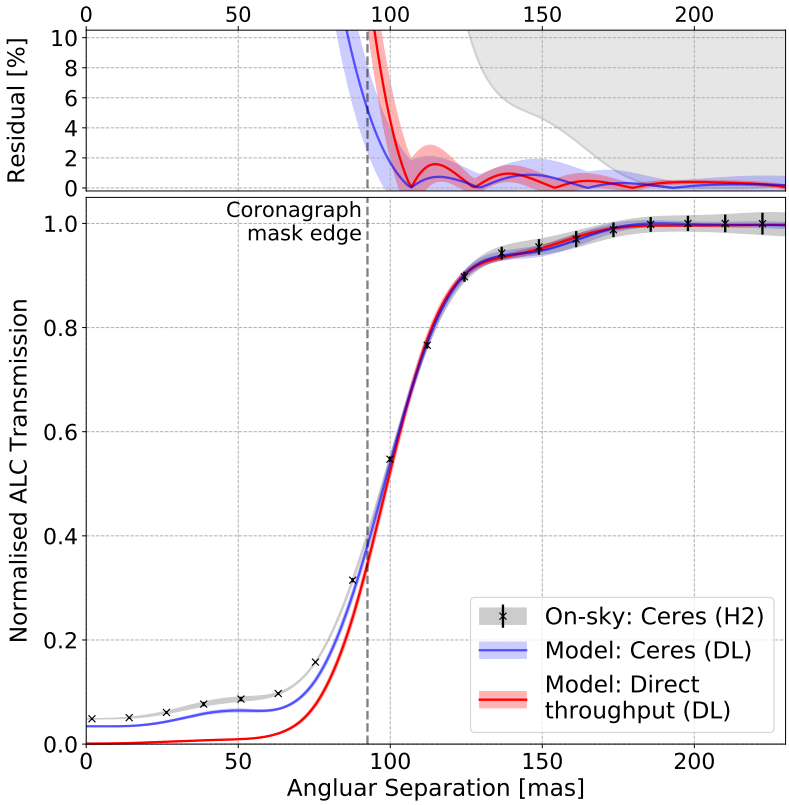


Figure 5.4: Comparison of the diffraction-limited (DL) ALC imaging simulation with the on-sky Ceres H2-band extinction profile measurement of Fig. 5.2. *Bottom:* Radial extinction profiles for on-sky Ceres data (black), simulated image of a Ceres-like object model (blue), and simulated direct throughput curve of the ALC system model (red) as a function of angular separation from the optical axis. *Top:* Absolute percentage residuals between both models and the Ceres on-sky data. Grey shading denotes the region in which systematic residuals are worse than the null hypothesis, where coronagraphic extinction effects are ignored.

which is fully constrained by the literature description of the IRDIS ALC coronagraph and hence has not been fitted to the Ceres observations at any stage. However, this model cannot fully account for all flux seen in the inner regions of the measured profile, which leads to an rapid increase in fractional error inside the IWA due to the small transmission values involved. This discrepancy could be due to a variety of instrumental effects such as non-common-path errors (NCPEs), and is addressed in Sec. 5.3.4.

Also plotted in Fig. 5.4 is the direct throughput of the coronagraphic system (red curve), measured as the relative transmission of pixel i, j for each coronagraphic PSF $p_{ij}(x, y)$ by analogy with the point-source calibration method discussed in Sec. 5.2.1. This direct throughput curve can be seen to tend towards zero at the coronagraphic axis, and is systematically lower than the full Ceres image model (blue curve) below approximately 110 mas. The difference between these two curves is simply the integrated signal that is diffracted inwards by the Lyot stop from un-occulted regions of the disk. It is noteworthy that there is still significant direct throughput for separations interior to the 92.5 mas edge of the focal-plane ALC mask. This would imply that it is in principle possible to extract useful flux information down to angular separations of

approximately 60 mas for this coronagraph, however in reality this will almost always prove impossible due to the dominance of speckle noise or instrumental polarisation artefacts at these small angular separations.

It is this direct throughput metric which constituents the useful throughput of the coronagraphic system, since it is the fraction of signal that the optical system correctly maps from object $O(i, j)$ to image $I(i, j)$, and hence provides spatially-resolved information about the target. The most straightforward calibration approach is therefore to normalise the central regions of all disk observations using the radial direct-throughput curve of the ALC coronagraph presented here, to correct for the relative useful throughput at all angular separations. As can be seen from Fig. 5.4 however, such an approach will result in a systematic flux overestimate by a factor corresponding to the fractional diffracted-light contribution at each pixel. This leads to a substantial error at small angular separations, which for the case of Ceres is equal to the red curve in the upper panel of Fig. 5.4. This issue is further complicated by the fact that the relative contribution of non-useful diffracted signal is inherently source-dependent, based on the surface brightness profile of the object under observation. This makes it impossible to determine a single limiting angular separation above which the accuracy of the calibrated signal can be trusted for all observations using the same coronagraph. A proper calibration for these innermost regions is therefore only possible by using an accurate forward model of the coronagraphic imaging system, which can determine this on a target-by-target basis. To this end we extend the diffraction-limited model in the following subsections to consider the impact of various instrumental factors on the extinction profile, in order to try to minimise the discrepancy between our model and the Ceres observations.

5.3.3 Effects of residual atmospheric turbulence

The most obvious factor affecting on-sky images is the presence of residual atmospheric turbulence after correction by the SPHERE XAO system (SAXO, Fusco et al. 2016). The reason for modelling this is two-fold: to assess the impact of residual phase errors on the inner regions of the extinction profile, and to explain the morphology of the outer regions of the measured Ceres curve as noted in 5.2.3.2.

For this work, residual XAO-corrected phase-screens were generated using SPHERE-like AO simulation code originally developed in Wilby et al. (2016b), based on a Kolmogorov spatial power spectrum which is suppressed for frequencies controllable by the SAXO 41×41 actuator deformable mirror. The final functional form of this XAO-corrected spatial power spectrum Φ_{AO} was chosen to be

$$\Phi_{\mathrm{AO}}(k, r_0) = \begin{cases} 5 \times 10^{-3}\Phi_{\mathrm{Kol}}(k, r_0) + 0.5\Phi_{\mathrm{Kol}}(k_{\mathrm{crit}}, r_0) & k < k_{\mathrm{crit}} \\ \Phi_{\mathrm{Kol}}(k, r_0) & k \geqslant k_{\mathrm{crit}} \end{cases} \tag{5.3}$$

where Φ_{Kol} is the Kolmogorov phase power spectrum as a function of spatial frequency k (in units of cycles per pupil, or equivalently λ/D) and Fried parameter r_0. The spatial frequency cutoff k_{crit} is set to be $20\ \lambda/D$, matching the control radius of the SAXO deformable mirror: beyond this cutoff the power spectrum is identical to Φ_{Kol}. The resulting AO-corrected power spectrum is shown in the top-left panel of Fig. 5.5 (solid line), compared to the equivalent Kolmogorov power spectrum Φ_{Kol} (dashed line). Variable seeing conditions can be emulated using this model by simply increasing or decreas-

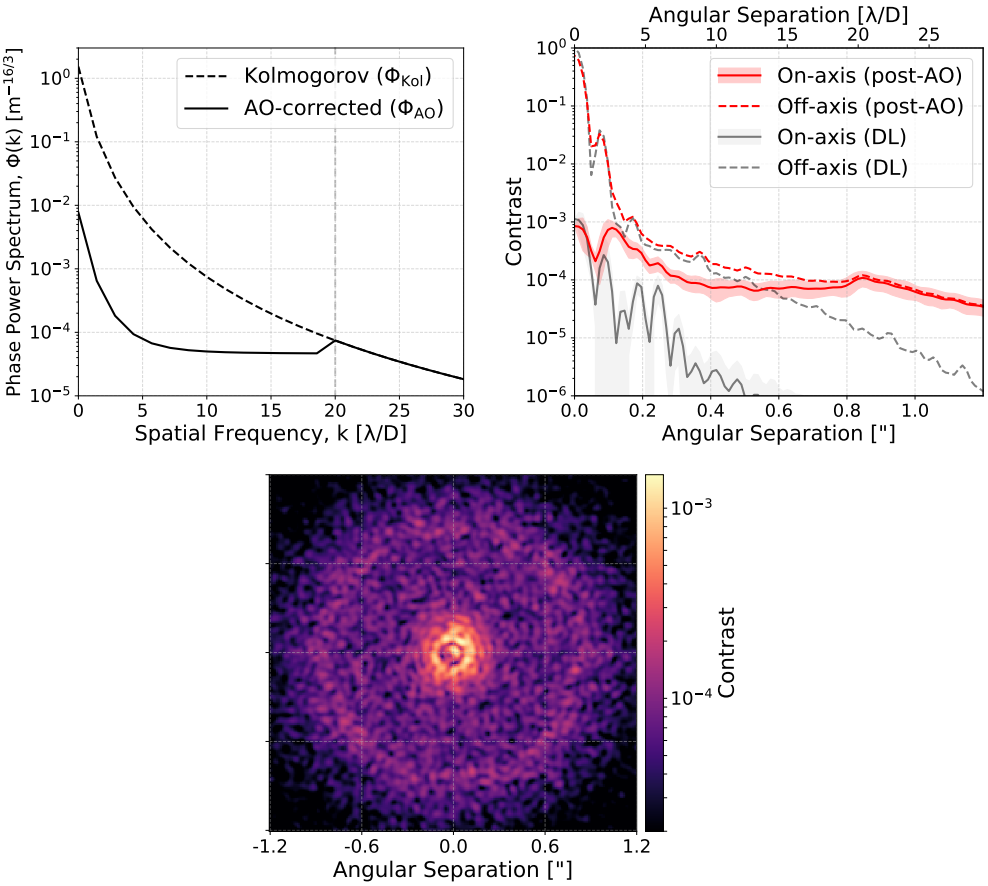


Figure 5.5: Key modelling parameters for SPHERE-like residual atmospheric phase errors. *Top left:* Power spectral density of XAO-corrected residual phase aberrations corresponding to Eq. 5.3 (solid line), with respect to purely Kolmogorov turbulence (dashed). *Top right:* The resulting raw coronagraphic contrast curve (on-axis, solid red line) and un-occulted PSF (off-axis, dashed red line), compared to the diffraction-limited case (DL, grey).*Bottom:* Corresponding on-axis coronagraphic PSF, comprising the sum of ten individual speckle realisations.

ing r_0, and hence the overall image Strehl ratio, whilst keeping the AO spatial filtering function constant.

The functional form of this AO-corrected atmospheric spatial power spectrum was chosen to allow for independent modification of the fractional suppression amplitude of the Kolmogorov spectrum, and a contrast saturation threshold which is reached at high spatial frequencies. These parameters are tailored in addition to r_0 in order to reproduce the typical on-sky ALC contrast profile under median seeing conditions, as reported by Fusco et al. (2016). Replication of the exact system response of SAXO is not crucial to this investigation, which simply seeks to probe how SPHERE-like post-XAO atmospheric phase residuals qualitatively affect the coronagraphic extinction curve measurements of Sec. 5.2: this approximate model is sufficiently realistic for this purpose.

An example short-exposure PSF comprising an average of ten speckle-field realisations is shown in the bottom panel, in which both the 20 λ/D control radius and elevated low-frequency residual speckles can be seen, matching the main features of

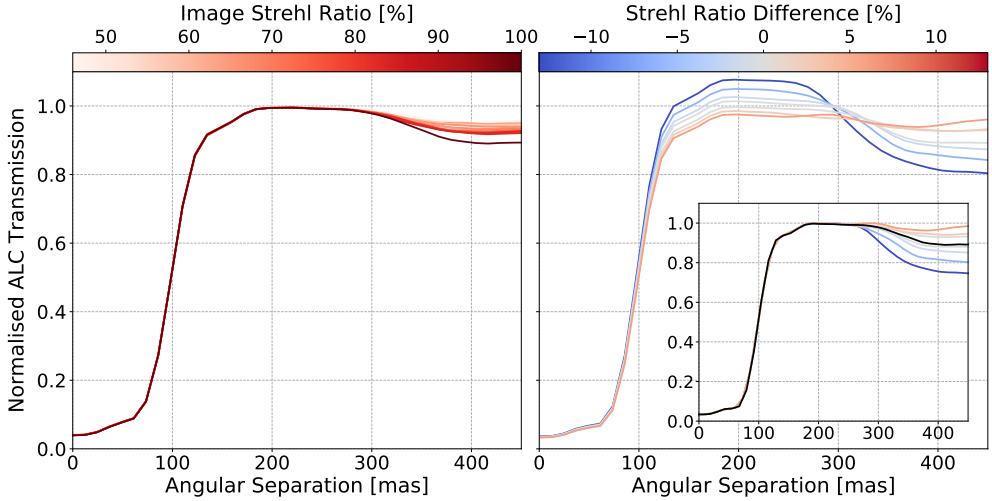


Figure 5.6: Simulation of the effect of varying Strehl ratio on the retrieved normalised intensity profile from Ceres observations. ***Left:*** Both on-axis and off-axis images are convolved with the same p_{ij}, and hence identical r_0 and seeing conditions. The darkest line corresponds to the diffraction-limited case. ***Right:*** On-axis image has a fixed Strehl ratio of approximately 75%, while the Strehl ratio of the off-axis image allowed to vary. A positive Strehl ratio difference corresponds to the off-axis image having a higher Strehl ratio. The inset panel shows the same data, normalised to unity between 200-250 mas. The diffraction-limited result is over-plotted in black, for comparison.

the chosen spatial power spectrum. The corresponding azimuthally-averaged contrast curve is shown in the top-right panel of Fig. 5.5 (solid red line), with the upper limit of the red shading denoting the 1σ sensitivity limit. Here the red dashed line shows the un-occulted (off-axis) PSF, while the grey curves show the equivalent diffraction-limited instrument performance for comparison purposes.

Figure 5.6 summarises the impact of these residual atmospheric phase aberrations on the coronagraphic extinction profile: the left-hand panel shows the simulated extinction profiles for various values of r_0, obtained in the same manner as in Sec. 5.3.2, where both the on-axis and off-axis images are produced under identical seeing conditions. This clearly shows that the inner coronagraphic extinction profile is not inherently dependent on residual atmospheric phase errors, with the only influence being minor variations beyond the 300 mas edge of the Ceres disk. This result can be anticipated, since the extinction profile in the science image is produced by near-diffraction-limited re-imaging of the coronagraphic focal plane mask itself. The morphology of this extinction profile is therefore dominated by diffraction from the oversized Lyot stop, making it insensitive to all phase errors upstream of the coronagraphic mask.

As the right-hand panel of Fig. 5.6 shows however, noticeable measurement artefacts are introduced when attempting to characterise the coronagraphic profile if even small variations in image Strehl ratio occur between the acquisition of the on-axis and off-axis datasets. These curves were computed by fixing the Strehl ratio of the on-axis image at 75%, and computing off-axis images for a variety of different r_0 values and hence image Strehl ratios. An average of ten instantaneous phase-screen realisations was used to produce each image, a number which was chosen because it provides some speckle averaging necessary to obtain a reliable estimate of image Strehl ratio for a given r_0, while still allowing relatively fast computation of the full spatial grid of

aberrated PSFs. It was not necessary for this investigation to perform the extremely computationally expensive task of simulating all $p_{ij}(x, y, t)$ required in order to properly model the speckle statistics for a single $8\ s$ Ceres exposure: for a speckle coherence time of 1 ms this would require the computation of approximately 10^7 individual PSFs for each simulated Ceres image.

The best qualitative match to the measured Ceres normalised intensity curves of Fig. 5.2 is obtained for an off-axis image Strehl ratio only 1% lower than that of the on-axis image. Such a small change can easily be expected to have occurred during the one-hour Ceres observing window due to natural changes in seeing conditions. This differential image Strehl ratio effect naturally results in a profile that exceeds unity by a few percent between 200 to 250 mas, before dropping outside the disk edge due to over-normalisation by additional scattered light in the off-axis image at these separations. The quantitative agreement between these seeing simulations and the Ceres observations is however mediocre: this is to be expected given the simplified AO model used in this work coupled with the lack of inclusion of additional factors such as NCPEs, which as noted in Sec. 5.2.3.1 are most likely responsible for specific features seen at the edge of the Ceres disk.

Most importantly for this work however, there is a simple solution for correcting this imaging artefact: the inset panel in the right-hand side of Fig. 5.6 shows the same data as the main panel, simply normalised to unity by the flattest region of the profile between 200-250 mas. This normalisation correction neatly recovers the diffraction-limited profile below angular separations of 250 mas, independent of all seeing artefacts. The same correction is therefore applied to the Ceres extinction profile measurements from Fig. 5.4 onwards, in order to accurately estimate the true coronagraphic extinction profile and fairly compare with simulations.

5.3.4 Additional considerations

In order to properly account for the discrepancy between Ceres observations and the diffraction-limited imaging model shown in Fig. 5.4, it is necessary to include one or more additional effects internal to the instrument which contribute additional signal underneath the coronagraph mask. Physical effects which potentially fulfil this requirement include the presence of non-common path aberrations downstream of the focal-plane mask, a mis-alignment of the ALC focal-plane mask itself along the optical axis, or the high-frequency vibration of downstream optics: each of these cases are explored in detail in the following subsections. Low levels of instrumental scattered light can also produce noticeable signal where ALC extinction is greatest: this last effect is however difficult to model as it is inherently dependent on the full description of reflective surfaces in the instrument, making its treatment beyond the scope of this paper.

Since it is difficult to comment on the various scenarios without additional information, we perform numerous simulations to estimate the magnitude of each effect which is required to address the observed discrepancy. This is then compared to SPHERE performance budgets, and hence assess the likelihood of each. Fig. 5.7 shows the best fit models to Ceres observations for each of the three considered effects, based on minimised least-squares residuals for angular separations above 60 mas. All can be seen to make significant improvements over the diffraction-limited case.

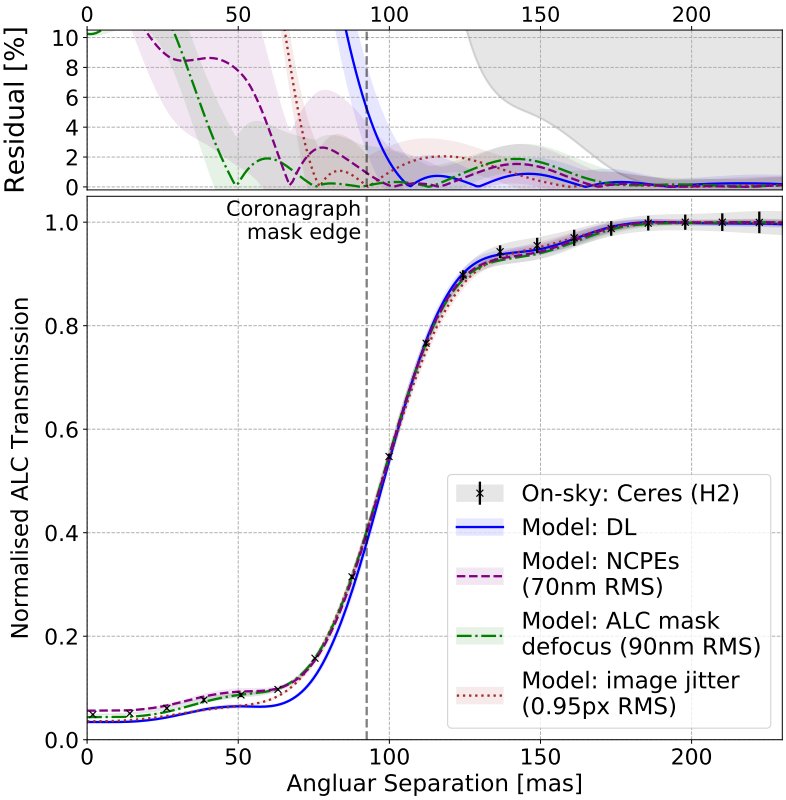


Figure 5.7: Comparison of the various non-diffraction-limited models presented in Sec. 5.3.4, compared to the on-sky Ceres extinction profile measurement. ***Main panel:*** ALC extinction profiles for H2-band Ceres measurements (black data points) and diffraction-limited image model (DL, solid blue curve) of Fig. 5.4. Also plotted are the best-fit image models which include the effects of non-common path errors (NCPE, dashed purple curve), ALC focal-plane mask defocus (dot-dashed green curve), and high-speed vibrational image jitter on the detector (dotted brown curve). ***Top panel:*** Percentage residuals between the aforementioned models and the measured Ceres extinction profile. Grey shading denotes the region in which systematic residuals are worse than the null hypothesis, where coronagraphic extinction effects are ignored.

5.3.4.1 Non-common path errors

Non-common path errors are present to some extent in all high-contrast imaging data, and contribute additional speckle noise which is strongest at small angular separations. If these aberrations occur downstream of the ALC focal-plane mask, re-imaging through the Lyot plane is no longer diffraction-limited and will correspondingly result in less efficient coronagraphic extinction. For SPHERE, these aberrations are typically estimated to contribute an average wavefront error of approximately 40 nm RMS (Fusco et al. 2016), which is consistent with the baseline requirement of 43 nm RMS for high-performance H-band operation. In addition, offline phase-diversity compensation routines have been measured to further reduce these NCPEs to 20 nm RMS using the internal calibration source (Sauvage et al. 2014). This performance budget does not however include the low-wind effect (Sauvage et al. 2016) which, while not itself due to NCPEs, significantly elevates image speckle noise under certain observing conditions.

We model these NPCEs by drawing randomised coefficients for the first 200 Zernike modes, with the resulting wavefront phase maps then spatially filtered in the Fourier domain to have a $1/f^2$-decreasing spatial power spectrum, as is often used to model NCPEs (Sauvage et al. 2007; Lamb et al. 2016). These randomised low-order phase aberrations are then injected in the downstream Lyot stop pupil plane, with five individual realisations being combined for each wavefront RMS value in order to somewhat average out speckle noise.

The best fit to Ceres observations was found to be for 70 nm RMS of $1/f^2$ NCPE, which provides a fit quality of better than 2% down to the 92.5 mas ALC mask edge, and better than 5% down to 60 mas. This is therefore significantly more accurate than the diffraction-limited model, being able to accurately account for the observed Ceres signal at all angular separations where the throughput of the ALC system is non-negligible. It should be noted that by requiring 70 nm RMS of NCPE to adequately explain the discrepancy, this scenario exceeds the 40 nm RMS performance tolerance reported during SPHERE calibration and testing. Such a 30 nm RMS increase in wavefront error is however not inconceivable, especially given that the presence of relatively strong, NCPE-like low-order aberrations was already noted during the reduction of Ceres data in Sec. 5.2.

By looking closely at the main panel of Fig. 5.7, it can be seen that the very innermost region of the profile below 60 mas is also slightly over-estimated compared to the region from 60 to 100 mas: this indicates that a $1/f^2$ spatial power spectrum may be marginally steeper than the optimal model for SPHERE NCPEs in this case, if it is indeed the underlying source of the signal discrepancy. While more complex NCPE models are likely to provide a significantly improved fit to the observed Ceres profile, the inclusion of more than a single additional free parameter is not justified here given the limited calibration data provided by a single measurement of the extinction profile. Most importantly, there is currently no information regarding the temporal variability of the ALC extinction profile, which makes it impossible to constrain a useful time-dependent NCPE model at this stage.

5.3.4.2 Out-of-focus ALC focal-plane mask

If the ALC focal-plane mask is accidentally mis-aligned along the optical axis, it will be out-of-focus with respect to the coronagraphic focal plane. This will increase transmission for resolution elements near the optical axis, because the mask now occults a marginally defocussed PSF which has more power at larger angular separations, hence a larger fraction lies outside the mask edge and is transmitted. Equally, resolution elements more than a diffraction-width from the mask edge will see a negligible change in throughput, making this an effect which only modifies the innermost regions of the extinction profile. The SPHERE longitudinal alignment tolerance for the focal-plane mask is ± 1 mm, which is equivalent to 23 nm RMS of defocus aberration in the coronagraphic focal-plane (Guerri et al. 2011).

A shift of the ALC mask along the optical axis is modelled here by applying the corresponding focus error in the upstream apodiser pupil plane, and then removing it in the downstream Lyot stop pupil plane by applying the opposite phase. As can be seen in Fig. 5.7 the best model for this scenario provides an excellent fit to the observed Ceres extinction curve, with an accuracy of better than 2% down to angular separations of 40 mas. However, the 90 nm RMS of focus error required to produce

this excellent fit corresponds to a 4 mm physical translation of the focal-plane mask from the nominal position (computed according to the optical description of Guerri et al. 2011), which is four times larger than the quoted alignment tolerance of this optic. It would be necessary to check the alignment of the SPHERE NIR coronagraph wheel in order to determine whether this is in fact a realistic scenario: such a mask defocus could easily be overlooked during instrument calibration, as it has a negligible effect on the extinction of an on-axis point source. The quality of fit which this model provides using only a single free parameter, combined with the fact that it is a time-independent effect, would make it an important factor to include in the optical system model were it shown to be the underlying cause of the discrepancy.

5.3.4.3 Image positional jitter

If optics downstream of the ALC focal-plane mask undergo mechanical vibrations or deformations on timescales of less than or equal to the exposure time, the coronagraphic image will jitter with respect to the detector plane, blurring the extinction profile by a corresponding amount. This effect is modelled here by shifting the diffraction-limited Ceres coronagraphic extinction profile model, using Gaussian-distributed displacement amplitudes in random directions to 1/50th of a pixel precision. An average of 200 individual realisations are combined for each RMS value, which when considering a single 8 s Ceres exposure corresponds a characteristic jitter frequency of 25 Hz.

Figure 5.7 shows that the best-fit for this jitter scenario still provides some improvement over the diffraction-limited case, although it is not as successful as either of the preceding effects in accounting for the overall morphology of the extinction profile. In particular, this positional jitter model cannot explain the observed flux excess for angular separations below approximately 75 mas. It should be noted that the best-fit model requires there to be a 0.95 pixel RMS jitter amplitude between the coronagraphic focal-plane mask and the detector. Such a significant instability in the imaging system would produce a noticeable degradation of the non-coronagraphic instrument PSF and, as noted in Sec. 5.2.2, there is also no indication of such displacements to the coronagraphic extinction profile on timescales similar to the 8 s exposure time in raw Ceres data frames: frame-to-frame positional variations of as little as 1/10th of a pixel RMS would already be noticeable by eye. It is therefore highly unlikely that either high-frequency instrumental jitter, or errors in the co-alignment and stacking of the Ceres frames during data reduction, are responsible for the observed discrepancy between the measured and modelled ALC extinction profiles.

5.3.4.4 Discussion and choice of final model

We can conclude from the above discussion that either NCPEs or an out-of-focus ALC mask can individually explain the observed discrepancy with Ceres observations to the precision required for accurate calibration of the coronagraphic system. However, this is only possible if these effects are allowed to significantly exceed their respective SPHERE performance tolerances. This naturally leads to a more likely alternative solution, where a combination of NCPEs and mask defocus can be invoked to accurately reproduce the observed extinction curve whilst remaining consistent with their respective tolerance limits.

However, for the remainder of this investigation we shall adopt the 90 nm RMS (4 mm) ALC focal-plane mask shift error model. Not only does this provide the best fit to on-sky data, but being a single-parameter and time-invariant effect it is also the simplest of the two models, which if correct is applicable to all observations. Since the two effects were found to be more-or-less identical in their effect on the ALC extinction curve, it is not helpful in this instance to construct a complex multi-parameter model which over-fits the available data. Most importantly, with only one on-sky dataset available there is no information on the time-variability of the extinction profile, making it difficult to invoke a NCPE model without first assessing its impact under changing instrument conditions, especially for observations where active NCPE compensation routines are used.

In order to build a more robust model for future calibration efforts, we strongly advise that a measurement of the full coronagraphic extinction profile be made as part of the standard calibration procedure for coronagraphic instruments. Ideally this measurement should be repeated over multiple nights in conjunction with NCPE wavefront sensor measurements, to determine the stability of the extinction profile and the underlying source of any degrading effects.

5.4 Calibrating Protoplanetary Disk Observations

The approach of coronagraphic extinction profile calibration will perform most accurately on polarimetry data, which contain significantly fewer stellar residuals and imaging artefacts than intensity data reduction techniques and is highly favoured for studying protoplanetary disks. When imaging disks using the dual-band polarimetric imaging (DPI) mode of SPHERE-IRDIS, there are however some additional caveats which are not applicable to the Ceres observations discussed up until this point.

5.4.1 Polarimetric image formation and instrumental polarisation

Polarimetric imaging artefacts predominately arise from differences in the image formation process for DPI observations, where the polarised intensity (PI) image is reconstructed based on sets of $Q_{\pm}$ and $U_{\pm}$ images, in which a linear polariser is angled at $0/90^{\circ}$ and $45/135^{\circ}$ respectively with respect to the image plane. A practical overview of PDI-mode imaging can be found in de Boer et al. (2017), while the core idealised equations for this reconstruction process are summarised below for convenience:

$$I = (Q_{+} + Q_{-} + U_{+} + U_{-})/2, \tag{5.4}$$

$$Q = Q_{+} - Q_{-}, \tag{5.5}$$

$$U = U_{+} - U_{-}, \text{and} \tag{5.6}$$

$$PI = \sqrt{Q^2 + U^2}. \tag{5.7}$$

The first caveat for DPI observations is that sources of instrumental polarisation, combined with the presence of a bright central source, often create significant spurious signals in PI images, particularly at small angular separations. A common approach

to overcome this limitation is to transform this polarisation signal into a polar coordinate system:

$$Q_\phi = Q\cos(2\theta) + U\sin(2\theta) \tag{5.8}$$
$$U_\phi = Q\sin(2\theta) - U\cos(2\theta). \tag{5.9}$$

It is then possible to analyse only the azimuthal polarisation contribution (positive Q_ϕ) since this contains the majority of astrophysical signal, while any residual U_ϕ signal is assumed to be instrumental in nature (Avenhaus et al. 2014b).

A more robust approach to calibrate for this instrumental polarisation is to apply a detailed model of the polarimetric crosstalk properties of the imaging system, in order to directly simulate and remove the residual instrumental polarisation contribution: such a model is provided for VLT-SPHERE by van Holstein (2016) and is applied in Sec. 5.5.

5.4.2 Convolutional depolarisation: an additional consideration for DPI-mode observations

For DPI-mode observations, the coronagraph is also no longer the only source of suppression for linearly polarised disk flux at small inner-working angles. The "butterfly polarisation erasure" effect (Heikamp et al., in prep.) results in an artificially depolarised hole in the center of the image, due to the polarimetric image formation process using a finite PSF.

Following Heikamp et al. (in prep.), the cause and impact of this polarisation erasure effect is illustrated in Fig. 5.8 for a constant-intensity, perfectly azimuthally polarised disk model with no central cavity. Polarimetric differential imaging of such a target in Stokes Q and U produces so-called butterfly patterns (lower panels), which display a characteristic quadrupole signal which tends towards a singularity at the stellar position. When this disk is imaged with a finite instrument PSF, convolutional blurring of the butterfly pattern mixes the positive and negative regions, an effect which is particularly pronounced within a few λ/D of the central singularity. This results in depolarisation and hence loss of disk signal in the total linearly polarised intensity (PI) image, creating the illusion that there is an inner cavity in the disk. The bottom panel of Fig. 5.8 shows that this convolutional polarisation erasure effect (green curve) has a significant impact on the retrieved radial profile out to angular separations as large as 1 arcsec, for the simplest case of a uniform disk of infinite extent. Its contribution to total disk signal loss is also larger than that of the coronagraph alone (blue curve) for separations below 100 mas, with the profile produced by coronagraphic DPI-mode observations (red curve) affected by both sources of extinction. Here the term "un-occulted" is used to refer to observations where the ALC focal-plane mask is not aligned with the target, in the same way as the off-axis Ceres images of Fig. 5.1. It should be noted that the convolutional polarisation erasure effect will be somewhat lower for observing modes which do not use the ALC coronagraphic system, since the Lyot mask and apodiser, which are still included in the green curve of Fig. 5.8, result in a significantly broader PSF and elevated diffraction structure when compared with the PSF of the VLT aperture alone.

The preceding discussion serves to emphasise that attempting to normalise DPI-mode disk observations with a coronagraphic throughput curve derived from intensity

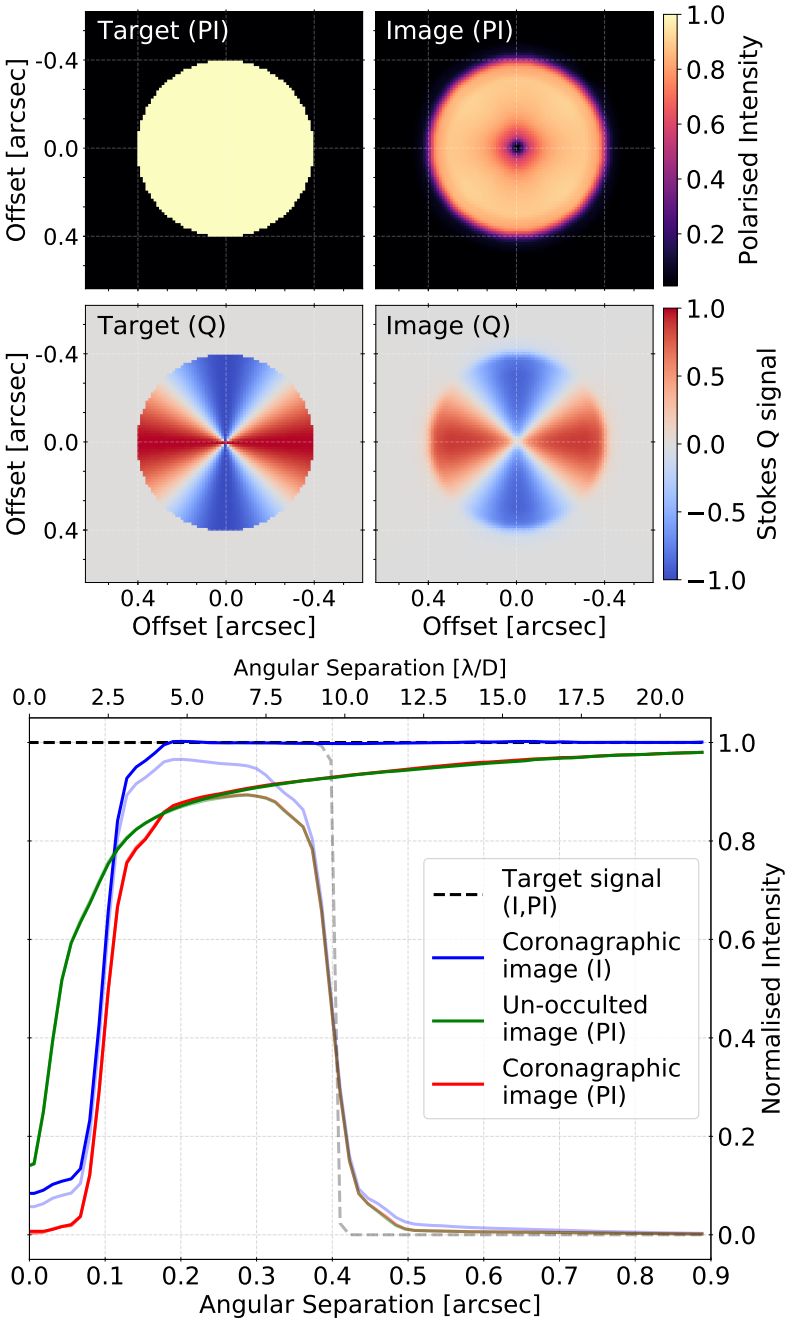


Figure 5.8: Impact of convolutional polarisation erasure on DPI-mode observations, under diffraction-limited VLT observing conditions. ***Top panels:*** Simulation of the polarimetric image formation process for a finite, uniformly 100% polarised disk target (upper left). The perfect Stokes Q (lower left) and convolved Stokes Q image (lower right) show a quadrupole butterfly pattern, the blurring of which leads to a depolarisation of the innermost regions of the polarised intensity (PI) image (upper right) when combined with the companion Stokes U image. ***Bottom panel:*** Radial intensity plot of signal extinction contributions for the various imaging modes. Faded lines denote the case of the finite 400 mas radius disk of constant-intensity shown in the upper panels, while opaque lines show the simpler case of a disk of infinite extent. The un-occulted PI image (green) is solely affected by polarisation erasure, while the coronagraphic PI image (red) is affected by both sources of extinction.

imaging alone (such as the red curve of Fig. 5.4) will still result in a systematic underestimate of the true surface brightness of the target. Due to the relative strength of this effect at larger radii, it should be properly accounted for not just when trying to recover signal in the innermost few hundred mas, but whenever there is significant disk signal inside 1 arcsec. Since the effect is inherently source-dependent, the only robust way do this is by forward modelling the disk in question through the raw $Q_{\pm}$ and $U_{\pm}$ images in which the effect is produced.

5.4.3 A recipe for disk calibration

The most straightforward approach to calibrating disk observations for coronagraphic extinction is to simply divide out the measured direct throughput of the coronagraphic system as a function of angular separation. Via this method it is in principle possible to correct for the influence of the coronagraph beyond approximately 100 mas to an accuracy of better than 2%, although the location of this radial limit depends somewhat on the flux distribution of the target in question.

As has been shown in Fig. 5.4 and Fig. 5.8 however, this normalisation approach will not perform adequately when the observer is either interested in accurately recovering the relative photometry of angular separations below approximately 100 mas, or is reducing DPI-mode polarimetric observations with significant disk signal interior to approximately 1 arcsecond. These factors, due to diffracted light and convolutional polarisation erasure respectively, are inherently dependent upon the flux distribution of the disk itself. Because of this there is no single profile which is capable of properly calibrating all coronagraphic observations via simple normalisation.

Instead, the most robust approach is to forward-model the disk in question, using the best model of the spatially-variable coronagraphic PSF from Sec. 5.3.4 with the polarimetric image formation procedure of Sec. 5.4.2, and hence properly account for both effects. This approach can be explicitly formulated in terms of the following steps:

1) Identify the centre of the focal-plane mask, which defines the zero-point for coronagraphic extinction calibration. This is best achieved using a flat-field reference image taken with the mask in the beam, but if not available it may be possible to estimate this from the inner regions of the disk image assuming no obvious asymmetric structures are present.

2) Generate coronagraphic instrument PSFs $p_{ij}(x, y)$ via Eq. 5.2 at the appropriate wavelength band centre for each resolution element of interest, with the simulated coronagraphic mask correctly centred with respect to step 1.

3) Identify the stellar position to define the origin for polarisation erasure calibration. This is ideally obtained using a centring frame, where artificial satellite copies of the star are created via the AO system in order to identify the location of the star under the coronagraphic mask.

4) Create an object model, preferably with as few free parameters as possible, which is capable of accurately describing the disk regions of interest. For inclined disks or those possessing complex features a 2-D or full 3-D model is most likely necessary, but for simpler face-on systems a 1-D radial model is sufficient. This is then projected and/or interpolated to create a 2-D model signal map, which constitutes a “perfect” observation of the target.

5) If reducing polarimetry data, create model $Q\pm$ and $U\pm$ frames: this is most simply achieved by assuming perfect azimuthal polarisation of the disk signal with respect to the stellar position from step 3. More realistic polarisation models including factors such as the angular-dependence of polarisation efficiency by scattering (e.g. Pinte et al. 2009) can also be performed in this step, to naturally account for such factors during the reconstruction procedure.

6) Propagate the resulting 2-D model signal map(s) through the imaging system by element-wise convolution with $p_{ij}(x, y)$ according to Eq.s 5.1, and perform Q_ϕ image reconstruction using Eq.s 5.4-5.9 if appropriate.

7) Apply a (preferably global) optimisation routine to derive the most likely disk model for the given observations, passing each iteration through the image formation process of steps 4 - 6.

Deconvolution algorithms (such as the Richardson-Lucy algorithm, Richardson 1972; Lucy 1974) could also in principle be used to invert the problem and directly estimate an object model based on the image data and full instrument model. However, most such algorithms are not capable of operating with a spatially variable PSF and hence coronagraphic observations. They are also prone to generating unphysical artefacts, including "ringing", due to image noise and small errors in the instrumental PSF model, given the inherent degeneracy of the convolution operation itself (see e.g. Mosleh et al. 2014).

In the following section we re-reduce the DPI-mode observations TW Hydrae presented in van Boekel et al. (2017), in order to illustrate the approach outlined above for calibrating the inner regions of protoplanetary disks.

5.5 TW Hydrae: calibration of a real disk observation

5.5.1 Background and observations

TW Hydrae is the closest T-Tauri star to the sun, located at a distance of $59.5^{+0.96}_{-0.93}$ pc (Gaia Collaboration et al. 2016a). It hosts a nearly face-on (i≈ 7°, Qi et al. 2004) transition disk with a very high degree of radial symmetry, which is particularly gas-rich for its estimated age (3-15 Myr, Sokal et al. 2018). Due to its close proximity to Earth, this object has been the subject of extensive study at infra-red through to sub-mm wavelengths, most notably with SPHERE (e.g. van Boekel et al. 2017, hereafter vB17) and ALMA (e.g. Andrews et al. 2016), the latter at an unprecedented spatial resolution of 1 AU. Scattered light observations reveal three prominent depressions in polarised intensity signal at 1.57" (94 AU), 0.39" (23 AU) and <0.11" (6.6 AU), when scaled with r^2 to correct for the decrease in stellar illumination as a function of radius. These variations have been successfully modelled using self-consistent radiative transfer (RT) code in vB17 as variations in surface density of micron dust grains in the upper layers of the disk. Andrews et al. (2016) reports the detection of depletion features in mm-sized dust grains at similar angular separations, in addition to an inner depletion feature on scales of 1 AU (20 mas), which could be an indication of planet formation: this last feature lies significantly below the 92.5 mas (5.5 AU) mask edge of the SPHERE ALC coronagraph.

In 2015 vB17 observed TW Hydrae in non-coronagraphic DPI mode in the R' and I bands with ZIMPOL, in addition to non-coronagraphic J-band and coronagraphic H-band observations with IRDIS. Here we use the coronagraphic H-band dataset for our analysis, which makes use of the N_ALC_YJH_S coronagraph system addressed in this work. The H-band Q_ϕ image of TW Hydrae is shown in the upper panel of Fig. 5.9, which displays two broad ring features within the central arcsecond and strong coronagraphic extinction in the centre-most region. This image has been corrected for instrumental polarisation following the approach of van Holstein (2016). Of particular interest to this work is the feature termed gap #3, which is located below approximately 110 mas. This is closely coincident with the shoulder of the ALC extinction profile and so provides a useful validation for our calibration approach. vB17 concluded that gap #3 is a real surface density variation: this is supported both by SED studies (Calvet et al. 2002) and their own RT modelling results, which were forward-modelled via $Q_\pm$ and $U_\pm$ to account for the DPI-mode extinction effect of Sec. 5.4.2.

The lower panel of Fig. 5.9 shows the radial intensity curve of this disk (black curve), azimuthally averaged about the disk centre, which is inferred based on the centroid of the ring #2 and ring #3 features outside the ALC region of influence. Since the disk of TW Hydrae is almost directly face-on and displays no significant azimuthal variations, we may readily perform the ALC calibration using a 1-D radial model, without the need for full 3-D modelling of the target.

The following subsections present the results of the two independent calibration approaches outlined in Sec. 5.4.3: both use the final model of the N_ALC_YJH_S coronagraphic imaging system developed in Sec. 5.3, which includes a 4 mm-defocussed focal-plane mask but is otherwise diffraction-limited.

5.5.2 A simple forward model of the TW Hydrae disk

In order to construct an appropriate object model for the disk, it is first important to estimate the likely components of the signal at each angular separation. This is shown by the coloured lines in the lower panel of Fig. 5.9, which simulate the coronagraphic DPI-mode image of a perfectly uniform r^0 intensity profile (similar to that used in Fig. 5.8), in addition to those of $1/r$ and $1/r^2$ intensity profiles. In the idealised case this last profile corresponds to a disk of constant surface density, since stellar illumination of the disk also falls off with $1/r^2$. The solid lines in this plot show the simulated images resulting from coronagraphic DPI-mode observation of each of the three disk profiles, which are scaled to match the most appropriate region of the observed TW Hydrae radial curve. From this it can immediately be concluded that a disk with constant surface density ($1/r^2$ intensity profile, red curve) cannot by itself explain the observed signal in the inner regions, since it does not match well with any part of the bright ring #3 feature interior to 250 mas. A better model for the outer shoulder of ring #3 from 175 to 250 mas is a 1/r radial intensity profile (green curve), however this also over-estimates the signal interior to the 175 mas peak of this feature, and so cannot continue any further in towards the star. A uniform intensity profile (blue curve) can however provide a near-perfect match to the innermost angular separations of the coronagraph, below this 175 mas peak. By contrast, the outer fall-off of ring #3 beyond 250 mas is significantly faster than $1/r^2$, before flattening out at approximately 400 mas.

This first-order analysis of the likely disk structure motivates a multi-component

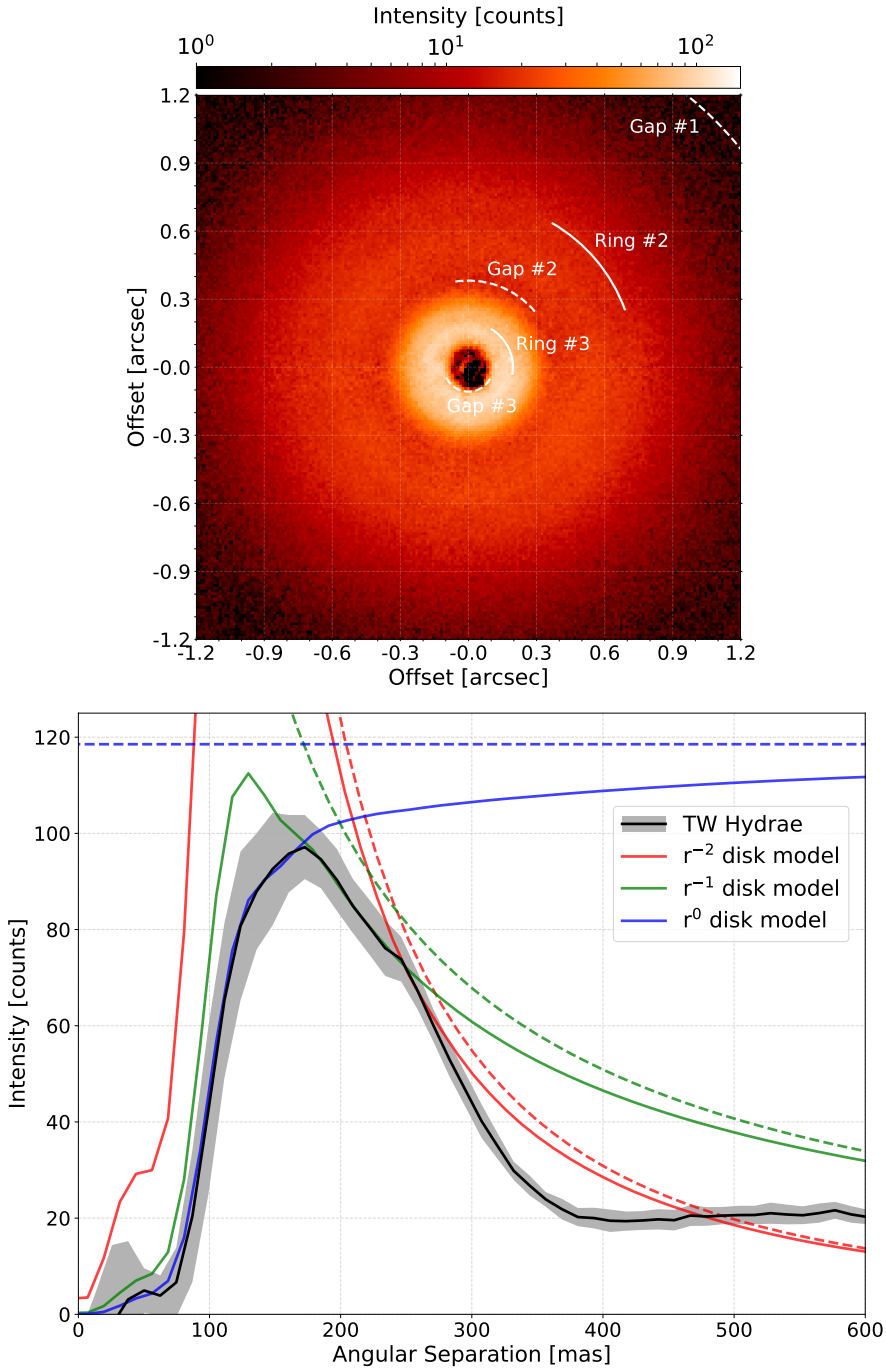


Figure 5.9: H-band coronagraphic Q_ϕ profile of the TW Hydrae disk, after correction for instrumental polarisation via the approach of van Holstein (2016). Intensity is not in this case normalised by r^2, so as to highlight the impact of coronagraphic extinction. *Top:* Logarithmically scaled Q_ϕ intensity image, with features labelled according to the convention of van Boekel et al. (2017). *Bottom:* Radial intensity plot of the Q_ϕ image (black curve), with the grey shaded region denoting the 1σ bound on azimuthal variability. Three simple disk models are over-plotted, with dashed lines corresponding to the true disk signal and solid lines to the simulated coronagraphic Q_ϕ image for each model, each of which are scaled to match appropriate regions of the TW Hydrae disk.

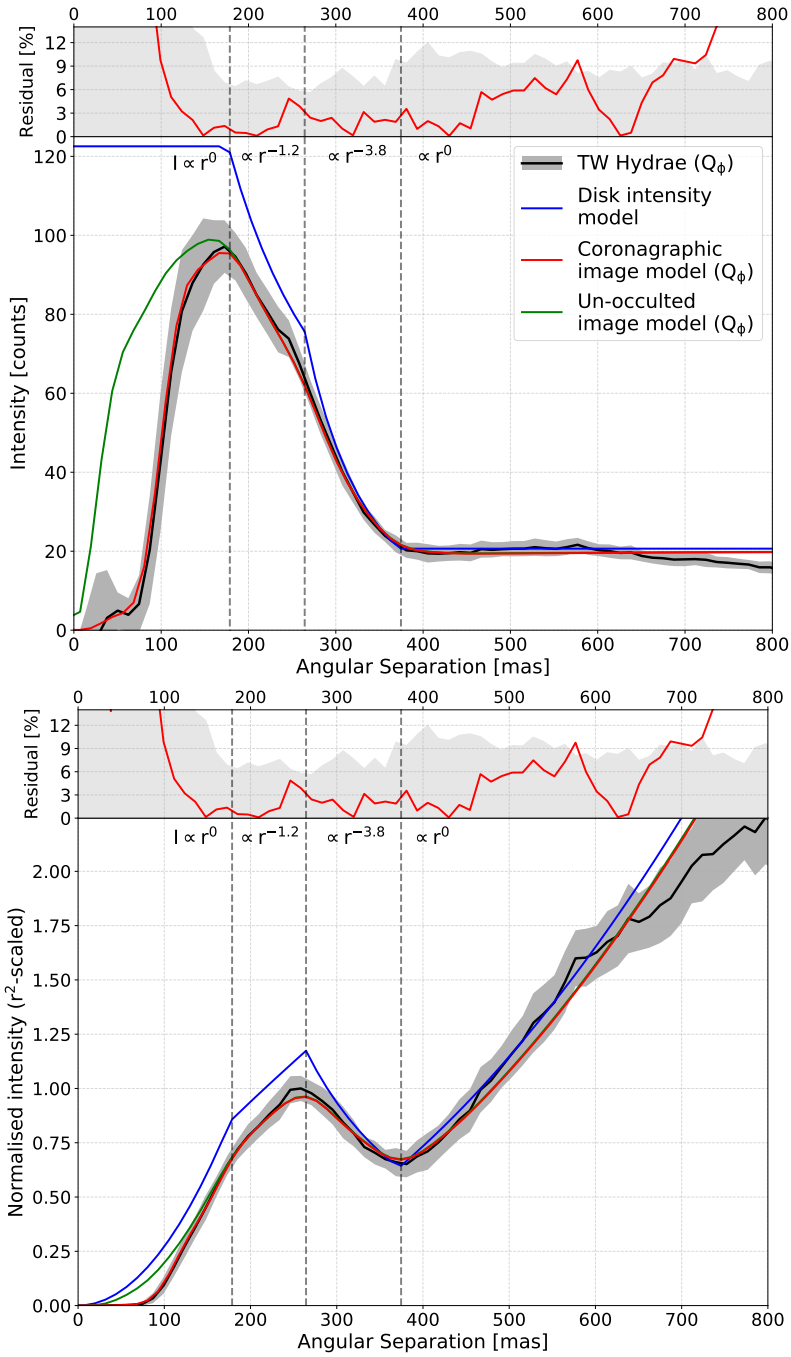


Figure 5.10: Comparison of a simplistic four-power law model of the inner 600 mas of the TW Hydrae disk (blue curve) with vB17 coronagraphic H-band observations (black curve). The upper figure shows this in an un-normalised intensity scaling as in Fig. 5.9, while the lower figure shows the same data scaled with r^2 to correct for the radially decreasing stellar illumination, normalised to unity by the peak of the ring #3 feature. The red and green curves in the main panel show the simulated coronagraphic and un-occulted Q_ϕ images, respectively, which result from the chosen four-power-law model. Percentage residuals between the coronagraphic image model and the observed radial curve are also shown in red in the upper inset panels, where the the grey shaded region denotes the 1σ bound on azimuthal variability of the TW Hydrae disk.

power-law model, with differing spectral index as a function of angular separation. One such model is shown in Fig. 5.10, which is constructed to reproduce all signal interior to 600 mas by using four separate power law components. The intensity profile $I(r)$ of each regime as a function or radius r is described in general by $I(r) = a_1 r^{a_2} + a_3$, where a_1, a_1 and a_3 are free parameters, the final coefficients of which are listed in Table 5.1. This model in fact possesses only three truly free parameters; the two non-zero spectral indices (a_2) and the constant offset (a_3) beyond 375 mas. All other values are constrained either to be zero, or by the requirement of continuity across the three radial boundaries, the locations of which are chosen to correspond with those identified in the previous paragraph based on Fig. 5.9. The resulting model consists of two r^0 regimes for the innermost and outermost regions of the disk, with an approximately $1/r$ profile describing the 175 to 250 mas region and a faster $r^{-3.8}$ describing the subsequent fall-off of the outer edge of the ring feature.

The upper panel of Fig. 5.10 shows that the Q_ϕ coronagraphic simulated image of this simple model (red curve) is capable of providing an excellent fit to the observed disk profile, remaining within the 1σ azimuthal variability limit (grey shaded region) for all separations below 700 mas. Most importantly, this shows that there is no evidence for a decrease in detected disk intensity below angular separations of 180 mas in the Q_ϕ image, once instrumental effects have been corrected for. The equivalent un-occulted Q_ϕ image model (green curve) is identical to the coronagraphic red curve except for the removal of the ALC focal-plane mask, and illustrates how significant the convolutional polarisation effect is for this disk. While the ALC coronagraph begins to cause instrumental signal suppression below 180 mas, noticeable systematic losses due to polarisation erasure already begin to occur at angular separations as large as 300 mas. This outer cutoff is smaller than that quoted in Sec. 5.4.2 since there is relatively little disk signal beyond 300 mas, but it remains arguably a more important effect than coronagraphic extinction for the correct inference of the underlying disk morphology at small angular separations.

The lower panel of Fig. 5.10 plots the same data, but using an r^2 intensity scaling to correct for stellar illumination and hence more fairly present the relative scattering signal from different regions of the disk. In this illumination-corrected scaling, all regions of the model with power-law indices shallower than r^{-2} are seen as radially-increasing, which in turn indicates that for this model there is a relative dearth of scattering signal from the innermost regions of the disk. If this feature may be assumed to originate from variations in the density of micron-sized dust grains in the upper layers of the disk (as opposed to shadowing effects due to surface scale height variations at smaller radii), then this model indicates that there must be a dearth of scattering medium for angular separations smaller than approximately 250 mas: as already shown by Fig. 5.9 it is impossible to construct an accurate model of the observed profile which does not contain such a depression feature from 250 mas down to at least 100 mas. This in turn confirms that the <6.6 AU gap #3 identified by vB17 cannot be purely instrumental in nature.

It should be noted that this forward-modelling process is inherently degenerate at small angular separations due to the high levels of extinction at these locations, and that this proposed model is only one of many possibilities. As may be expected, it was seen that almost any flux distribution underneath the 92.5 mas radius ALC focal-plane mask provides a good match to the observations. For detailed characterisation studies it is therefore recommended that a more thorough forward-modelling approach such

r_0 [mas]	r_1 [mas]	a_1 [counts/mas]	a_2	a_3 [counts]
0	179	0	0	123
179	264	6.10×10^4	-1.2	0
264	375	1.21×10^{11}	-3.8	0
375	∞	0	0	20.6

Table 5.1: Coefficients for the four-power-law model ($I(r) = a_1 r^{a_2} + a_3$) of the inner 600 mas of the TW Hydrae disk. r_0 and r_1 denote the inner and outer boundaries of each regime, respectively.

as a Monte Carlo Markov chain should be used, in order to properly constrain the distribution of possible object models. Such a detailed and computationally expensive analysis is however beyond the scope of this paper.

5.5.3 Assessing the accuracy of ALC extinction calibration by normalisation

The preceding analysis has shown how, by using full model of the spatially variable PSF of the imaging system, it is possible to account for both coronagraphic extinction and convolutional polarisation erasure and determine a simple but accurate model of the underlying disk signal. The final goal of this investigation is to evaluate the relative effectiveness of the simpler, more readily-applicable calibration technique; normalising for the direct throughput of the coronagraphic system. This is significantly faster to implement since it does not require forward modelling as in the previous results, where each step in the computation consists of a computationally expensive element-wise convolution with the spatially-variable instrument PSF. The downsides of this approach are that, as discussed in Sec. 5.4.3, the normalisation calibration does not take polarisation erasure into account and will also over-estimate the innermost regions underneath the ALC mask, where the signal is dominated by non-useful diffracted signal. It is therefore important to determine how far this approach can be trusted for the example of a real protoplanetary disk.

Calibration was achieved by co-aligning a 2-D interpolation of the direct throughput curve model with the ALC mask centre for the TW Hydrae observations, which was itself estimated by minimizing the azimuthal variation of the retrieved radial profile between 80 and 125 mas. The TW Hydrae image was then normalised using this 2-D profile and the radial profile and azimuthal standard deviation computed around the measured disk centre. The normalised image and resulting radial profile are shown in Fig. 5.11: these can be seen to display significant residuals in the inner regions of the disk, which are most likely due to the amplification of low-level instrumental polarisation residuals. Nonetheless, when the median profile of this normalised dataset (solid brown line) is compared with the non-coronagraphic simulated Q_ϕ image from Fig. 5.10 (shown in green in both figures), an extremely close match can be seen down to an angular separation of 75 mas.

Such a close correspondence is in fact expected, since both calibration approaches use the same underlying model of the coronagraphic system and the forward model is tailored to reproduce the coronagraphic TW Hydrae image. It does however clearly illustrate that the simple normalisation approach can be statistically trusted in this case for all angular separations above 75 mas, which in turn implies that there is no

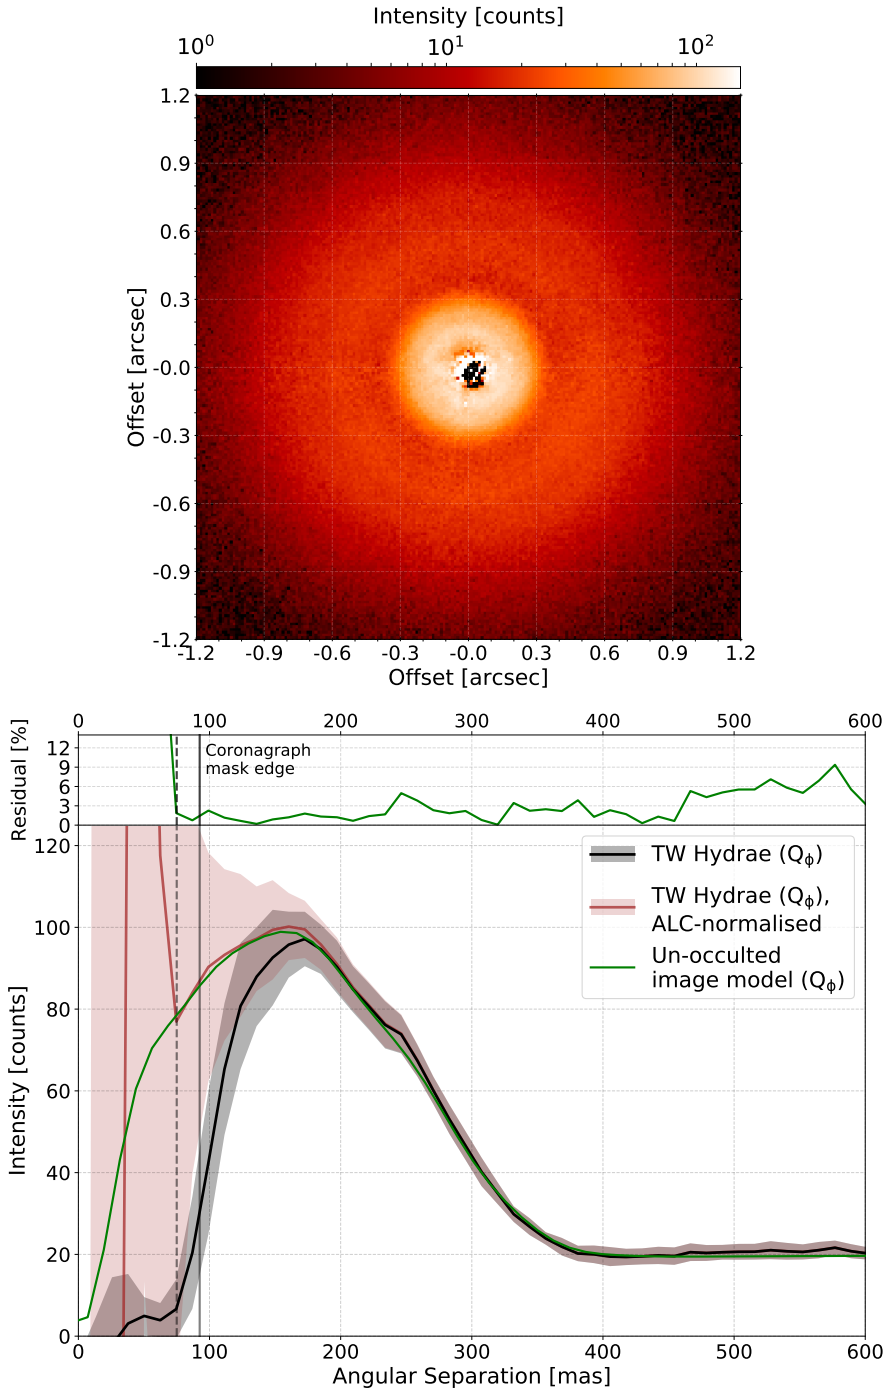


Figure 5.11: Correction of the TW Hydrae Q_ϕ image for coronagraphic extinction, by normalising with the simulated direct-throughput curve of the best-fit ALC system model from Sec. 5.3.4.4. ***Top:*** Q_ϕ image on the same logarithmic scale as Fig. 5.9, normalised with a 2-D interpolation of the ALC extinction curve. ***Bottom:*** Comparison of the radial curve of the normalised TW Hydrae data (brown) with the un-occulted image model simulation of Fig. 5.10 (green). The upper panel shows the percentage residual between these two curves (green), with the left-most dashed line denoting the point of divergence at 75 mas. Brown shading denotes the 1σ bound on azimuthal variability of the ALC-corrected TW Hydrae disk.

significant contribution of non-useful diffracted signal at these separations. Such a small inner limit on the calibration accuracy, which lies below the ALC mask edge, is notably lower than the general 100 mas limit derived in Sec. 5.3 based on Fig. 5.4. This most likely stems from the fact that the TW Hydrae disk displays significantly less extended signal at wider separations than the 300 mas-radius, uniform-intensity Ceres disk, resulting in a correspondingly lower diffracted signal contribution from these resolution elements. This is true of most protoplanetary disk observations, and so one may reasonably expect the normalisation approach to be capable of accurately calibrating for coronagraphic extinction down to the 92.5 mas edge of the ALC mask, unless large regions of high surface brightness are observed within the central few hundred milli-arcseconds of the disk.

Despite the fact that the median profile is seen to be accurately corrected down to this 75 mas limit, such small angular separations are however of limited value in the case of TW Hydrae: it is difficult to draw strong conclusions about the radial disk structure below approximately 175 mas based on the brown curve alone, due to the explosion in relative uncertainty that comes with rapidly decreasing ALC throughput. This is of course notwithstanding the systematic modification made to the apparent disk signal by polarisation erasure: in this case one could erroneously claim that, rather than a uniform-intensity (r^0) signal interior to 180 mas, there is in fact evidence for a further central depression in the un-normalised Q_ϕ image.

It is also important to note that these two ALC-corrected model profiles are still not directly comparable to the flux distribution which would be observed via equivalent non-coronagraphic DPI-mode observations. The subtle difference between the two is that non-coronagraphic observation modes do not include the ALC apodiser and Lyot stop, whereas un-occulted (or ALC extinction-corrected) simulations still include the impact of these pupil-plane optics on the instrument PSF. This is crucial when considering the impact of convolutional polarisation erasure, since the effect is significantly worse in the latter case due to the broader instrument PSF. The simple solution for this when performing forward modelling of the target is to use an equivalent optical model in which these two pupil-plane optics have been removed; in this case leaving only the VLT aperture mask. No such solution is however available for the normalisation calibration, which makes it challenging to directly compare datasets taken in different observing modes even after calibrating for coronagraphic extinction.

The authors therefore strongly recommend that all studies using DPI-mode observations should perform a full analysis of the observed disk signal by forward-modelling through $Q_\pm$ and $U_\pm$ frames, in order to take into account the highly non-linear effects of convolutional polarisation erasure. For intensity-mode observations it should in most cases be sufficient to simply normalise for coronagraphic extinction, provided that the contribution of non-useful diffracted signal is shown to be negligible compared to the direct throughput of the coronagraphic system at the angular separations of interest.

5.6 Conclusions

In this work we have demonstrated that observing solar system targets such as Ceres, which are small enough on which to lock the adaptive optics system of SPHERE but extended enough to cover the full radial extent of the coronagraphic extinction profile,

provide a useful method for calibrating the coronagraphic imaging system. This approach provides the full two-dimensional extinction profile of the coronagraph and is shown to be insensitive to changes in seeing conditions and upstream wavefront aberrations, but with the caveat that it inherently combines the direct (useful) throughput of the coronagraph with the (non-useful) diffracted signal contribution of the extended target. This makes it highly complementary to the more common point-source calibration approach, which measures the one-dimensional direct throughput profile of the system and is more sensitive to temporal instrumental or atmospheric variations.

When such calibration measurements are used to inform instrument simulation efforts, the resulting model of the spatially-variable point-spread function (PSF) of the instrument provides a powerful tool for properly calibrating observations of circumstellar material. We conclude that in the case of the SPHERE-IRDIS apodised Lyot coronagraph (ALC), the observed extinction profile is best reproduced by invoking a 4 mm longitudinal shift of the ALC mask from the coronagraphic focal plane, in an otherwise diffraction-limited system. It is important to determine whether this is a feasible scenario, or whether other factors such as non-common path error (NCPE) are in fact responsible for the minor observed degradation in coronagraphic efficiency. Most importantly, if a time-variable effect such as NCPE is found to be responsible, a more detailed instrument model would need to be developed to accurately calibrate future observations.

For dual-band polarimetric imaging (DPI) observations, the loss of polarised intensity signal due to convolutional polarisation erasure is seen to be of arguably greater importance than coronagraphic extinction, especially at angular separations larger than 100 mas. The effect is however less severe in non-coronagraphic imaging modes, since it is the broadened instrument PSF produced by the use of apodising pupil-plane optics during ALC observations that significantly worsens the convolutional properties of the imaging system. Due to the highly source-dependent and non-linear nature of this polarisation erasure effect, it can only be effectively accounted for via appropriate forward modelling of the surface brightness distribution of the target through the raw $Q_{\pm}$ and $U_{\pm}$ frames in which the effect is produced.

The application of this forward-modelling calibration approach to coronagraphic DPI-mode observations of TW Hydrae has shown that it is possible to adequately explain all observed signal from the inner 600 mas of the disk using a four-component one-dimensional power-law model. The innermost signal depression feature (gap #3) observed by vB17 is in this way confirmed to be a real property of the disk, and not due to artefacts of the coronagraphic or polarimetric imaging process. The simpler calibration approach of direct-throughput normalisation is also capable of self-consistently correcting the median radial profile of the disk for coronagraphic extinction down to an angular separation of 75 mas, significantly below the 92.5 mas edge of the ALC focal-plane mask. While this makes the approach suitable for the calibrating the majority of intensity-mode disk observations, its major limitation for DPI mode is that it cannot be used to correct for polarisation erasure effects, which also prevents a direct comparison with non-coronagraphic observations. Finally, it was seen for both calibration approaches that the choice of disk model remained largely unconstrained below the 100 mas inner-working-angle of the coronagraph, due to the rapid increase in fractional uncertainty with decreasing system throughput at smaller angular separations.

Such a calibration approach is nonetheless indispensable when attempting to correctly interpret the relative surface brightness profiles of the innermost regions of protoplanetary disks, especially between angular separations of approximately 100 to 300 mas. The approach presented in this paper is applicable not only to the SPHERE-IRDIS ALC system, but equally to any other current or future high-contrast imaging instruments which make use of attenuating coronagraphic optics. This is particularly true for JWST-NIRCam, which will incorporate asymmetric wedge coronagraphs (Krist et al. 2009) which cannot be calibrated using conventional 1-D radial extinction profile approaches. Such a detailed calibration effort will ultimately allow these instruments to maximise their science yield at the smallest possible angular separations, where the discovery space is as-yet largely unexplored.

Bibliography

Adams, D. 1986, The Hitch Hiker's Guide to the Galaxy: A Trilogy in Four Parts, Hitchhiker's Guide to the Galaxy Series (Heinemann)

Aime, C. & Soummer, R. 2004, ApJ, 612, L85

Alexander, R. D. & Armitage, P. J. 2007, Monthly Notices of the Royal Astronomical Society, 375, 500

Amara, A. & Quanz, S. P. 2012, MNRAS, 427, 948

Andrews, S. M., Wilner, D. J., Zhu, Z., et al. 2016, ApJ, 820, L40

Anglada-Escudé, G., Amado, P. J., Barnes, J., et al. 2016, Nature, 536, 437

Avenhaus, H., Quanz, S. P., Meyer, M. R., et al. 2014a, ApJ, 790, 56

Avenhaus, H., Quanz, S. P., Schmid, H. M., et al. 2014b, ApJ, 781, 87

Baraffe, I., Chabrier, G., Barman, T. S., Allard, F., & Hauschildt, P. H. 2003, A&A, 402, 701

Barata, J. C. A. & Hussein, M. S. 2012, Brazilian Journal of Physics, 42, 146

Baudoz, P., Dorn, R. J., Lizon, J.-L., et al. 2010, in Proc. SPIE, Vol. 7735, Ground-based and Airborne Instrumentation for Astronomy III, 77355B

Baudoz, P., Rabbia, Y., & Gay, J. 2000, A&AS, 141, 319

Benisty, M., Juhasz, A., Boccaletti, A., et al. 2015, A&A, 578, L6

Benisty, M., Stolker, T., Pohl, A., et al. 2017, A&A, 597, A42

Beuzit, J.-L., Feldt, M., Dohlen, K., et al. 2008, in Proc. SPIE, Vol. 7014, Ground-based and Airborne Instrumentation for Astronomy II, 701418

Biller, B. A., Close, L. M., Lenzen, R., et al. 2006, in IAU Colloq. 200: Direct Imaging of Exoplanets: Science and Techniques, ed. C. Aime & F. Vakili, 571–576

Biller, B. A., Close, L. M., Masciadri, E., et al. 2007, ApJS, 173, 143

Biller, B. A., Liu, M. C., Rice, K., et al. 2015, MNRAS, 450, 4446

Boehle, A., Glauser, A. M., Kenworthy, M. A., et al. 2018, in Proc. SPIE, Vol. 10702, Ground-based and Airborne Instrumentation for Astronomy VII, 10702 – 10702 – 8

Bonnefoy, M., Zurlo, A., Baudino, J. L., et al. 2016, A&A, 587, A58

Booth, M. J. 2003, Proc. SPIE, 5162, 79

Borucki, W. J., Koch, D., Basri, G., et al. 2010, Science, 327, 977

Boss, A. P. 2001, ApJ, 563, 367

Brandl, B. R., Feldt, M., Glasse, A., et al. 2014, in Proc. SPIE, Vol. 9147, Ground-based and Airborne Instrumentation for Astronomy V, 914721

Brauer, F., Dullemond, C. P., Johansen, A., et al. 2007, A&A, 469, 1169

Brogi, M., Keller, C. U., de Juan Ovelar, M., et al. 2012a, A&A, 545, L5

Brogi, M., Snellen, I. A. G., de Kok, R. J., et al. 2012b, Nature, 486, 502

Brown, B. R. & Lohmann, A. W. 1969, IBM Journal of Research and Development, 13, 160

Calvet, N., D'Alessio, P., Hartmann, L., et al. 2002, ApJ, 568, 1008

Canovas, H., Ménard, F., Hales, A., et al. 2013, A&A, 556, A123

Carrano, C. J., Olivier, S. S., Brase, J. M., Macintosh, B. A., & An, J. R. 1998, in Proc. SPIE, Vol. 3353, Adaptive Optical System Technologies, ed. D. Bonaccini & R. K. Tyson, 658–667

Chabrier, G., Johansen, A., Janson, M., & Rafikov, R. 2014, Protostars and Planets VI, 619

Chambers, J. E. 2006, ApJ, 652, L133

Changhai, L., Fengjie, X., Shengyang, H., & Zongfu, J. 2011, Appl. Opt., 50, 1631

Chauvin, G., Lagrange, A.-M., Dumas, C., et al. 2004, A&A, 425, L29

Cheetham, A. C., Girard, J., Lacour, S., et al. 2016, in Proc. SPIE, Vol. 9907, Optical and Infrared Interferometry and Imaging V, 99072T

Cheney, E. & Kincaid, D. 2009, Linear Algebra: Theory and Applications (Jones and Bartlett Publishers)

Clarke, C. J., Gendrin, A., & Sotomayor, M. 2001, MNRAS, 328, 485

Codona, J. L. 2013, Optical Engineering, 52, 097105

Codona, J. L. & Kenworthy, M. 2013, The Astrophysical Journal, 767, 100

Codona, J. L., Kenworthy, M. A., Hinz, P. M., Angel, J. R. P., & Woolf, N. J. 2006, in Proc. SPIE, Vol. 6269, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 62691N

Crepp, J. R., Pueyo, L., Brenner, D., et al. 2011, ApJ, 729, 132

Currie, T., Cloutier, R., Brittain, S., et al. 2015, ApJ, 814, L27

Côté, O., Allain, G., Brousseau, D., et al. 2018, in Proc. SPIE, Vol. 10702, Ground-based and Airborne Instrumentation for Astronomy VII, 10702 – 10702 – 8

de Boer, J., Girard, J. H., Canovas, H., et al. 2017, MNRAS, 466, L7

de Boer, J., Salter, G., Benisty, M., et al. 2016, A&A, 595, A114

de Juan Ovelar, M., Pinilla, P., Min, M., Dominik, C., & Birnstiel, T. 2016, MNRAS, 459, L85

de Kok, R. J., Brogi, M., Snellen, I. A. G., et al. 2013, A&A, 554, A82

Doelman, D. S., Snik, F., Warriner, N. Z., & Escuti, M. J. 2017, in Proc. SPIE, Vol. 10400, Techniques and Instrumentation for Detection of Exoplanets VIII, 10400 – 10400 – 12

Doelman, D. S., Tuthill, P., Norris, B., et al. 2018, in Proc. SPIE, Vol. 10701, Optical and Infrared Interferometry and Imaging VI, 10701 – 10701 – 11

Dohlen, K., Langlois, M., Saisse, M., et al. 2008, in Proc. SPIE, Vol. 7014, Ground-based and Airborne Instrumentation for Astronomy II, 70143L

Dohlen, K., Vigan, A., Mouillet, D., et al. 2016, in Proc. SPIE, Vol. 9908, Ground-based and Airborne Instrumentation for Astronomy VI, 99083D

Dong, S., Haist, T., Osten, W., Ruppel, T., & Sawodny, O. 2012, Appl. Opt., 51, 1318

Doucet, C., Habart, E., Pantin, E., et al. 2007, A&A, 470, 625

Draine, B. T. 2006, ApJ, 636, 1114

ESO. 2015, SPHERE user manual, Period 96, Phase 2 (4th release), https://www.eso.org/sci/facilities/paranal/instruments/sphere/doc.html, accessed: 2016-06-21

ESO. 2016, SPHERE user manual, Period 99, Phase 1 (6th release), https://www.eso.org/sci/facilities/paranal/instruments/sphere/doc.html, accessed: 2017-02-06

ESO. 2018, SPHERE user manual, Period 101, Phase 1 (9th release), https://www.eso.org/sci/facilities/paranal/instruments/sphere/doc.html, accessed: 2018-04-13

Fitzgerald, M. P. & Graham, J. R. 2006, ApJ, 637, 541

Fortney, J. J., Marley, M. S., Saumon, D., & Lodders, K. 2008, ApJ, 683, 1104
Fressin, F., Torres, G., Charbonneau, D., et al. 2013, ApJ, 766, 81
Fusco, T., Petit, C., Rousset, G., et al. 2006, in Proc. SPIE, Vol. 6272, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 62720K
Fusco, T., Sauvage, J.-F., Mouillet, D., et al. 2016, in Proc. SPIE, Vol. 9909, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 99090U
Fusco, T., Sauvage, J.-F., Petit, C., et al. 2014, in Proc. SPIE, Vol. 9148, Adaptive Optics Systems IV, 91481U
Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2016a, A&A, 595, A2
Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016b, A&A, 595, A1
Garufi, A., Quanz, S. P., Avenhaus, H., et al. 2013, A&A, 560, A105
Ginski, C., Benisty, M., van Holstein, R. G., et al. 2018, A&A, 616, A79
Ginski, C., Stolker, T., Pinilla, P., et al. 2016, A&A, 595, A112
Giorgini, J. et al. 2015, NASA JPL Horizons On-Line Ephemeris System, https://ssd.jpl.nasa.gov/horizons.cgi, last accessed: 2018-04-11
Give'on, A., Belikov, R., Shaklan, S., & Kasdin, J. 2007, Opt. Express, 15, 12338
Give'on, A., Kasdin, N. J., Vanderbei, R. J., & Avitzour, Y. 2006, J. Opt. Soc. Am. A, 23, 1063
Gonsalves, R. A. 1982, Optical Engineering, 21, 215829
Gonsalves, R. A. 2001, Opt. Lett., 26, 684
Gonsalves, R. A. 2002, in European Southern Observatory Conference and Workshop Proceedings, Vol. 58, European Southern Observatory Conference and Workshop Proceedings, ed. E. Vernet, R. Ragazzoni, S. Esposito, & N. Hubin, 121
Gould, A., Udalski, A., Shin, I.-G., et al. 2014, Science, 345, 46
Grady, C. A., Muto, T., Hashimoto, J., et al. 2013, ApJ, 762, 48
Guerri, G., Daban, J.-B., Robbe-Dubois, S., et al. 2011, Experimental Astronomy, 30, 59
Guyon, O. 2003, A&A, 404, 379
Guyon, O., Martinache, F., Cady, E. J., et al. 2012, in Proc. SPIE, Vol. 8447, Adaptive Optics Systems III, 84471X
Haffert, S. Y., Por, E. H., Keller, C. U., et al. 2018a, ArXiv e-prints [arXiv:1803.10693]
Haffert, S. Y., Wilby, M. J., Keller, C. U., & Snellen, I. A. G. 2016, in Proc. SPIE, Vol. 9908, Ground-based and Airborne Instrumentation for Astronomy VI, 990867
Haffert, S. Y., Wilby, M. J., Keller, C. U., et al. 2018b, in Proc. SPIE, Vol. 10703, Adaptive Optics Systems VI, 10703 – 10703 – 11
Heap, S. R., Lindler, D. J., Lanz, T. M., et al. 2000, The Astrophysical Journal, 539, 435
Herbig, G. H. 1960, ApJS, 4, 337
Hoeijmakers, H. J., Schwarz, H., Snellen, I. A. G., et al. 2018, A&A, 617, A144
Howard, A. W., Marcy, G. W., Johnson, J. A., et al. 2010, Science, 330, 653
Howard, A. W., Sanchis-Ojeda, R., Marcy, G. W., et al. 2013, Nature, 503, 381
Hubickyj, O., Bodenheimer, P., & Lissauer, J. J. 2005, AGU Fall Meeting Abstracts, P42A
Jovanovic, N., Absil, O., Baudoz, P., et al. 2018, 10703, 107031U
Jovanovic, N., Martinache, F., Guyon, O., et al. 2015, PASP, 127, 890

Kasdin, N. J., Vanderbei, R. J., Littman, M. G., Carr, M., & Spergel, D. N. 2004, in Proc. SPIE, Vol. 5487, Optical, Infrared, and Millimeter Space Telescopes, ed. J. C. Mather, 1312–1321

Kasper, M., Beuzit, J.-L., Verinaud, C., et al. 2010, Proc. SPIE, 7735, 77352E

Käufl, H.-U., Ballester, P., Biereichel, P., et al. 2004, in Proc. SPIE, Vol. 5492, Ground-based Instrumentation for Astronomy, ed. A. F. M. Moorwood & M. Iye, 1218–1227

Keller, C. U. 2016, in Proc. SPIE, Vol. 9908, Ground-based and Airborne Instrumentation for Astronomy VI, 99089V

Keller, C. U., Korkiakoski, V., Doelman, N., et al. 2012, in Proc. SPIE, Vol. 8447, Adaptive Optics Systems III, 844721

Keller, C. U., Schmid, H. M., Venema, L. B., et al. 2010, in Ground-based and Airborne Instrumentation for Astronomy III, Vol. 7735, 77356G–77356G–13

Kenworthy, M. A., Codona, J. L., Hinz, P. M., et al. 2007, ApJ, 660, 762

Kenworthy, M. A., Hinz, P. M., Codona, J. L., et al. 2010a, in Proc. SPIE, Vol. 7734, Optical and Infrared Interferometry II, 77342P

Kenworthy, M. A., Quanz, S., Meyer, M., et al. 2010b, The Messenger, 141, 2

Kenworthy, M. A., Quanz, S. P., Meyer, M. R., et al. 2010c, in Proc. SPIE, Vol. 7735, Ground-based and Airborne Instrumentation for Astronomy III, 773532

Keppler, M., Benisty, M., Müller, A., et al. 2018, A&A, 617, A44

Kipping, D. M. 2009a, MNRAS, 392, 181

Kipping, D. M. 2009b, MNRAS, 396, 1797

Knutson, H. A., Charbonneau, D., Allen, L. E., et al. 2007, Nature, 447, 183

Komanduri, R. K., Lawler, K. F., & Escuti, M. J. 2013, Opt. Express, 21, 404

Korkiakoski, V., Doelman, N., Codona, J., et al. 2013, Appl. Opt., 52, 7554

Korkiakoski, V., Keller, C. U., Doelman, N., et al. 2012, in Proc. SPIE, Vol. 8447, Adaptive Optics Systems III, 84475Z

Korkiakoski, V., Keller, C. U., Doelman, N., et al. 2014, Appl. Opt., 53, 4565

Kraus, A. L. & Ireland, M. J. 2012, ApJ, 745, 5

Krist, J. E., Balasubramanian, K., Beichman, C. A., et al. 2009, in Proc. SPIE, Vol. 7440, Techniques and Instrumentation for Detection of Exoplanets IV, 74400W

Krist, J. E., Stapelfeldt, K. R., Bryden, G., et al. 2010, AJ, 140, 1051

Kuhn, J. R., Potter, D., & Parise, B. 2001, ApJ, 553, L189

Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, É. 2007, ApJ, 660, 770

Lagrange, A.-M., Bonnefoy, M., Chauvin, G., et al. 2010, Science, 329, 57

Lamb, M., Correia, C., Sauvage, J.-F., et al. 2016, in Proc. SPIE, Vol. 9909, Adaptive Optics Systems V, 99096D

Lawson, P. R., Poyneer, L., Barrett, H., et al. 2012, in Proc. SPIE, Vol. 8447, Adaptive Optics Systems III, 844722

Lenzen, R., Hartung, M., Brandner, W., et al. 2003, in Proc. SPIE, Vol. 4841, Instrument Design and Performance for Optical/Infrared Ground-based Telescopes, ed. M. Iye & A. F. M. Moorwood, 944–952

Lovis, C., Snellen, I., Mouillet, D., et al. 2017, A&A, 599, A16

Lozi, J., Guyon, O., Jovanovic, N., et al. 2018, in Proc. SPIE, Vol. 10703, Adaptive Optics Systems VI, 10703 – 10703 – 12

Lucy, L. B. 1974, AJ, 79, 745

Lyot, B. 1939, MNRAS, 99, 580
Macintosh, B. A., Graham, J. R., Barman, T., et al. 2015, Science, 350, 64
Macintosh, B. A., Graham, J. R., Ingraham, P., et al. 2014, Proceedings of the National Academy of Science, 111, 12661
Macintosh, B. A., Graham, J. R., Palmer, D. W., et al. 2008, in Proc. SPIE, Vol. 7015, Adaptive Optics Systems, 701518
Maire, A.-L., Bonnefoy, M., Ginski, C., et al. 2016, A&A, 587, A56
Marino, S., Casassus, S., Perez, S., et al. 2015, ApJ, 813, 76
Marley, M. S., Saumon, D., Cushing, M., et al. 2012, ApJ, 754, 135
Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, The Astrophysical Journal, 641, 556
Marois, C., Macintosh, B., Barman, T., et al. 2008, Science, 322, 1348
Marois, C., Macintosh, B., & Véran, J.-P. 2010a, in Adaptive Optics Systems II, Vol. 7736, 77361J–77361J–12
Marois, C., Zuckerman, B., Konopacky, Q. M., Macintosh, B., & Barman, T. 2010b, Nature, 468, 1080
Martín, E. L. & Zapatero Osorio, M. R. 2003, ApJ, 593, L113
Martinache, F., Guyon, O., Jovanovic, N., et al. 2014, PASP, 126, 565
Martinez, P., Boccaletti, A., Kasper, M., Baudoz, P., & Cavarroc, C. 2007, A&A, 474, 671
Martinez, P., Kasper, M., Costille, A., et al. 2013, A&A, 554, A41
Martinez, P., Loose, C., Aller Carpentier, E., & Kasper, M. 2012, A&A, 541, A136
Masuda, K. 2014, ApJ, 783, 53
Mawet, D., Pueyo, L., Lawson, P., et al. 2012, in Proc. SPIE, Vol. 8442, Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave, 844204
Mawet, D., Serabyn, E., Liewer, K., et al. 2010, ApJ, 709, 53
Mayer, L., Quinn, T., Wadsley, J., & Stadel, J. 2002, Science, 298, 1756
Mayer, L., Wadsley, J., Quinn, T., & Stadel, J. 2005, MNRAS, 363, 641
Mayor, M. & Queloz, D. 1995, Nature, 378, 355
Meheut, H., Meliani, Z., Varniere, P., & Benz, W. 2012, A&A, 545, A134
Merín, B., Brown, J. M., Oliveira, I., et al. 2010, ApJ, 718, 1200
Miller, K., Males, J. R., Guyon, O., et al. 2018, in Proc. SPIE, Vol. 10703, Adaptive Optics Systems VI, 10703 – 10703 – 17
Milli, J., Mawet, D., Mouillet, D., Kasper, M., & Girard, J. H. 2016, in Astrophysics and Space Science Library, Vol. 439, Astronomy at High Angular Resolution, ed. H. M. J. Boffin, G. Hussain, J.-P. Berger, & L. Schmidtobreick, 17
Milli, J., Mouillet, D., Lagrange, A.-M., et al. 2012, A&A, 545, A111
Miskiewicz, M. N. & Escuti, M. J. 2014, Opt. Express, 22, 12691
Molnár, L., Plachy, E., Juhász, Á. L., & Rimoldini, L. 2018, ArXiv e-prints [arXiv:1805.11395]
Monnier, J. D., Kraus, S., Buscher, D., et al. 2014, in Proc. SPIE, Vol. 9146, Optical and Infrared Interferometry IV, 914610
Mordasini, C., Alibert, Y., Benz, W., & Naef, D. 2009, A&A, 501, 1161
Morzinski, K. M., Close, L. M., Males, J. R., et al. 2014, in Proc. SPIE, Vol. 9148, Adaptive Optics Systems IV, 914804

Mosleh, A., Langlois, J. M. P., & Green, P. 2014, in Computer Vision – ECCV 2014, ed. D. Fleet, T. Pajdla, B. Schiele, & T. Tuytelaars (Cham: Springer International Publishing), 247–262

Mulders, G. D., Min, M., Dominik, C., Debes, J. H., & Schneider, G. 2013, A&A, 549, A112

Muro-Arena, G. A., Dominik, C., Waters, L. B. F. M., et al. 2018, A&A, 614, A24

Muto, T., Grady, C. A., Hashimoto, J., et al. 2012, ApJ, 748, L22

N'Diaye, M., Dohlen, K., Caillat, A., et al. 2014, Proc. SPIE, 9148, 91485H

N'Diaye, M., Dohlen, K., Tisserand, S., et al. 2011, in Proc. SPIE, Vol. 8169, Optical Fabrication, Testing, and Metrology IV, 81690G

N'Diaye, M., Martinache, F., Jovanovic, N., et al. 2018, A&A, 610, A18

N'Diaye, M., Vigan, A., Dohlen, K., et al. 2016, A&A, 592, A79

Neil, M. A. A., Booth, M. J., & Wilson, T. 2000, J. Opt. Soc. Am. A, 17, 1098

Nielsen, E. L., Close, L. M., Biller, B. A., Masciadri, E., & Lenzen, R. 2008, ApJ, 674, 466

Noll, R. J. 1976, J. Opt. Soc. Am., 66, 207

Otten, G. P. P. L. 2016, Suppressing a Sea of Starlight : enabling technology for the direct imaging of exoplanets (PhD thesis) (Leiden University)

Otten, G. P. P. L., Snik, F., Kenworthy, M. A., et al. 2017, ApJ, 834, 175

Otten, G. P. P. L., Snik, F., Kenworthy, M. A., Miskiewicz, M. N., & Escuti, M. J. 2014a, Opt. Express, 22, 30287

Otten, G. P. P. L., Snik, F., Kenworthy, M. A., et al. 2014b, in Proc. SPIE, Vol. 9151, Advances in Optical and Mechanical Technologies for Telescopes and Instrumentation, 9151 – 9151 – 10

Paul, B., Sauvage, J.-F., Mugnier, L. M., et al. 2014a, in Proc. SPIE, Vol. 9147, Ground-based and Airborne Instrumentation for Astronomy V, 91479O

Paul, B., Sauvage, J.-F., Mugnier, L. M., et al. 2014b, A&A, 572, A32

Perrin, M. D., Duchene, G., Millar-Blanchaer, M., et al. 2015, ApJ, 799, 182

Perrot, C., Boccaletti, A., Pantin, E., et al. 2016, A&A, 590, L7

Petigura, E. A., Howard, A. W., & Marcy, G. W. 2013, Proceedings of the National Academy of Science, 110, 19273

Petit, C., Sauvage, J.-F., Fusco, T., et al. 2016, Journal of Astronomical Telescopes, Instruments, and Systems, 2, 2

Pinilla, P., de Boer, J., Benisty, M., et al. 2015, A&A, 584, L4

Pinte, C., Harries, T. J., Min, M., et al. 2009, A&A, 498, 967

Pollack, J. B., Hubickyj, O., Bodenheimer, P., et al. 1996, Icarus, 124, 62

Polo, A., Haber, A., Pereira, S. F., Verhaegen, M., & Urbach, H. P. 2013, Journal of the European Optical Society, 8

Por, E. H. 2017, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10400, Techniques and Instrumentation for Detection of Exoplanets VIII, 104000V

Por, E. H. & Haffert, S. Y. 2018, ArXiv e-prints [arXiv:1803.10691]

Por, E. H. & Keller, C. U. 2016, in Proc. SPIE, Vol. 9909, Adaptive Optics Systems V, 990959

Qi, C., Ho, P. T. P., Wilner, D. J., et al. 2004, ApJ, 616, L11

Quanz, S. P., Amara, A., Meyer, M. R., et al. 2015, ApJ, 807, 64

Quanz, S. P., Amara, A., Meyer, M. R., et al. 2013a, ApJ, 766, L1
Quanz, S. P., Avenhaus, H., Buenzli, E., et al. 2013b, ApJ, 766, L2
Quanz, S. P., Meyer, M. R., Kenworthy, M. A., et al. 2010, ApJ, 722, L49
Rappaport, S., Levine, A., Chiang, E., et al. 2012, ApJ, 752, 1
Rice, W. K. M. & Armitage, P. J. 2003, ApJ, 598, L55
Richardson, W. H. 1972, Journal of the Optical Society of America (1917-1983), 62, 55
Ridden-Harper, A. R., Snellen, I. A. G., Keller, C. U., et al. 2016, A&A, 593, A129
Roberts, Jr., L. C., Perrin, M. D., Marchis, F., et al. 2004, in Proc. SPIE, Vol. 5490, Advancements in Adaptive Optics, ed. D. Bonaccini Calia, B. L. Ellerbroek, & R. Ragazzoni, 504–515
Roddier, F. & Roddier, C. 1997, PASP, 109, 815
Rodenhuis, M., Canovas, H., Jeffers, S., & Keller, C. 2011, in Astronomical Society of the Pacific Conference Series, Vol. 449, Astronomical Polarimetry 2008: Science from Small to Large Telescopes, ed. P. Bastien, N. Manset, D. P. Clemens, & N. St-Louis, 33
Roelfsema, R., Bazzon, A., Schmid, H. M., et al. 2016, in Proc. SPIE, Vol. 9909, Adaptive Optics Systems V, 990927
Ros, K. & Johansen, A. 2013, A&A, 552, A137
Sallum, S., Follette, K. B., Eisner, J. A., et al. 2015, Nature, 527, 342
Sauvage, J.-F., Fusco, T., Guesalaga, A., et al. 2015, in Adaptive Optics for Extremely Large Telescopes 4 – Conference Proceedings
Sauvage, J.-F., Fusco, T., Lamb, M., et al. 2016, in Proc. SPIE, Vol. 9909, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 990916
Sauvage, J.-F., Fusco, T., LeMignant, D., et al. 2011, in Second International Conference on Adaptive Optics for Extremely Large Telescopes, 48
Sauvage, J.-F., Fusco, T., Petit, C., et al. 2010, in Adaptive Optics Systems II, Vol. 7736, 77360F–77360F–10
Sauvage, J.-F., Fusco, T., Petit, C., et al. 2014, in Adaptive Optics Systems IV, Vol. 9148, 9148 – 9148 – 10
Sauvage, J.-F., Fusco, T., Rousset, G., & Petit, C. 2007, Journal of the Optical Society of America A, 24, 2334
Schwarz, H. 2017, Spinning worlds (PhD thesis) (Leiden University)
Schwarz, H., Ginski, C., de Kok, R. J., et al. 2016a, A&A, 593, A74
Schwarz, K. R., Bergin, E. A., Cleeves, L. I., et al. 2016b, ApJ, 823, 91
Shvartzvald, Y., Yee, J. C., Calchi Novati, S., et al. 2017, ApJ, 840, L3
Smith, W. H. 1987, PASP, 99, 1344
Snellen, I. A. G., Brandl, B. R., de Kok, R. J., et al. 2014, Nature, 509, 63
Snellen, I. A. G., de Kok, R., Birkby, J. L., et al. 2015, A&A, 576, A59
Snellen, I. A. G., de Kok, R. J., de Mooij, E. J. W., & Albrecht, S. 2010, Nature, 465, 1049
Snik, F., Otten, G., Kenworthy, M., et al. 2012, in Proc. SPIE, Vol. 8450, Modern Technologies in Space- and Ground-based Telescopes and Instrumentation II, 84500M
Sokal, K. R., Deen, C. P., Mace, G. N., et al. 2018, ApJ, 853, 120
Soummer, R., Ferrari, A., Aime, C., & Jolissaint, L. 2007a, ApJ, 669, 642
Soummer, R., Pueyo, L., & Larkin, J. 2012, ApJ, 755, L28

Soummer, R., Pueyo, L., Sivaramakrishnan, A., & Vanderbei, R. J. 2007b, Opt. Express, 15, 15935
Spangenberg, D.-M., Dudley, A., Neethling, P. H., Rohwer, E. G., & Forbes, A. 2014, Opt. Express, 22, 13870
Spiegel, D. S., Burrows, A., & Milsom, J. A. 2011, ApJ, 727, 57
Stam, D. M., Hovenier, J. W., & Waters, L. B. F. M. 2004, A&A, 428, 663
Stevenson, K. B., Désert, J.-M., Line, M. R., et al. 2014, Science, 346, 838
Stolker, T., Min, M., Stam, D. M., et al. 2017, A&A, 607, A42
Strom, K. M., Strom, S. E., Edwards, S., Cabrit, S., & Skrutskie, M. F. 1989, AJ, 97, 1451
Stuik, R., Bailey, J. I., Dorval, P., et al. 2017, A&A, 607, A45
Thalmann, C., Schmid, H. M., Boccaletti, A., et al. 2008, in Proc. SPIE, Vol. 7014, Ground-based and Airborne Instrumentation for Astronomy II, 70143F
Tuthill, P., Lacour, S., Amico, P., et al. 2010, in Proc. SPIE, Vol. 7735, Ground-based and Airborne Instrumentation for Astronomy III, 77351O
van Boekel, R., Henning, T., Menu, J., et al. 2017, ApJ, 837, 132
van Holstein, R. 2016, Master's Thesis, TU Delft
van Holstein, R. G., Snik, F., Girard, J. H., et al. 2017, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10400, Techniques and Instrumentation for Detection of Exoplanets VIII, 1040015
Vigan, A., Bonnefoy, M., Ginski, C., et al. 2016a, A&A, 587, A55
Vigan, A., Postnikova, M., Caillat, A., et al. 2016b, in Proc. SPIE, Vol. 9909, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 99093F
Wagner, K., Apai, D., Kasper, M., & Robberto, M. 2015, ApJ, 813, L2
Wang, J. J., Graham, J. R., Pueyo, L., et al. 2016, AJ, 152, 97
Weidenschilling, S. J. & Marzari, F. 1996, Nature, 384, 619
Wilby, M. J., Keller, C. U., Haffert, S., et al. 2016a, in Proc. SPIE, Vol. 9909, Adaptive Optics Systems V, 990921
Wilby, M. J., Keller, C. U., Sauvage, J.-F., et al. 2016b, in Proc. SPIE, Vol. 9909, Adaptive Optics Systems V, 99096C
Williams, J. P. & Cieza, L. A. 2011, ARA&A, 49, 67
Winn, J. N. & Fabrycky, D. C. 2015, ARA&A, 53, 409
Wolszczan, A. & Frail, D. A. 1992, Nature, 355, 145
Wyatt, M. C. 2008, Annual Review of Astronomy and Astrophysics, 46, 339
Zurlo, A., Vigan, A., Galicher, R., et al. 2016, A&A, 587, A57

List of Abbreviations

ADI	Angular Differential Imaging
ADU	Analog-to-Digital Unit
ALMA	Atacama Large Millimeter Array
A(P)LC	Apodised (Pupil) Lyot Coronagraph
APP	Apodising Phase Plate
AU	Astronomical Unit
CCD	Charge Coupled Device
cMWS	coronagraphic Modal Wavefront Sensor
CO	Carbon Monoxide
COFFEE	COronagraphic Focal-plane wave-Front Estimation for Exoplanet detection
CRIRES	CRyogenic high-resolution InfraRed Echelle Spectrograph
DL	Diffraction-Limited
DM	Deformable Mirror
dOTF	differential Optical Transfer Function
DPI	Dual-band Polarimetric Imaging
DTTS	Differential Tip-Tilt Sensor
EFC	Electric Field Conjugation
ELT	Extremely Large Telescope
EPICS	ExoPlanet Imaging Camera and Spectrograph
ERIS	Enhanced Resolution Imager and Spectrograph
ESO	European Southern Observatory
F&F	Fast & Furious
FF-GS	Fast & Furious Gerchberg-Saxton
FFT	Fast Fourier Transform
FoV	Field of View
FP	Focal Plane
FWHM	Full Width at Half Maximum
GMC	Giant Molecular Cloud
GPI	Gemini Planet Imager
GUI	Graphical User Interface
HAM	Holographic Aperture Masking
HCI	High-Contrast Imaging
HDS	High-Dispersion Spectroscopy
HMWS	Holographic Modal Wavefront Sensor
HRS	High-Resolution Spectroscopy

HST	Hubble Space Telescope
IFS	Integral Field Spectrograph
IFU	Integral Field Unit
IRDIS	Infra-Red Dual-beam Imager and Spectrograph
IWA	Inner-Working Angle
JWST	James Webb Space Telescope
LAM	Laboratoire d'Astrophysique de Marseille
LBT	Large Binocular Telescope
LCOS-SLM	Liquid Crystal on Silicon Spatial Light Modulator
LEXI	Leiden EXoplanet Instrument
LO	Low Order
LOCI	Locally-Optimised Combination of Images
LWE	Low-Wind Effect
M_{Earth}	Earth Masses
M_{jup}	Jupiter Masses
M2	Secondary Mirror
MagAO	Magellan Adaptive Optics
METIS	Mid-infrared E-ELT Imager and Spectrograph
MITHIC	Marseille Imaging Testbed for HIgh Contrast
NACO	NAOS-CONICA: Nasmyth Adaptive Optics System (NAOS) Near-Infrared Imager and Spectrograph (CONICA)
NCP	Non-Common Path
NCPA	Non-Common Path Aberrations
NCPE	Non-Common Path Errors
NIR	Near-InfraRed
NIRC2	Near-InfraRed Camera 2
NOVA	Nederlandse Onderzoekschool Voor Astronomie
PCA	Principal Component Analysis
PFI	Planet Formation Imager
PI	Polarised Intensity
PIAA	Phase-Induced Amplitude Apodisation
PP	Pupil Plane
PPD	ProtoPlanetary Disk
PSF	Point-Spread Function
PTT	Piston-Tip-Tilt
PVE	Peak-to-Valley Error
QSS	Quasi-Static Speckles
RDI	Reference Differential Imaging
RMS	Root-Mean-Square
RT	Radiative Transfer
RV	Radial Velocity
SAM	Sparse Aperture Masking

SAXO	Sphere Adaptive optics for eXoplanet Observation
SCAR	Single-mode Complex Amplitude Retrieval coronagraph
ScExAO	Subaru coronagraphic Extreme Adaptive Optics
SDI	Simultaneous/Spectral Differential Imaging
SED	Spectral Energy Distribution
SH-WFS	Shack-Hartmann WaveFront Sensor
SLM	Spatial Light Modulator
S/N	Signal-to-Noise Ratio
SNR	Signal-to-Noise Ratio
SPHERE	Spectro-Polarimetric High-contrast Exoplanet REsearch instrument
SR	Strehl Ratio
TMT	Thirty-Meter Telescope
TTV	Transit Timing Variation
UV	Ultra-Violet
vAPP	vector Apodising Phase Plate
VLT	Very Large Telescope
WHT	William Herschel Telescope
WFS	WaveFront Sensor
XAO	eXtreme Adaptive Optics
ZELDA	Zernike sensor for Extremely Low-level Differential Aberrations
ZIMPOL	Zurich IMaging POLarimeter

Chapter 6

Nederlandse Samenvatting

Al eeuwen worden filosofen, wetenschappers en science fiction schrijvers gefascineerd door het mogelijk bestaan van planetenstelsels rondom andere sterren dan de zon. En, zoals Giovanni Schiaparelli claimde kunstmatige kanalen op Mars te zien in 1877 of zoals de fantasieën over het bestaan van tropische paradijzen onder het dikke wolkendeken van Venus uit het begin van de 19de eeuw, gaat deze fascinatie bijna altijd gepaard met de begeerte om te weten: zou er leven kunnen bestaan op deze planeten?

Sinds deze oude mijmeringen, is onze wetenschappelijk kennis over planeten buiten ons eigen zonnestelsel (exoplaneten) bijzonder veel gegroeid. Te beginnen met de ontdekking van de eerste exoplaneet in 1992, kennen we nu bijna 4,000 planeet kandidaten en kunnen we steeds beter hun compositie, atmosfeer en oppervlakte bestuderen.

6.1 Hoe vinden we exoplaneten?

6.1.1 Indirecte detectie methodes

Het hoofdeel van de tot nu toe gedetecteerde exoplaneten zijn gevonden met indirecte methodes, waar de invloed van de planeet op het licht van zijn ster wordt gebruikt. De meest succesvolle van deze methodes zijn de overgang methode, waar de planeet een klein beetje van het ster light blokkeert als het voor zijn ster beweegt, en de rood verschuiving methode, waar gebruik wordt gemaakt van periodieke verschuivingen van kenmerken in het spectrum van de hoofdster doordat deze om een gemeenschappelijk middelpunt draait met de planeet.

Deze indirecte methodes hebben een groot aantal planeten gevonden doordat het mogelijk is om hoge precisie te halen met relatief kleine telescopen en instrumenten, wat ze zeer geschikt maakt voor zoektochten op grote schaal. Ze zijn ook meer gevoelig voor zware planeten die dicht om hun hoofdster draaien, wat een gedeelte is van de exoplaneten parameter ruimte waar meer planeten voorkomen dan het gedeelte dat toegankelijk is voor andere methoden, zoals door middel van directe detectie.

De indirecte methodes hebben echter een nadeel, ze zijn namelijk alleen gevoelig zijn voor de planeten waarvan de baan overeenkomt met ons gezichtspunt, en daardoor wordt een groot gedeelte van de exoplaneten waarvoor dit niet geldt gemist. De detecteerbaarheid van planeten met deze methodes is ook intrinsiek gebonden aan de omlooptijd, wat betekent dat er over meerder jaren of zelfs meerdere decennia moeten worden geobserveerd wanneer er wordt gezocht voor systemen zoals ons eigen zonnestelsel.

6.1.2 Directe detectie

Als we in plaats daarvan licht van een aardachtige planeet willen onderscheiden van zijn hoofdster, hebben we instrumenten nodig die hoeken kunnen onderscheiden tot op een fractie van een boogseconde (1/3600ste van een graad) voor planeet licht dat ongeveer een miljard keer zwakker is dan dat van de ster. Deze technische uitdaging is vergelijkbaar met het vinden van een vuurvliegje dat een paar centimeter van een vuurtoren vliegt op een afstand van 200 km. Als dit kan worden bereikt, zal directe detectie een heel krachtig hulpmiddel zijn om planeten in ongeëvenaard detail te bestuderen en karakteriseren.

Deze precisie kan het beste worden bereikt door gebruik te maken van de grootst mogelijke telescopen, en geavanceerde coronagrafische optica dat het ongewenste licht van de ster kan wegfilteren en het signaal van de planeet intact laat. Deze optische technieken worden aangevuld met observatie technieken die het vermogen hebben om planeet licht te onderscheiden van sterlicht door gebruik te maken van fundamentele verschillen tussen het planeet en sterlicht. Technieken zoals polarimetrie en spectroscopie worden veel gebruikt om het zwakke planeet licht uit de zee van sterlicht waarin het zich bevindt te ontwaren.

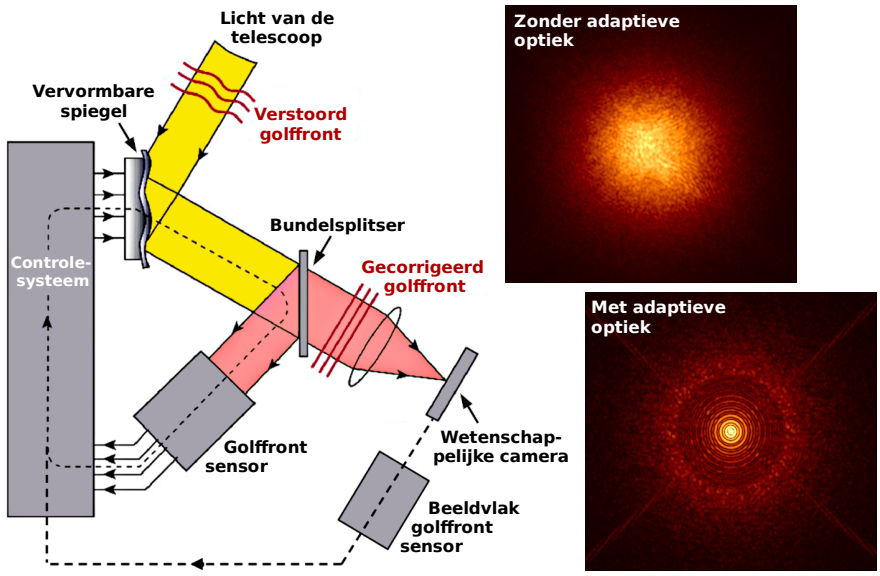
Instrumenten die bevestigd zijn aan telescopen op de Aarde hebben de extra uitdaging om te observeren door de turbulente, vervormende effecten van de atmosfeer van de Aarde. Adaptieve optische technologieën, waar een of meerdere vervormbare spiegels duizenden keren per seconden worden aangestuurd, zijn nu geavanceerd genoeg om effectief te compenseren voor deze vervormingen. In Figuur 6.1 is een schematisch overzicht van adaptieve optiek geschetst. Naarmate we licht zwakkere planeten met kleinere banen proberen te detecteren, worden vervormingen door imperfecties in de optica van het instrument zelf ook belangrijk. In het bijzonder de zogenoemde niet gemeenschappelijke pad aberraties (NGPAs), die worden veroorzaakt door gebieden van het instrument die niet goed gecontroleerd kunnen worden door het adaptieve optiek systeem, zijn op dit moment een van de grootste limiterende factoren voor de planeet zoekende instrumenten.

Een oplossing voor deze NGPAs is de techniek van beeldvlak golffront metingen, waar de information van de camera in het beeldvlak wordt gebruikt om de precieze correctie voor een perfecte afbeelding wordt bepaald. Een groot gedeelte van deze proefschrift is toegewijd aan het ontwikkelen van effectieve methodes om deze techniek toe te passen.

6.2 Hoe characteriseren we exoplaneten?

Het detecteren van planeetachtige objecten rondom andere sterren geeft ons al een idee over het aantal exoplaneten in ons sterrenstelsel. Echter, om ons begrip te vergroten van deze complexe lichamen voorbij een enkel data punt, moeten we een gedetailleerde karakterisatie studie doen met gebruik van de vele beschikbare observatie technieken.

Overgang en roodverschuiving observaties zorgen respectievelijk voor de mogelijkheid tot het meten van de straal en massa van deze planeten relatief ten opzichte van hun ster. Door het combineren van deze twee metingen kunnen we de gemiddelde dichtheid van de planeet schatten, en daarmee concluderen of deze een rotsachtige, ijsachtige of gasachtige compositie heeft.



Figuur 6.1: Een schematische weergave van een adaptief optiek (AO) systeem: een gedeelte van het licht van een telescoop wordt afgesplitst van het wetenschappelijke pad naar een golffront sensor, die een vervormbare spiegel aandrijft om te compenseren voor de verstoringen van de atmosfeer. De inzetpanelen laten gesimuleerde afbeeldingen zien van een ster voor (boven) en na (onder) AO correctie. Niet gemeenschappelijke pad aberraties treden op in het rood aangeven gebied van het instrument, en kunnen niet correct worden gemeten door het adaptieve optiek systeem. Deze kunnen worden gecontroleerd door een beeldvlak golffront sensor toe te voegen, die de afbeeldingen gebruikt van de wetenschappelijke camera om de goede correctie te meten. Dit figuur is aangepast van http://lyot.org.

Echter, hoge resolutie spectroscopische studies verschaffen verreweg het meest krachtige en veelzijdige gereedschap om bijkomende eigenschappen te bepalen van deze planeten. Niet alleen kan het gebruikt worden om de spectrale vingerafdruk van moleculen in de atmosfeer te identificeren, maar de analyse van de spectrale kenmerken zelf kan gebruikt worden om eigenschappen zoals de temperatuur en druk in de bovenste lagen van de atmosfeer te bepalen, en zelfs de draaisnelheid van de planeet. For onopgeloste planeten in een overgang, kan dit worden gedaan door het analyseren van het kleine gedeelte sterlicht dat door de planeet atmosfeer gaat, waardoor het de spectrale vingerafdruk van de planeet oppikt. Het is ook mogelijk om een spectrograaf direct te koppelen aan een hoog contrast camera systeem, wat ervoor zorgt dat ook planeten die geen overgang veroorzaken met hetzelfde detail kunnen worden gekarakteriseerd als ze via de directe methode gedetecteerd kunnen worden.

6.3 Protoplanetaire schijven en planeet formatie

Protoplanetaire schijven vormen tijdens het begin van de ineenstorting van stof en gas tijdens de eerste fase van ster formatie, en zijn op hun beurt de geboorteplaats van planeten. Daarom zal het beter begrijpen van de fysische processen die plaats vinden in deze schijven leiden tot beter begrip van planeet formatie en hoe deze systemen daarna evolueren tot volwassen zonnestelsels. Door een grote hoeveelheid van deze jonge ster systemen met verschillende leeftijden te bestuderen, kunnen we

de stukje bij elkaar leggen om een tijdlijn over de formatie van ons eigen zonnestelsel te schetsen.

Afbeeldingen van deze schijven op nabij infrarode golflengtes vertellen ons over hun buitenste lagen: het detecteren van openingen en spiralen in deze lagen geeft ons een indicatie van hun leeftijd, en potentiele regio's waar op dit moment planeet formatie plaats vind. Echter, door de zwakheid van deze schijven vergeleken met het licht van hun ster, is het maken van afbeeldingen technisch bijna net zo uitdagend als het vinden van planeten, vooral omdat het moet gebeuren over een groot oppervlak zonder dat er beeld artefacten optreden.

Polarimetrie is een extreem effectieve techniek om deze uitdaging te overkomen, doordat sterlicht van nature ongepolariseerd is maar gedeeltelijk gepolariseerd raakt wanneer het wordt verstrooid door stof deeltjes. Het meten en aftrekken van al het ongepolariseerde licht laat daardoor een afbeelding achter van de regio's waar sterlicht wordt verstrooid. Het huidige generatie hoge contrast instrument SPHERE op de VLT is erg succesvol op dit terrein, en de opgave om de prestatie van dit instrument verder te verhogen is het onderwerp van drie hoofdstukken van deze thesis.

6.4 Dit proefschrift

In dit proefschrift richt ik me op twee hoofddoelen: het ontwikkelen van nieuwe optische technieken om de uiteindelijke bereikbare contrast ratio van directe detectie methodes te verbeteren, en het adresseren van sommige openstaande beperkingen van huidige generatie instrumenten. Dit werk is verdeeld over de volgende hoofdstukken:

Hoofdstuk 2: De grootste uitdaging van golffront metingen met het beeldvlak is het effectief gebruik maken van de afbeeldingen van de wetenschappelijke camera, doordat de information over de aberraties die de lichtstraal verstoren verloren raakt tijdens het nemen van de afbeeldingen. In dit hoofdstuk presenteren we de theorie, de laboratorium implementatie en de eerste validatie op een echte ster van de coronagraphic Modal Wavefront Sensor (cMWS): een optische element dat gebruikt maakt van holografische technieken om het licht te manipuleren op een zo danige manier dat het tegelijk functioneert als coronagrafische camera en makkelijk te gebruiken lage orde golffront sensor. Na de validatie van het concept in numerieke simulaties, hebben we een prototype ontwerp van de cMWS ingezet op de 4.2 m William Herschel Telescoop (WHT) op La Palma. We hebben aangetoond dat deze cMWS gebruikt kan worden om passief de lage orde golffront fouten te meten op hoge snelheid (50 Hz frame-rate) en over een grote spectrale bandbreedte (50% in R-band), beide zijn grote uitdagingen voor de meeste beeldvlak golffront metingen. Nadat dit werk gepubliceerd was, is de cMWS verder gevalideerd als onderdeel van de Leiden EXoplanet Instrument (LEXI), inclusief succesvolle terugkoppeling voor actieve correctie tijdens observaties. Daarnaast zijn er verschillende andere cMWS optische elementen geinstalleerd in telescopen over heel de wereld, inclusief een recente succesvolle vlucht op grote hoogte met de HiCIBaS ballon.

Hoofdstuk 3: Terwijl niet gemeenschappelijke pad fouten de meest geciteerde

limiterende factor voor directe detectie instrumenten is, is het niet altijd het meest significante effect. In dit hoofdstuk ontwikkelen en testen wij een potentiele controle oplossing voor het zogenoemde langzame wind effect (LWE) dat optreedt in de hoge contrast camera SPHERE: dit is een golffront controle probleem wat de prestatie significant vermindert tijdens optimale observatie condities. In dit hoofdstuk passen wij het zogeheten “Fast & Furious” (F&F) beeld golffront meting controle algoritme toe op het specifieke geval van het LWE, en simuleren we de techniek met terugkoppeling onder realistische observatie condities die de condities van SPHERE nabootsen. Wij merken op dat dit algoritme extreem stabiel is tegen alle gesimuleerde observatie condities, waardoor het een zeer effectieve belooft te zijn om het LWE te verwijderen en als groot voordeel is gelijk implementeerbaar als softwarematige oplossing voor SPHERE.

Hoofdstuk 4: Met gelijk een vervolg op hoofdstuk 3, valideren wij in dit hoofdstuk F&F op het MITHIC hoge contrast test systeem in het Laboratoire d’Astrophysique de Marseille, om de effectiviteit te evalueren in het gevecht tegen de LWE in een realistische lab omgeving. Wij merken op dat de laboratorium prestatie van F&F consistent zijn met de simulaties, en dat het in staat is om op een robuuste wijze kunstmatig geïnjecteerde LWE effecten te verwijderen binnen vijf terugkoppeling iteraties, zelfs wanneer lage signaal ruis afbeeldingen worden gebruikt als input. Hoewel het algoritme nog steeds gevalideerd moet worden wanneer er ook tegelijk met een adaptief optiek systeem wordt gewerkt dat atmosferische turbulentie corrigeert, concluderen wij dat F&F een uitmuntende oplossing is voor het LWE in het SPHERE instrument, en dat het in staat is om op een robuuste manier in real-time golffront controle uit te voeren zelfs onder de meest uitdagende observationele condities zonder dat het de wetenschappelijk observaties verslechterd.

Hoofdstuk 5: Optimale data reductie technieken zijn net zo cruciaal als hoge precisie optica wanneer je het meeste uit de data wilt halen dat geproduceerd wordt door hoge contrast systemen. Dit hoofdstuk presenteert de inspanningen om voor de apodised Lyot coronagraaf van het IRDIS nabij-infrarood subsysteem van SPHERE een kalibratie methode te ontwikkelen dat het mogelijk maakt om op een juiste manier coronagrafische en polarimetrische data te reduceren. Dit is belangrijk want de binnenste regio van circumstellaire schijf observaties, die cruciaal zijn voor het identificeren van de centrale holten in protoplanetaire transitie schijven, worden vaak gedomineerd door artefacten van het camera systeem. Kalibratie observaties met behulp van de kleine planeet Ceres zijn gebruikt om het extinctie profiel van de coronagraaf te bepalen, en deze zijn gecombineerd met uitgebreide optische modellen om het geobserveerde signaal volledig te begrijpen. Wij concluderen dat de coronagrafische en polarimetrische observaties van protoplanetaire schijven een volledig gemodelleerde analyse vereisen om op een goede manier de niet-lineaire diffractie effecten en polarisatie effecten mee te nemen: het is niet voldoende om simpelweg te normaliseren voor de coronagrafische verliezen. Wij valideren de accuraatheid van onze kalibratieroutine op polarimetrische observaties van de goed bestudeerde protoplanetaire schijf TW Hydrae, waar we met succes de bekende centrale holte kunnen terugvinden na het corrigeren voor instrumentele effecten.

6.5 Algemene conclusies

Het werk in hoofdstukken 2 tot en met 4 kaarten het eerste doel van dit proefschrift aan, door middel van demonstraties van meerdere technieken om beeldvlak golffront metingen uit te voeren met behulp van de wetenschappelijke camera afbeeldingen. Als deze worden geïmplementeerd in de beste instrumenten, kunnen deze technieken het uiteindelijk gehaalde contrast verbeteren met meerdere magnitudes door betere beeld stabilisatie te leveren. Het gebruik van holografische technieken is een krachtige en flexibele technologie om de informatie die een afbeelding bevat aan te passen, wat op dit moment ook al toepassingen vind in andere gebieden van hoge contrast instrumenten. Voor de volgende generatie van hoge contrast instrumenten voor de ELT klasse telescopen waarvan het eerste licht wordt verwacht eind 2020 en begin 2030, laten hoofdstukken 3 & 5 zien hoe de combinatie van expertise in optica en data reductie bijna zeker nodig zal zijn om de onvoorziene uitdagingen voor deze instrumenten aan te gaan. Het streven is daarom dat dit proefschrift zal bijdrage aan een geïnformeerd ontwerp van de volgende generatie instrumenten, waardoor deze in staat gesteld worden om zwakkere planeten met kleinere banen en dus meer aardachtige planeten direct te detecteren en karakteriseren.

Chapter 7

English Summary

For centuries, philosophers, scientists, and science fiction writers alike have been fascinated by the idea of planetary systems existing around stars other than the sun. And, like Giovanni Schiaparelli claiming to see artificial canals on Mars in 1877 or early 19th Century fantasies about tropical paradises existing below the thick clouds of Venus, this interest is almost always combined with the desire to know: could there be life on these planets?

Since these early musings, our scientific understanding of these extra-solar planets (exoplanets) has grown extraordinarily rapidly. Starting with the discovery of the first exoplanet in 1992, we now know of almost 4,000 planetary-mass companions and are increasingly able to characterise their composition, atmospheres and likely surface conditions.

7.1 How do we find exoplanets?

7.1.1 Indirect detection methods

The majority of exoplanets detected to-date have been found using indirect methods, via the influence of the planet on the light of its host star. The most prolific of these are the transit method, where the planet blocks a small but detectable part of the star's light as it passes in front, and the radial velocity method, where the orbital motion of the host star about the common centre of mass of the star-planet system causes features in its spectrum to shift periodically in wavelength.

These indirect techniques have produced large numbers of detections primarily because it is possible to achieve high precision using relatively small telescopes and simple instrument designs, making them highly suited to carry out large-scale surveys. They are also more sensitive to massive close-orbiting planets, which is a more populated area of the exoplanet parameter space than those accessible by other detection methods, including direct imaging.

These indirect methods do however have the disadvantage that are only sensitive only if the planet's orbit aligns with our line of sight, and so they naturally miss a large fraction of the total exoplanet population where this is not the case. The detectability of planets via these methods is also intrinsically tied to their orbital periods, meaning that surveys must span years or even decades when looking for solar system-like planets in order to obtain detectable signals.

7.1.2 Direct imaging

If we instead wish to directly resolve the light of an Earth-like exoplanet from its host star, we require instruments which are capable of spatially resolving angles on the sky of a fraction of an arcsecond (1/3600th of a degree), and teasing out planet light which is approximately one billion times fainter than that of the star. This technical challenge

is equivalent to trying to detect a firefly fluttering just a few centimetres from a lighthouse, from a distance of over 200 km. If this can be achieved however, direct imaging offers a powerful tool to study and characterise these planets in unprecedented detail.

Such precision is best achieved by using the largest available telescopes, and advanced coronagraphic optics which filter out unwanted starlight while preserving the signal of the planet. A variety of complementary techniques which are capable of differentiating planet light from starlight by their fundamental properties, such as polarimetry and spectroscopy, are now also widely used to help tease out the faint planet light from the sea of starlight in which it is embedded.

Instruments attached to ground-based telescopes also have the added challenge of looking through the turbulent, distorting effects of the Earth's atmosphere. Adaptive optics technologies, which adjust one or more deformable mirrors thousands of times per second, are now advanced enough to effectively compensate for this distortion. The general optical layout and performance of such an adaptive optics system is shown in Fig. 7.1. As we push towards detecting fainter and closer-orbiting planets however, distortions due to imperfections in the optics of the instrument itself now also becomes a significant consideration. In particular, so-called non-common path aberrations (NCPAs), which are produced in regions of the instrument which are not properly controlled by the adaptive optics system, are currently a major limiting factor of these planet-hunting instruments.

One solution to these NCPAs is a technique called focal-plane wavefront sensing, where information from the science imaging camera is used to determine the exact correction which must be made to perfect the image. A large fraction of this thesis is dedicated to developing effective ways to perform this technique.

7.2 How do we characterise exoplanets?

Simply detecting the existence of planetary mass companions around other stars already gives us an idea of how abundant exoplanets are in our galaxy. However, in order to develop our understanding of these complex bodies beyond the level of a single data point, we need to perform detailed characterisation studies using the many available observation techniques.

Transit and radial velocity observations respectively provide estimates of the radius and mass of these planets compared to their host star. By combining these two measurements we can estimate the planet's overall density, and hence infer whether they are rocky, icy or gaseous in composition.

However, high-resolution spectroscopic study provides by far the most powerful and versatile tool for determining additional properties of these planets. Not only can it be used to identify the spectral fingerprint of molecules present in the atmosphere, but analysis of the spectral features themselves can be used to determine properties such as the temperature and pressure in the upper layers of the atmosphere, and even the spin rate of the planet. For unresolved transiting planets, this can be achieved by analysing the small fraction of starlight that passes through the planet's atmosphere, picking up its spectral fingerprint. It is also possible to couple a spectrograph directly to high-contrast imaging systems, allowing non-transiting planets to be characterised in the same detail if they can be directly imaged.

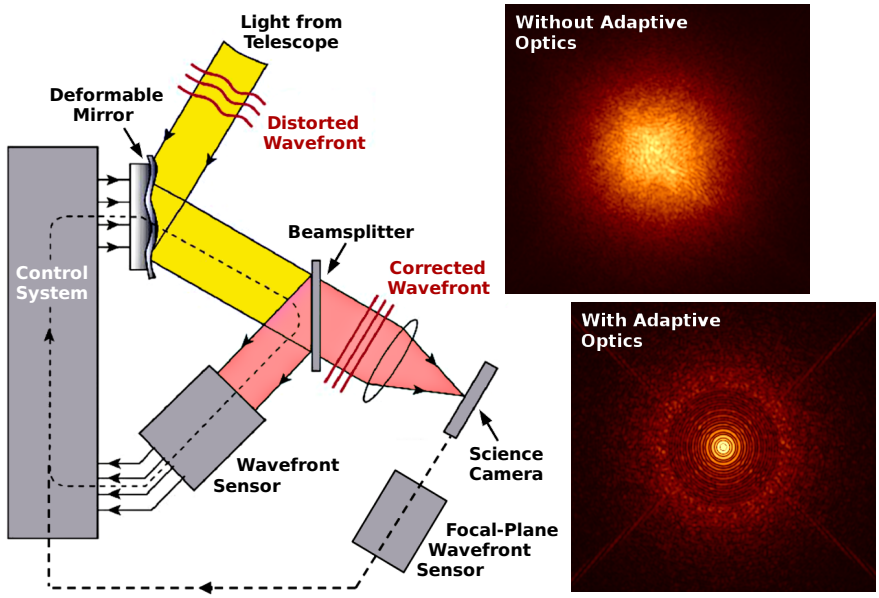


Figure 7.1: Schematic of an adaptive optics (AO) system: part of the light from the telescope is split off from the science beam into a wavefront sensor, which controls a deformable mirror in order to compensate for atmospheric distortions. The image panels show the simulated image of a star before (top) and after (bottom) AO correction. Non-common path aberrations occur in the red-shaded regions of the instrument, which are not correctly sensed by the adaptive optics system. These can be controlled by adding a focal-plane wavefront sensor, which uses images from the science camera to determine the right correction. Figure adapted from http://lyot.org.

7.3 Protoplanetary disks and planet formation

Protoplanetary disks form during the initial collapse of dust and gas that constitutes the first stage of star formation, and in turn are the birthplace of planets. Therefore, understanding the physical processes operating in these disks allows us to better determine how planets form and evolve into the objects we detect in mature star systems. By observing a large number of these young systems with different ages, we can also piece together a timeline for the formation of our own solar system.

Imaging these disks at near-infrared wavelengths tells us about their outermost layers: detecting gaps and spiral features in this surface gives us an indication of their age, and also potential regions of ongoing planet formation. However, due to the faintness of these disks compared to the light of their host star, imaging them is almost as technically challenging as finding planets themselves, especially as it must also be achieved over a large area without creating unphysical image artefacts.

Polarimetry has proved an extremely effective technique for overcoming this challenge, since starlight is naturally un-polarised but becomes partially polarised when scattered by dust particles. Measuring and subtracting all un-polarised light therefore leaves behind an undistorted image of the regions where starlight is being scattered. The current-generation high-contrast imaging instrument SPHERE at the VLT has been particularly successful in this field, and the task of further optimising the performance of this instrument is the subject of three chapters of this thesis.

7.4 This thesis

In this thesis I focus two main goals: developing new optical techniques to improve the final contrast ratios achievable by direct imaging, and addressing some of the outstanding limitations of current-generation instruments. This work is split into the following chapters:

Chapter 2: The main challenge of focal-plane wavefront sensing is in effectively utilising images from the science camera, since information about the aberrations which are distorting the light beam is fundamentally lost during the normal image formation process. In this chapter we present the theory, laboratory implementation and first on-sky validation of the coronagraphic Modal Wavefront Sensor (cMWS): an optic which uses holographic techniques to engineer the light falling on the science camera in such a way as to provide simultaneous coronagraphic imaging and straightforward low-order wavefront retrieval. After validating the concept in numerical simulations, we deployed a prototype cMWS design at the 4.2 m William Herschel Telescope (WHT) in La Palma. We show that this cMWS is capable of passively sensing low-order wavefront aberrations at high speeds (50 Hz frame-rate) and over a wide observing bandwidth (50 % in R-band), both of which are major challenges for most focal-plane sensing techniques. Since the work in this chapter was published, the cMWS has been further validated as part of the Leiden EXoplanet Instrument (LEXI), including successful on-sky closed-loop operation. In addition, a number of cMWS optics have been installed at telescopes around the world, including a recent successful flight on the HiCIBaS high-altitude balloon pathfinder mission.

Chapter 3: While non-common path aberrations are a commonly-cited limiting factor of direct imaging instruments, this is not always the most significant effect. In this chapter we develop and test a potential control solution for the so-called low-wind effect (LWE) seen in the SPHERE high-contrast imager: this is a wavefront control issue which is seen to significantly degrade the imaging performance of the instrument under otherwise optimal observing conditions. In this chapter we adapt the so-called “Fast & Furious” (F&F) focal-plane wavefront control algorithm to the specific case of the LWE, and simulate its closed-loop performance under realistic observing conditions emulating those of the SPHERE instrument. We find that the algorithm is extremely stable against all simulated observing conditions, offering an effective method of eliminating the LWE which is in principle immediately implementable as a software-only solution for SPHERE.

Chapter 4: Following on directly from Chapter 3, in this chapter we validate F&F on the MITHIC high-contrast testbench at the Laboratoire d’Astrophysique de Marseille, in order to evaluate its effectiveness in combating the LWE in a realistic laboratory environment. We find that the laboratory performance of F&F is highly consistent with simulations, and is capable of robustly eliminating artificially injected LWE aberrations within five closed-loop iterations, even when using low-signal-to-noise images as input. Although it remains necessary to validate the algorithm in parallel with a live adaptive optics system performing atmospheric correction, we conclude that F&F represents an excellent solution to the LWE in the SPHERE instrument, capable of robust real-time wavefront control under even the most challeng-

ing observing conditions without degrading the image feed for science observations.

Chapter 5: Optimal data reduction techniques are just as crucial as high-precision optics when it comes to making the most of the data produced by current high-contrast imaging facilities. This chapter presents a characterisation effort of the apodised Lyot coronagraph system of the IRDIS near-infrared subsystem of SPHERE, in order to develop a calibration algorithm capable of properly reducing coronagraphic, polarimetric image data. This is important since the innermost regions of circumstellar disk observations, which are crucial for the identification of central cavities in transitional protoplanetary disks, are often dominated by artefacts of the imaging system. Calibration observations were made of the minor planet Ceres in order to accurately determine the extinction profile of the coronagraph, and combined with extensive optical modelling in order to fully understand the observed signal. We conclude that coronagraphic, polarimetric observations of protoplanetary disks require full forward-modelling in order to properly account for non-linear diffraction and polarimetric effects: it is not sufficient to simply normalise for coronagraphic throughput losses. We validate the accuracy of our calibration routine on polarimetric observations of the well-studied TW Hydrae protoplanetary disk, successfully recovering the known central cavity feature after correcting for instrumental effects.

7.5 Overall conclusions

The work in Chapters 2 to 4 addresses the first goal of this thesis, by demonstrating multiple valid techniques for performing focal-plane wavefront sensing using science camera images. If implemented in cutting-edge instruments, techniques such as these will allow us to gain multiple orders of magnitude in final contrast performance by providing better image stabilisation. The use of holographic techniques to customise the information content provided by an image is also a powerful and highly flexible tool, which is already finding applications in other areas of high-contrast imaging. With the next generation of planet-hunting instruments for ELT-class telescopes currently due to see first-light in the late 2020s and early 2030s, Chapters 3 to 5 also highlight how combined expertise in both optics and data reduction will almost certainly be required to tackle unforeseen challenges faced by these instruments. It is therefore hoped that the work in this thesis will help to inform the design of these next-generation instruments, ultimately enabling them to directly image and characterise fainter, closer-orbiting and hence more Earth-like planets.

List of Publications

Refereed Publications

1. *The coronagraphic Modal Wavefront Sensor: a hybrid focal-plane sensor for the high-contrast imaging of circumstellar environments*
   M. J. Wilby, C. U. Keller, F. Snik, V. Korkiakoski, A. G. M. Pietrow
   2017, A&A 597, A112
2. *Laboratory verification of Fast & Furious phase diversity: Towards controlling the low wind effect in the SPHERE instrument*
   M. J. Wilby, C. U. Keller, J.-F. Sauvage, K. Dohlen, T. Fusco, D. Mouillet, J.-L. Beuzit
   2018, A&A 615, A34
3. *These are not the voids you're looking for: Using Ceres to calibrate SPHERE-IRDIS coronagraphic DPI observations of TW Hydrae*
   M. J. Wilby, J. de Boer, R. G. van Holstein, C. Ginski, J. Girard, A. Boccaletti, C. U. Keller
   In preparation
4. *First direct detection of a polarized companion outside a resolved circumbinary disk around CS Chamaeleonis*
   C. Ginski, M. Benisty, R. G. van Holstein, A. Juhasz, T. O. B. Schmidt, G. Chauvin, J. de Boer, **M. J. Wilby**, C. F. Manara, P. Delorme, F. Ménard, P. Pinilla, T. Birnstiel, M. Flock, C. U. Keller, M. Kenworthy, J. Milli, J. Olofsson, L. Perez, F. Snik, N. Vogt
   2018, A&A 616, A79

Proceedings Articles

1. *Designing and testing the coronagraphic Modal Wavefront Sensor: a fast non-common path error sensor for high-contrast imaging*
   M. J. Wilby, C. U. Keller, S. Y. Haffert, F. Snik, V. Korkiakoski, A. G. M. Pietrow
   2016, Proc. of SPIE Vol. 9909, Adaptive Optics Systems V, 990921
2. *A "Fast and Furious" solution to the low-wind effect for SPHERE at the VLT*
   M. J. Wilby, C. U. Keller, J.-F. Sauvage, T. Fusco, D. Mouillet, J.-L. Beuzit, K. Dohlen
   2016, Proc. of SPIE Vol. 9909, Adaptive Optics Systems V, 99096C
3. *The Leiden EXoplanet Instrument (LEXI): a high-contrast high-dispersion spectrograph*
   S. Y. Haffert, **M. J. Wilby**, C. U. Keller, I. A. G. Snellen
   2016, Proc. of SPIE Vol. 9908, Ground-based and Airborne Instrumentation for Astronomy VI, 990867

4. *On-sky results of the Leiden EXoplanet Instrument (LEXI)*
   S. Y. Haffert; **M. J. Wilby**; C. U. Keller; I. A. G. Snellen; D. S. Doelman; E. H. Por; M. van Kooten; S. P. Bos; J. Wardenier
   2018, Proc. of SPIE Vol. 10703, Adaptive Optics Systems VI, 1070323

5. *Multiplexed holographic aperture masking with liquid-crystal geometric phase masks*
   D. S. Doelman, P. Tuthill, B. Norris, **M. J. Wilby**, E. H. Por, C. U. Keller, M. J. Escuti, F. Snik
   2018, Proc. of SPIE Vol. 10701, Optical and Infrared Interferometry and Imaging VI, 107010T

6. *Focal plane wavefront sensing and control strategies for high-contrast imaging on the MagAO-X instrument*
   K. Miller, J. R. Males, O. Guyon, L. M. Close, D. Doelman, F. Snik, E. Por, **M. J. Wilby**, C. Bohlman, J. Lumbres, K. van Gorkom, M. Kautz, A. Rodack, J. Knight, N. Jovanovic, K. Morzinski, L. Schatz
   2018, Proc. of SPIE Vol. 10703, Adaptive Optics Systems VI, 107031T

7. *Review of high-contrast imaging systems for current and future ground-based and space-based telescopes II. Common path wavefront sensing/control and Coherent Differential Imaging*
   N. Jovanovic, O. Absil, P. Baudoz, M. Beaulieu, M. Bottom, E. Cady, B. Carlomagno, A. Carlotti, D. S. Doelman, K. Fogarty, R. Galicher, O. Guyon, S. Haffert, E. Huby, J. Jewell, C. U. Keller, M. A. Kenworthy, J. Knight, J. Kuhn, K. Miller, J. Mazoyer, M. N'Diaye, L. Pueyo, A.J. Riggs, G. Ruane, D. Sirbu, F. Snik, J. K. Wallace, **M. J. Wilby**, M. Ygouf
   2018, Proc. of SPIE Vol. 10703, Adaptive Optics Systems VI, 107031U

Curriculum Vitae

I was born on the 25th of July 1992 in Hemel Hempstead, a small commuter town in Hertfordshire, UK. When I was seven my parents and I moved to the coastal town of Paignton, in Devon, and it was under the dark skies of the Dartmoor national park that I developed a fascination with astronomy. While attending Torquay Boys' Grammar School I was also fortunate enough to have access to the school's own (sadly much under-used) on-site observatory: I will never forget the thrill of seeing the surface details of Mars with my own eyes for the first time, whilst balancing precariously on a step-ladder in order to reach the eyepiece.

In 2010 it was this passion for astronomy that inspired me to move to the other end of the country to start my undergraduate degree in Physics with Astronomy at Durham University. During this time I met a fantastic group of like-minded amateur astronomers, with whom I spent many a night standing on observatory hill, trying to maintain enough feeling in my fingers to operate the AstroSoc telescopes. The degree itself was also fascinating: alongside high-quality theoretical lectures and programming courses, one of my bachelor projects involved tracking Jovian and Martian trojan asteroids from night to night with the department's roof-mounted telescopes, in order to compute their orbits. For my master's thesis I tried my hand at a more theoretical topic, modelling the infall and merger of satellite galaxies with their host under the supervision of Prof. Shaun Cole in the Institute for Computational Cosmology.

Two lecture series which particularly inspired me during this period covered the topics of exoplanets and the optical instrumentation which enables their detection. For this reason, in September 2014 I once again upped sticks and this time left the country for Leiden, to pursue a PhD with Prof. Christoph Keller in the high-contrast imaging group. The purpose of this thesis has been to address some of the current technical limitations of direct exoplanet imaging, in order to pave the way for the next generation of ELT-class planet-finding instruments. This has included many fantastic (but exhausting) observing trips to the Roque de Muchachos Observatory on La Palma to test out new instrumentation, as well as a memorable side-project observing the northern lights in arctic Norway.

With my doctorate in hand, I now intend to move back to the UK and turn to the dark side: having recently completed the month-long *Science to Data Science* programme in London, I plan to leave academia for the glitz and glamour of the data science industry. I will nonetheless very much miss being a professional astronomer, as it remains a subject that is very close to my heart.

Acknowledgements

This thesis is the product of four years of intense study, hard work and self-development, which would not have been possible without the support of a huge number of people both in Leiden and further afield. I am hugely grateful to the network of PhDs, postdocs, staff members and non-sterrewacht friends who have enlightened, challenged, assisted, and generally kept me (mostly) sane throughout this journey. It is difficult to thank everyone properly in just a couple of pages, so I apologise in advance if I accidentally miss anyone out.

First on this list is of course Christoph: as my supervisor and guide for these past four years, you took a chance and hired the odd combination of a computational cosmologist and amateur astronomer for an instrumentation PhD position. Thank you for the opportunity: it was a steep learning curve, but the mix seems to have worked and I got there in the end! I'd also like to thank my Durham pastoral tutor Tom Theuns for being the one to originally point me in the direction of Leiden: if we hadn't had that one conversation in your office five years ago, my life would now probably be very different.

I have been very fortunate to be part of the Leiden high-contrast imaging group, and to see it grow substantially during my time at the sterrewacht. Frans, Gilles, Emiel, David, Maaike and Steven, as (pseudo-)office mates over the years you have ensured that there has always been a crazy idea or four bouncing around to keep us entertained. I'm also grateful to Jos, Christian and Rob, for keeping me in touch with the SPHERE observational community, and the various ways in which the instruments we build can be (mis-)used. Michiel and Sebastiaan: you were the ones who gave me my day-one crash course in optical lab work with ExPo and its AO system (though I'm sorry for my lack of success in mastering the latter!). Our regular observing trips to the WHT have been a real highlight of my PhD, so thanks to everyone who has been part of the ExPo and LEXI observing teams over the years, and helped keep our oxygen- and sleep-deprived brains functioning after the 3 AM dip. A big thanks as well to the WHT staff and support astronomers, especially Jürg, Luis, Neil, Émilie, Frank, Fiona and Richard, for always being ready to lend a hand when it all starts going wrong!

A huge thank you should of course also go to the observatory support staff: to Erik and Eric, for putting up with my many computer-related issues, and also to Marjan, Evelijn, Els, and Alexandra for keeping the observatory running smoothly.

I am also immeasurably grateful to the many fantastic people I have met during my time in Leiden, through borrels, tea/coffee breaks, film nights, spontaneous trips, and general gezelligheid. Heather, Jorryt, Ricardo, Niels, Gaby, Eleonora, Henriette, Tiffany, Mason, Margot: you have all made the sterrewacht far more than just a workplace. Chris, Zoe, Ann-Sofie, Nico, Aayush, Mieke and Christian: we had an absolute blast in Barcelona that first year, and it has been a pleasure getting to know you all better since then!

There are of course a number of particular personal mentions to make: David, as a fellow Trevelyanite and my flat-mate for the first year of my PhD, thank you for taking me under your wing and giving me a crash course in the quirks of life in the Nether-

lands, including (but certainly not limited to) the idiosyncrasies of the Dutch tax system. Chris, thank you for many enlightening discussions about the difference between the British and Canadian culture. I'm definitely also going to miss your unique shared sense of humour with Ann-Sofie. Eva and Valeria: it's a shame we didn't start the baking club earlier in the PhD (or maybe it was for the best, at least for my waistline!). Geert Jan and Andrew: I wish you both all the best for your North American adventure - Andrew, good luck with Trump! Luke, thank you for bringing a touch of Yorkshire to the sterrewacht. Ann-Sofie, Jeroen & Andrej: you have introduced me to an unexpected new passion in the form of swing dance, which I will definitely keep up after the PhD. Allison & Alex, thank you for the many lovely meals, kitty sitting, and for soundly thrashing me at a wide variety of board games. Nico & Wijnand: it's been great getting to know you guys (and Shaffy!) and I wish you all the best for the future.

I hope we will all stay close friends in the coming years, despite all ending up in different parts of the globe!

Ben, Jenny, Jasmine, Gideon, Amy, Vicky, Bryony, Natasha and Tim: it has also been great keeping in touch with the old "Durham crew" over the past four years. Hopefully once I'm back in the UK you'll be seeing a lot more of me! Especially to Ben, Gideon, Jasmine, Natasha, James and Piran: thank you for the many geeky physics conversations during our time in Durham, which definitely stood me in good stead for the PhD.

Finally, to my parents, who have always encouraged me to strive and fulfil my ambitions, and who are a source of sound, experienced advice whenever the path ahead is not clear. I can imagine how tough it must have been to say goodbye to your only child when I first moved to Durham to start my university education, and I have only moved further away since! You'll be glad to hear I'm not thinking of moving to New Zealand any time soon. Thank you for being a pillar of certainty upon which I can always rely. And of course, to Charlotte: thank you so much for taking the leap of faith and moving to Leiden with me, and for your unwavering support and understanding over these past four years. Completing this PhD has at times seemed like a insurmountable task, and I really could not have done it without you by my side. If nothing else you must have proof-read pretty much all of this thesis at one point or another! Here's to pastures new, and many more years (and adventures) to come. And lastly I have to thank Áine, for two years of stimulating (if mostly uni-directional) conversations, and in helping us both to define a new level of insanity.

To everyone I have had the pleasure of getting to know over the last four years: bedankt, en tot ziens!