Lasers, lenses and light curves

adaptive optics microscopy and peculiar transiting 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 zondag
1 september 2013 klokke 10.00 uur

door

Tim van Werkhoven
geboren in 1985 te Rhenen

Promotoren:	prof. dr. Christoph U. Keller prof. dr. Hans C. Gerritsen
Promotiecommissie:	prof. dr. Christoph U. Keller prof. dr. Hans C. Gerritsen

This research is supported by the Dutch Technology Foundation STW, which is part of the Netherlands Organisation for Scientific Research (NWO) and partly funded by the Ministry of Economic Affairs (project number 10 433).

Η περιέργεια είναι η αρχή της σοφίας

Contents

Contents ix

1 Introduction to adaptive optics microscopy 3
1.1 Imaging . . . 3
1.2 Optical Aberrations . . . 6
1.3 Adaptive optics . . . 8
1.4 Microscopy . . . 9
1.5 Current status & outlook . . . 13
1.6 This thesis . . . 18

2 Introduction to transiting exoplanets 23
2.1 Detection methods . . . 24
2.2 Transit analysis . . . 25
2.3 Characterising exoplanets . . . 28
2.4 This thesis . . . 28

I Adaptive optics microscopy 33

3 Coherence-gated direct wavefront sensing 35
3.1 Introduction . . . 36
3.2 Method . . . 37
3.3 Results . . . 49
3.4 Discussion . . . 54
3.5 Conclusions . . . 56

4 Sensorless adaptive optics for non-linear microscopy 59
4.1 Introduction . . . 60
4.2 Definition of the basis functions for the control of the deformable mirror . . . 62
4.3 Least-squares estimation of the unknown aberration . . . 65

4.4 Experimental results 69
4.5 Conclusion 82

II Peculiar transiting exoplanets **85**

5 Analysis and interpretation of KIC 12557548 b **87**
5.1 Introduction 88
5.2 Data reduction 89
5.3 KIC 1255b analysis 94
5.4 Cloud model 101
5.5 Conclusions 106

6 Analysis of J1407 transit fine-structure: hints of exomoons **109**
6.1 Introduction 110
6.2 Data reduction 111
6.3 Revised Stellar Parameters 117
6.4 Ring model 119
6.5 Results 122
6.6 Initial ring fitting 125
6.7 Discussion 126
6.8 Conclusions 129

Bibliography **133**

Preface

I started my PhD project in Utrecht in 2009, where I was hired to work on the implementation of adaptive optics in a microscopy setting, collaborating with the molecular biophysics group at Utrecht University. With my background in the application of adaptive optics in astronomy, this sounded like an interesting challenge, and proved to be an environment where I would learn many new things. Later, I got involved in adding adaptive optics to the *Extreme Polarimeter*, an instrument for exoplanet imaging developed in Utrecht, to improve the resolution in the presence of seeing. This instrument is now routinely used on sky at the William Herschel Telescope. After our move to Leiden Observatory, my interests diverged again and I got involved in two projects where we investigated the two peculiar transiting exoplanets.

Since these topics are after all quite different, this thesis is split in two parts. Part I deals with the implementation of adaptive optics in microscopy, part II considers two peculiar exoplanets we have investigated. To acknowledge the difference between these subjects, the introduction is also split in two, one for each part of this thesis.

Tim van Werkhoven, February 15, 2014

1 Introduction to adaptive optics microscopy

With vision being one of the senses, light is an important tool for studying one's surroundings. The behaviour of light has been studied extensively throughout history, ranging from Euclid in the 4th century BC to Newton's Opticks in 1704. To help in their studies, people built tools, starting with simple lenses possibly as far back as the 8th century BC with the Assyrian Nimrud lens (Layard, 1853). The invention of the telescope took place somewhere in the 16th century, but it is not clear whom is to be credited with the invention, if that is even applicable here[1]. Indeed, various names have been mentioned as inventors at different times, including persons from the Netherlands, England, France, Spain and Italy (Zuidervaart, 2010; Kriss & Kriss, 1998). The story of the first microscope is similar, also here it is unclear who *invented* the microscope. For investigation of minute objects, Dutch draper Antoni van Leeuwenhoek was one of the first to build microscopes, although others were active in this field as well, such as Robert Hooke (Bardell, 2004). Since then, both telescopes and microscopes have come a long way, with the ability to distinguish on one end dim, distant planets orbiting their host stars to on the other hand synaptic vesicles tens of nanometers in diameter (Marois et al., 2008; Kalas et al., 2008; Willig et al., 2006).

1.1 Imaging

When creating images by recording light, the quality of the image is influenced by everything the light finds on its path. This can be a human eye,

[1] Zuidervaart (2010) argues that the telescope was not *invented* per se, but was more a recognition of the capabilities of devices already in existence.

imaging the letters of a newspaper on the retina, or a telescope, feeding light of extrasolar planets onto a camera. In neither case is the image perfect, people often need glasses at old age and even before, visual acuity is finite. In the latter example, besides the telescope's optics, the Earth's atmosphere is of great influence on the image quality, blurring details that would otherwise reveal important information.

One solution to this is to go to space void of atmospheric influence, and indeed many telescopes operate from space, with *Hubble Space Telescope* perhaps being one of the most well-known examples. While space-based telescopes have great advantages, they are expensive and limited in size. Another solution is to build telescopes in places where the influence of the atmosphere is minimal. Often this is on high mountains, with less atmosphere between the observer and the stars, or on an island in a lake where the surrounding water creates stable air, or in places with steady, clean wind providing a clear atmosphere. Nevertheless, even the best places on Earth suffer from artefacts which ultimately limit the resolution of ground-based telescopes. These artefacts are called *seeing*, and are due to the turbulent mixing in the atmosphere of layers with different temperatures. Because refractive index is a function of temperature, these bubbles cause refraction of the incoming light, degrading the quality when imaging (Fig. 1.1). This problem was known as early as the 19th century when people built telescopes large enough to notice the effect (Hough, 1885) and indeed people suggested methods to alleviate this to some extent (Langley, 1903), but no solution for this problem existed for a long time. The effect of seeing is not unlike the image deformation seen when looking underwater, where the wavy water surface bends the light and deforms the image of what lies underneath (Fig. 1.2).

Active correction of aberrations caused by the atmosphere is therefore necessary to allow diffraction-limited imaging from ground-based telescopes. For this reason, *adaptive optics* was developed in the fields of defence and astronomy independently, one group was trying to look up, the other was looking down. In both cases, a 20 km thick layer of turbulent air was in between the object and the observer, which required active compensation of aberrations this air created.

The technique was first publicly proposed by Babcock (1953) in the United States, and independently discovered by Linnik (1957) in the U.S.S.R. While the technique had been known, it was not until the 1980's that the U.S.military built the first working adaptive optics systems at the Starfire Optical Range (SOR) in New Mexico and the Advanced Electro-Optical System (AEOS) on Maui (Reed et al., 1990; McCray, 2000). This delay was

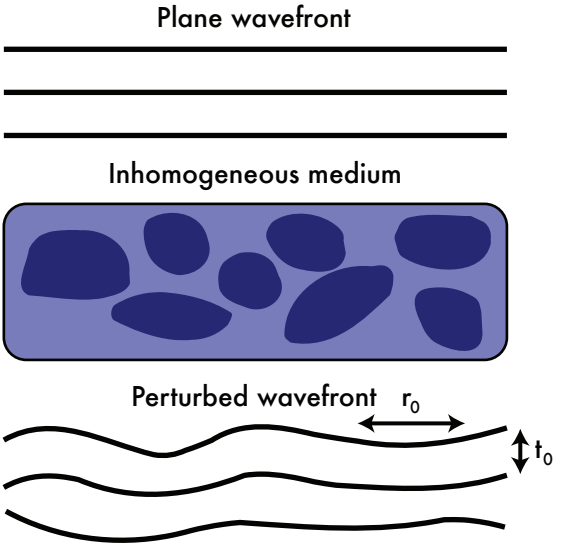


Figure 1.1: A schematic illustration of wavefront perturbation by turbid media. The plane wavefront enters at the top, traverses through the inhomogeneous medium, and comes out perturbed. The different shades in the figure illustrate differences in the index of refraction, n. r_0 and τ_0 indicate the spatial and temporal correlation of the aberration, respectively.

due to the computational requirements for such a system to operate, and without computers it was unfeasible. The declassification of these military systems began in 1988, with an important milestone being the AAS meeting in 1991 where the laser guide star assisted adaptive optics work was presented (Duffner, 2009).

Two decades after the declassification, adaptive optics is common on telescopes, for example on Gemini South (GPI, Macintosh et al., 2008), the Subaru Telescope (HiCIAO, Hodapp et al., 2008), and the Very Large Telescope (NACO, Lenzen et al., 2003). While many telescopes use adaptive optics, it is often implemented as an add-on system and not fully integrated in the telescope optics. This is not ideal as it requires extra optical components for sensing and correction. On the contrary, the future 40-meter class telescopes to be constructed the coming decade are designed with adaptive optics in mind, where the necessary optics will be integrated (McPherson et al., 2012; Bouchez et al., 2012; Ellerbroek, 2013).



Figure 1.2: A trout underwater, showing severe optical aberrations due to the non-flatness of the water surface. The light pattern on the surface are called 'caustics', and are similar to scintillation due the atmospheric turbulence (Rocca et al., 1974). The fish was not harmed during the experiment. *Photo credit: Andrew Kirk.*

As adaptive optics is evolving, it is finding more applications in various fields benefiting from this development. These fields include, but are not limited to, ophthalmology (Liang et al., 1997; Roorda et al., 2002), laser writing (Mauclair et al., 2008; Jesacher et al., 2010; Čižmár et al., 2010), as well as microscopy (Hermann et al., 2004; Booth, 2007b).

1.2 Optical Aberrations

The *points spread function* (PSF) of an optical system is the image it makes of a point source. For an instrument with a circular aperture d, the PSF is an *Airy disc* (Airy, 1834) described by

$$I(x) = I_0 \left(\frac{2 J_1(x)}{x} \right)^2 \tag{1.1}$$

with $x = \pi d \sin(\theta)/\lambda$, where θ is the angular position in the image, λ the wavelength, and J_1 the Bessel function of the first kind order one. This is the diffraction-limited performance of the instrument, and the first minimum of (1.1) is given by

$$\theta \approx 1.22\, \lambda/d. \tag{1.2}$$

The limit in (1.2) is a theoretical resolution of the system; if two objects are closer together than this, they cannot be distinguished[2]. This limit, known as the Rayleigh criterion, was developed by Lord Rayleigh (1902) at end of the 19th century (Hecht, 2002, Chap. 10). Around the same time, this diffraction limit in the context of optical microscopes was formulated by Helmholtz & Fripp (1876) and later by Abbe (1883), with the latter work also including empirical verification.

In the presence of aberrations, the diffraction limit is not reached and an inferior resolution is obtained. These aberrations perturbing the wavefront can either be due to the instrument itself, e.g. misalignment, or because of external factors, such as the atmosphere. Some aberrations are common and have been known since antiquity, including defocus, astigmatism, coma, or spherical aberration which can occur in the human eye, for example. To quantify these errors, they are often represented using Zernike polynomials (von Zernike, 1934; Noll, 1976). An example of instrumental aberrations is the *Hubble Space Telescope's* initial spherical aberrations due to a defect in the primary mirror (Burrows et al., 1991), which was repaired during the first service mission (Trauger et al., 1994).

While instrument aberrations are usually static and can be prevented by careful design, external factors such as aberrations due to the atmosphere are dynamic and unknown a priori. Mitigation of the latter therefore requires characterisation of the magnitude and kind of aberrations, and active correction of the changing conditions. Kolmogorov (1941a,b) describes the local structure of turbulence in fluids, a description also applicable to the behaviour of atmospheric turbulence, which describes a power spectrum of turbulence at various length scales. Later, Fried (1965) characterised the spatial frequency on the wavefront aberrations induced by turbulence with a coherence length scale r_0, the Fried parameter. r_0, typically on the order of centimetres for visible light, is the length over which a wavefront can be considered flat, and apertures larger than this will notice resolution degradation due to seeing. Other parameters that characterise the aberrations are θ_0 and f_G. The former is the angular equivalent of r_0, called the isoplanatic patch, and indicates an angular patch on the sky over which the aberrations are invariant, typically on the order of arcseconds, again for visible light (Korff et al., 1975). The latter is the Greenwood frequency and indicates the bandwidth requirement for correction of atmospheric turbulence, which was determined to be on the order of a hundred hertz, also for visible light (Greenwood, 1977). Together, these

[2] Although somewhat arbitrary, this definition is convenient and simple. As Lord Rayleigh (1902, p. 85) writes: *"This rule is convenient on account of its simplicity; and it is sufficiently accurate in view of the necessary uncertainty as to what exactly is meant by resolution."*.

parameters give a rough statistical description of the effect of atmospheric seeing on observations.

1.3 Adaptive optics

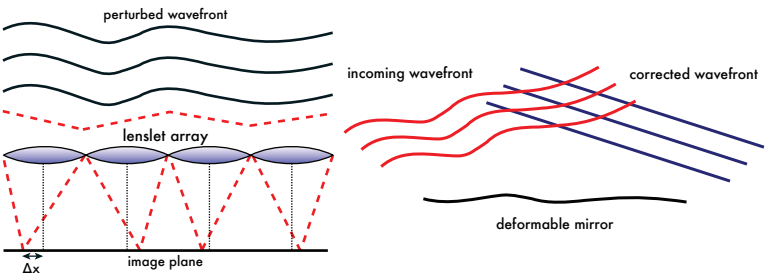


Figure 1.3: *Left:* A schematic illustration a Shack-Hartmann wavefront sensor. The telescope aperture is split up in smaller subapertures by a lenslet array. Each lens images only a part of the wavefront, such that the local phase gradient causes a displacement of the image (Δx). By measuring the spot positions one can reconstruct the wavefront. *Right:* An illustration of a deformable mirror (DM). The mirror is deformed such that after reflection the wavefront is flat.

Measuring aberrations is called wavefront sensing in adaptive optics parlance, and is the first step of correction. Because of the high frequency of visible light, there are no devices that can measure the phase of the electromagnetic wave directly, and instead all wavefront sensing devices measure the intensity of the light. Various methods exist to retrieve the phase by measuring the intensity, one common method is the Shack-Hartmann (or less common but more chronologically correct: Hartmann-Shack) wavefront sensor (Hartmann, 1900; Shack & Platt, 1971), which measures the local gradient of the phase as a displacement of a focal spot (left panel in Fig. 1.3). Using these local phase gradients the phase can be reconstructed up to a certain precision, mainly determined by the number of sub-apertures and the amount of light available. To optimally sample the aberration, the sub-apertures should be smaller than r_0 and the temporal sampling should be faster than f_G.

Once the aberration is known, the aberration is canceled by introducing the inverse shape on a correcting device with half the amplitude such that the sum of the aberration and correction yields the desired flat wavefront (right panel in Fig. 1.3). Devices to do this include deformable mirrors; a thin mirror that can be deformed using mechanical actuators, or spatial

light modulators; arrays of liquid crystals that can modify the phase. The former typically has on the order of 10 to 1000 individual actuators, and operates at speeds of kilohertz, the latter has 10^4 to 10^6 of actuators but operates at a slower rate of tens of hertz. For optimal correction the correcting device should be matched to the aberrations in both spatial and temporal resolution.

This simplified view ignores many details as various things need to be taken into account for reliable operation, including finite response speed, hysteresis, limited stroke, non-linear deformation and others. One particular issue that is becoming more important is computing power. Converting the measured aberration to actuating control is usually a matrix multiplication and therefore scales with the square of the number of degrees of freedom (typically actuators). With the number of actuators going towards the ten thousands (Kasper et al., 2010), this problem is becoming significant.

1.4 Microscopy

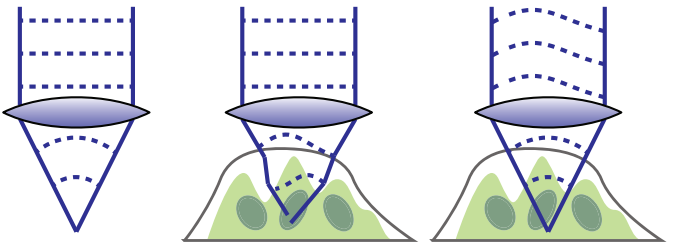


Figure 1.4: Illustration of sample-induced aberrations in microscopy. These aberrations deform the point spread function, degrading both the resolution as well as the returned fluorescence. *Left:* in absence of any aberrations, the focus is diffraction limited. *Middle:* when focussing inside a sample, the light is aberrated due to heterogeneity in the sample, giving a smeared out focus with loss of signal and resolution. *Right:* to mitigate this, one can pre-compensate the incident wavefront on the sample such that the compensation and aberration together again yield a diffraction-limited focus inside the sample. The challenge is to determine the shape to optimally compensate aberrations with.

As mentioned before, adaptive optics is also finding applications in other fields, one of these being microscopy. While it is the opposite of astronomy in terms of scale, the instruments used in both fields often deal with similar visible light, and both are affected by aberrations. When imaging with a microscope, light travels through a sample, which spatially varies

in refractive index as the Earth's atmosphere does. In Fig. 1.4 this issue is schematically illustrated. With light with similar properties, similar optics and comparable aberrations, indeed microscopy is a field where adaptive optics can be used for correction.

While the parameter space of microscopes is similar to optical telescopes (spatial, temporal, spectral resolution, field of view, collection power, etc.), there are several important differences. In microscopy, the proximity of samples under study allows illumination of those objects, and choosing the kind of illumination therefore provides a new set of parameters, such as depth penetration, contrast generation, photobleaching, and phototoxicity. Depth penetration indicates how deep light can penetrate in a sample before scattering degrades the signal and resolution significantly, and therefore how deep one can image. Contrast generation indicates that one can chemically or biologically alter a sample such that only specific parts of interest emit light when illuminated. Finally, photobleaching indicates the destruction of a dye or fluorophore, and phototoxicity describes light-induced cell death due to excessive incident laser power.

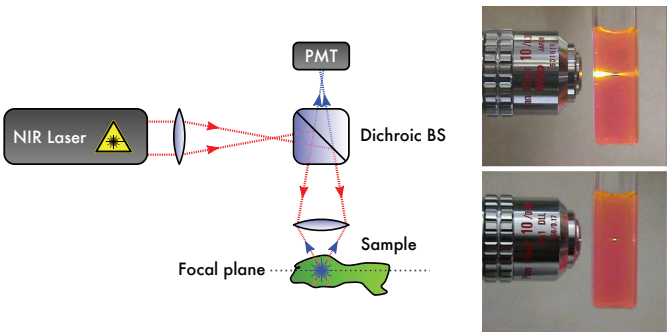


Figure 1.5: *Left:* a schematic illustration of a (multi-photon) scanning microscope. A near-infrared laser beam forms a focus inside a sample. At the focus, one or more photons to excite a fluorescent molecule which subsequently emits light at a different wavelength. This fluorescence is chromatically separated using a beamsplitter (BS) and collected by a photo multiplier tube (PMT). By moving the laser focus with respect to the sample while recording the fluorescence, an image is formed on a point-by-point basis. *Right:* comparison of single-photon fluorescence (bottom objective) and two-photon fluorescence (top), illustrating the extended emitting conical volume in the former case and the much more confined volume in the latter case due to the non-linear excitation effect. Because of the point-by-point image construction of scanning microscopes, this confinement is required to obtain axial resolution.

A specific kind of microscope is a scanning microscope, shown in Fig. 1.5, which scans the light focused by an objective through a sample (Minsky, 1961; Shotton, 1989; Pawley, 2006). The object emits light at the focus at a different wavelength, usually through fluorescence, which can be separated and subsequently collected. To form an image, the focal spot is moved with respect to the sample, either by scanning the laser beam using mirrors, or moving the sample physically with a motorized stage. By recording the fluorescence intensity as well as the location of the focus in the sample, an image is constructed in a point-by-point fashion. The resolution of such images is determined by the focal volume illuminated, which in perfect conditions is a diffraction-limited spot with lateral size

$$d \sim 0.61\,\lambda\,/\,n\sin(\theta) = 0.61\,\lambda\,/\,\mathrm{NA}, \tag{1.3}$$

with n the refractive index of the medium in which the objective is used, θ the half-angle of the objective acceptance cone and NA the numerical aperture of the objective used. For multi-photon fluorescence, where two or more photons are simultaneously absorbed to excite a fluorophore, the emitted intensity scales as $I_{\mathrm{fluor}} \propto I_{\mathrm{laser}}^{n}$, with n the number of photons absorbed. This non-linearity strongly localises the fluorescence generated to a small focal volume, creating intrinsic three-dimensional resolution (Denk et al., 1990; Zipfel et al., 2003). To obtain maximum fluorescence signal with minimal average incident power, multi-photon microscopy works best when using (sub-picosecond) pulsed lasers which deliver high instantaneous power and thus fluorescence, but low average power on the sample, thereby reducing phototoxicity.

In scanning microscopes, aberrations are caused by the sample under study perturbing the incident laser beam, among other things. While travelling through the sample towards a focus, the wavefront is aberrated, degrading the focus (Fig. 1.4 middle panel). Because the volume of the focus determines both the fluorescence intensity as well as the resolution, both quantities are degraded by these aberrations. With a spatial correlation in samples on the order of micrometers (the size of organelles in cells) aberrations become important when imaging at a depth of tens of micrometers. To solve this, one needs to pre-correct the incident laser wavefront to compensate for the sample-induced aberrations.

Adaptive optics microscopy

In scanning microscopes, it is not possible to put wavefront sensing and correcting optics in between the aberrating medium and the imaging device, as is done in telescopes. Since the aberrations occur inside the sample

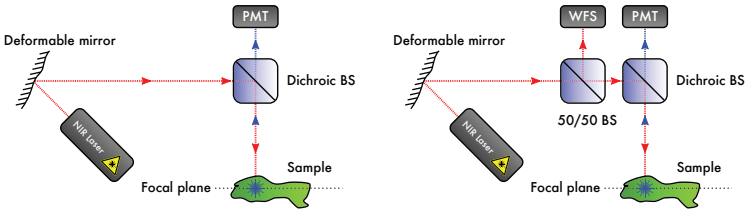


Figure 1.6: Schematic illustration of sensorless (left) and direct wavefront sensing (right) correction methods. *Left:* sensorless implementations have an extra deformable mirror compared to regular scanning microscopes. By using the fluorescence signal as optimisation metric, the deformable mirror shape can be iteratively improved, correcting aberrations. *Right:* in direct wavefront sensing additional optics direct the back-reflected laser light to the wavefront sensor (WFS). One beamsplitter (BS) transmits the fluorescence while reflecting the back-reflected laser light, the second BS directs parts of this light to the WFS. The wavefront sensor then reconstructs the aberrations, allowing direct correction with the deformable mirror.

(Fig. 1.4, right panel), the incident laser beam has to be pre-compensated such that adding the aberrations due to the specimen will yield a perfect focus inside the sample. For this reason, different wavefront sensing methods have been developed in the context of microscopy (Booth, 2007a; Booth et al., 2012), which can be divided in the following categories:

- 'Sensorless', using a proxy for the aberration such as image intensity or sharpness;
- Fluorescence wavefront sensing, measuring the wavefront off the fluorescent light;
- Direct wavefront sensing, using the backscattered laser light from the sample.

The first method has been most widely used, and uses the fluorescence intensity collected as a proxy for the aberrations (Sherman et al., 2002; Marsh et al., 2003; Booth, 2007b; Débarre et al., 2009a, 2007a; Facomprez et al., 2012; Tang et al., 2012). Since wavefront aberrations decrease the fluorescence intensity, the inverse implies that a maximum intensity gives a flat wavefront. Using either a model or heuristics, one can use the intensity of an image as feedback into a correction algorithm[3]. The simplest example

[3]For an insightful illustration of various sensorless correction schemes, see Fig. 1 in Facomprez et al. (2012).

is to actuate each wavefront mode on the deformable mirror sequentially, and set that mode such that the fluorescence is maximised. This method is quite robust albeit slow, typically requiring on the order of $mN + 1$ measurements, with N the degrees of freedom to be corrected and m a method-dependent factor in the range of 1 to 5.

The second method uses part of the fluorescence in a wavefront sensor (Ji et al., 2009; Azucena et al., 2010, 2011; Tao et al., 2013b, 2012, 2011b,a; Ji et al., 2011), which has some similarity to laser guide star methods in astronomical situations. One can either use autofluorescence, but because this is generally quite weak, a faster method is to use fluorescent beads injected in the sample, acting as artificial fiducial markers. The advantage is that this measures the wavefront unambiguously, instead of returning a metric that is an indirect proxy of the wavefront. The drawback is that this could require modification of the sample, and that it uses part of the fluorescence signal for wavefront sensing instead of imaging, thus requiring longer integration times for imaging or increasing phototoxicity and photobleaching.

The last method utilises the otherwise discarded back-reflected laser light for wavefront sensing (Feierabend et al., 2004; Rückel et al., 2006; Tuohy & Podoleanu, 2010; Cha et al., 2010; Wang & Gh, 2012). This method combines the advantages of the former two methods, it does not use fluorescence which could otherwise be used for imaging and it provides a direct measurement of the wavefront. The disadvantage is optical complexity, as the laser light passes through the sample as well as parts of the optics twice, and selecting only the back-reflected light from the focal plane is non-trivial. Odd and even symmetry aberration modes therefore behave differently as they pass through the sample (Rückel & Denk, 2007; Booth et al., 2012; Rahman & Booth, 2013).

1.5 Current status & outlook

Astronomical adaptive optics systems are used in daily operation on various instruments, and will soon have thousands of degrees of freedom and multiple mirrors and sensors for measurement and correction (e.g. GPI (Macintosh et al., 2008), SPHERE (Beuzit et al., 2008), HiCIAO (Hodapp et al., 2008), Project 1640 (Hinkley et al., 2008)). In ophthalmology adaptive optics plays a significant role as well (Wolfing et al., 2006; Martin & Roorda, 2005; Ooto et al., 2010, 2011) and is being used in clinical applications. Compared to these fields, adaptive optics in microscopy is still in relative infancy. The different methods described above have been used by various groups over the last decade, but have focussed on development of the

technology (Booth et al., 2012). While still at an early stage, already vendors are supplying off-the-shelf adaptive optics systems for microscopy, although their scope is limited (Bifano, 2011). Despite this progress, several issues need to be addressed for adaptive optics to become as common in microscopy as it is elsewhere today. Below I discuss possible opportunities for this field in the next decade.

Aberration characterisation

While turbulence characterisation in astronomy is common, this has not been done thoroughly in microscopy. Although some studies have looked into specimen-induced aberrations in scanning microscopy (Schwertner et al., 2004a,b, 2007; Zeng et al., 2012), much work needs to be done. In scanning microscopy, the geometry of the beam is important for the aberration correction (see Fig. 1.7). When imaging deep inside a sample with relative narrow field of view, the overlap of the incident laser light beam between the two outer-most points in the field of view is significant, such that a common correction is possible for the whole field of view. Conversely, when imaging shallow or over a large field of view, one expects the aberration to be strongly position dependent, requiring new correction every time. This effect has been reported by Simmonds & Booth (2013), where the authors propose multi-conjugate adaptive optics as solution. However, before resorting to more complex solutions, it might be preferable to further characterise and validate aberration energy spectra in specimen as done in astronomy.

Besides this consideration of aberration correlation-length, another important aspect is the spatial frequency of the aberrations. When imaging inside a sample, one would expect spatial frequencies on the order of $1/r_0$, where r_0 the typical length scale of the aberrations. The incident beam footprint on the sample (in air) is given by

$$d_{\text{beam}} = 2\,z\,\text{NA}/(1-\text{NA}^2)^{0.5}, \tag{1.4}$$

where z is the depth of the focus inside the sample and NA the objective numerical aperture (see Fig. 1.7). Combining these two quantities gives the degrees of freedom required for optimal correction of the aberrations, i.e. d_{beam}/r_0. For an imaging depth of 100 µm and an NA of 0.90, this gives a footprint of ~400 µm. With a typical size of organelles on the order of $r_0 = 1$ µm, this requires 400 correcting elements along one axis, or 400^2 in a plane. Relaxing these requirements to a depth of 50 µm with an NA of 0.75 gives a footprint of 113 µm, and taking $r_0 = 10$ µm the typical size of a cell still requires $10^2 = 100$ correcting elements. Of course this analysis

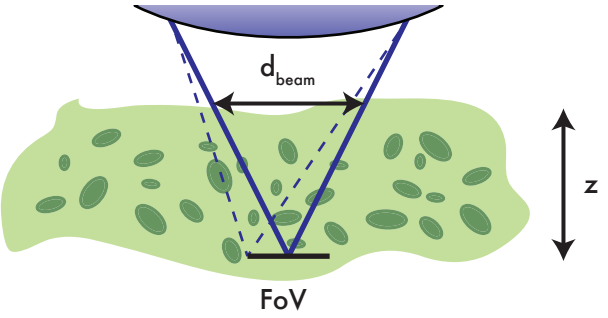


Figure 1.7: Illustration of beam and sample geometry. The focal depth z and the numerical aperture of the objective determine the beam footprint d_{beam}. Together with the typical size of refractive index variations, this determines the spatial frequency of aberrations. When scanning the beam over a field of view in the specimen (FoV, dashed lines), it samples different aberrations. The overlap between the beams determines how much the aberrations across the field of view are expected to vary, and depends on NA, depth, and the field of view.

does not include the relative strength of each spatial frequency contributing to the aberration, but this back-of-the-enveloped calculation is insightful to get a grasp on the requirements for correction. Modelling of these geometrical effects as initiated by Simmonds & Booth (2013) should therefore be done more extensively for better characterisation of aberrations, allowing for improved systems design and performance.

System aberrations

There is a distinction between system and sample aberrations, where the former are due to setup misalignments, mirror non-flatness, etc. and the latter are due to the sample, as described before. Ideally, the setup used is well-aligned and delivers diffraction-limited performance. In practice, this is not always the case, which is why significant gain can already be obtained by correcting the system aberrations with adaptive optics. After correction of these system aberrations, specimen-induced aberrations can be corrected to improve the resolution further.

A kind of aberration that is a combination of these is a depth-dependent aberration due to refractive index mismatch (Hell et al., 1993; Sheppard & Török, 1997; Török et al., 1997; Booth & Wilson, 2001). Samples are usually aqueous, with a refractive index around 1.3. When focussing a beam through a plane-parallel refractive index change, this causes spherical

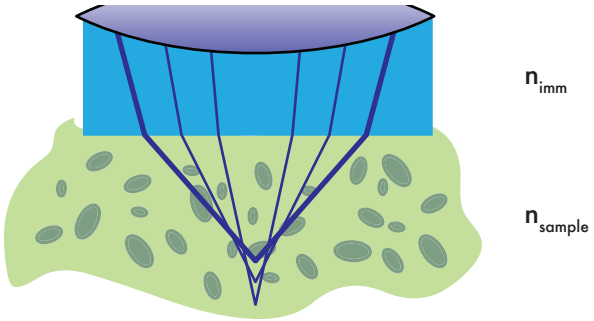


Figure 1.8: Because of the difference in refractive index between immersion fluid (n_{imm}) and the sample (n_{sample}), the converging beam suffers from spherical aberration, where focus position differs between paraxial and peripheral rays, the strength of which increases with depth. Such aberrations are not strictly due to the sample, but to a sample-system mismatch, and therefore only vary axially. The effect is exaggerated here for clarity.

aberration (see Fig. 1.8). For this reason water-immersion objectives are used where water is placed on the exit lens of the objective to fill the gap to the sample. Although this reduces the refractive index difference, some mismatch may remain. Because of this, there is a depth-dependent aberration that is due to the combination of objective and sample used, but not due to the fine-structure of the sample.

When correcting static system aberrations, these effects may not be taken into account. Therefore, significant improvement can be gained by correcting these aberrations in the case of considerable refractive index mismatch. Considering that this aberration is only depth-dependent, a single correction can be used over a large field of view where would expect no correlated in sample-induced aberrations, as described earlier and shown in Fig. 1.7. Indeed several papers report corrections of aberrations that only vary with depth, or use a single correction over a large field of view (Débarre et al., 2009a; Jesacher et al., 2009; Olivier et al., 2009). While correction of these aberrations is required for optimal resolution, the above distinction between axially and laterally-varying aberrations is often not made in literature.

Super-resolution

With the advent of super-resolution microscopy, where optical techniques are used to obtain spatial resolution beyond the classic diffraction limit,

adaptive optics can be beneficial. For example in stimulated emission depletion (STED) microscopy (Hell & Wichmann, 1994; Hell, 2003), where two beams have to overlap to within a fraction of the diffraction limit, adaptive optics can give significant advantage. In this kind of microscopy, an excitation beam is used to illuminate a sample, where it will be excited to a fluorescent state, while a depletion beam is used to deplete excited molecules except in a confined region. To this end, the depletion beam is modulated by a phase plate to create a toroidal point spread function, such that the center of the focus has zero intensity, which is the confined region where depletion does not occur. This way only molecules in this sub-diffraction-limited volume remain excited, which will subsequently fluoresce, thus achieving a resolution better than the diffraction limit.

While regular STED microscopy uses a beam shape that is not deteriorated much by aberrations (Deng et al., 2009), this shape provides super-resolution laterally, but not axially. Using instead a π-step phase plate, also axial super-resolution can be obtained (Hell, 2003), however in this case the exact beam shape is more severely influenced by aberrations. For this reason, people are using spatial light modulators (SLMs) to optimise not only beam alignment, but also correct low-order aberrations (Kromann et al., 2012; Gould et al., 2012).

In addition, because of the sensitivity to aberrations, using STED to image deep ($\sim$ 100 µm) is possible, but challenging, and indeed these applications use low order aberration correction through an objective's correction collar (Urban et al., 2011). The potential for adaptive optics here is obvious, with a closed-loop aberration correction scheme, correction far greater than only using a correction collar can be used. This will not only help to align excitation and depletion beams, but also to alleviate aberrations deeper in the sample. Compared to regular scanning microscopy, in this application adaptive optics is almost a necessity, and a suitable field for new applications.

The next decade

Although not exhaustive, the possibilities iterated above provide some ideas on where to go in this field. Characterisation of aberrations should be the initial focus for adaptive optics microscopy, which requires reliable wavefront sensing techniques. With knowledge of aberration spectra adaptive optics systems can be tailored to correct those errors that are most detrimental for image quality. Furthermore, collaboration between different fields can be beneficial when solving common (adaptive optics-related) problems. AO for super-resolution microscopy might, for example, bor-

row from extreme-AO applications in astronomy, where deformable mirrors with thousands of actuators are required. In return, adaptive optics in microscopy could provide a new market for corrector devices, such that with increasing demand there will hopefully be more incentive to develop better and cheaper deformable mirrors, benefiting all applications.

1.6 This thesis

In this thesis we implemented a direct wavefront sensing method, where we measure the aberrations from the back-scattered laser light, as well as an improved method for sensorless wavefront sensing with optimal calibration and correction.

Direct wavefront sensing

One method we have developed is a direct wavefront sensing method, where the wavefront is obtained directly from the back-reflected laser light (Chap. 3). When using the back-reflected light for wavefront sensing, one needs to select only the light that reflects back from the focal volume, and reject much stronger ghosts due to the cover glass and objective optics, for example. For this reason the direct wavefront sensing method uses a confocal pinhole and coherence gating (Feierabend et al., 2004) to suppress out of focus light.

Because the lasers used in multi-photon microscopy are pulsed, we can exploit this to use coherence gating where only light from a selected plane of interest is interfered with a reference beam. Light reflecting off different layers not arriving co-temporally with the reference beam will thus not interfere. By selecting only the light in the interference pattern, we can select the light back-reflected from a confined region in the specimen. To obtain the wavefront in a single measurement, we tilt the reference beam, generating a interference fringe pattern in the back-aperture plane of the objective (Takeda et al., 1982). The deformation of the fringe pattern encodes the wavefront aberrations. Obtaining the measurement in a single exposure improves the robustness against vibration over multi-frame wavefront sensing methods, since exposure is much shorter than the frame cadence (Rückel et al., 2006).

To measure the wavefront in a robust way, we scan the laser beam over an area while measuring the wavefront, such that artefacts such as speckles are averaged out. Additionally, we use a virtual Shack-Hartmann to smooth the retrieved phase from the fringe pattern (Rückel & Denk, 2007). This step is required because the back-reflection from the focus is not a

single point but a volume of scattering point sources, where each source contributes to the interference pattern on the camera. Because these point sources are not in one plane, this creates singularities in the phase which is problematic for phase unwrapping techniques. Using a virtual Shack-Hartmann overcomes this problem.

We implemented and verified this method successfully in a custom-built two-photon fluorescence microscope in both rat rail collagen tendons and MCF-7 cancer-cell spheroids using different objectives. We are able to correct both even and odd modes, which we interpret as the back-reflection not being purely specular which would cancel out odd aberrations due to the double-pass through the sample.

Sensorless correction

Complementary to this method, we developed a model-driven, sensorless correction scheme with robust model identification (Chap. 4). Because sensorless correction is a feedforward algorithm, the deformable mirror calibration is more important than in a closed-loop feedback system. To this end, we developed a method to construct an orthogonal set of basis modes that the deformable mirror can produce and removed modes that were poorly represented or had strong mechanical cross-coupling. Additionally, we removed the tip, tilt and defocus modes as these represent lateral and axial translations of the focal spot, and not aberrations. Since these modes cannot be detected in a sensorless system, the correction should not introduce tip, tilt and defocus, which results in undesirable translation of the image.

After obtaining a suitable basis set to drive the deformable mirror with, we developed and validated a robust method to identify the dependence of the metric on the wavefront aberrations (Antonello et al., 2012). In our case, we used the mean image intensity as the metric, which inversely scales with the square of the aberration amplitude for small aberrations (Débarre et al., 2009a). Because the metric has a global maximum (i.e. the intensity is maximum in the absence of aberrations), this implies that the model used for the metric must be positive semidefinite. While acknowledged by others, their proposed methods for computation of the model did not take this into account (Débarre et al., 2008a; Facomprez et al., 2012). Here we do enforce a positive semidefinite solution for the model, making the identification more robust against noise and providing a more accurate description of the metric.

To correct for sample aberrations, in model-based sensorless correction schemes one sets trial aberrations on the deformable mirror and records

the corresponding metric value. Using the metric model one then computes the unknown aberration vector for each mode used for correction. Various methods have been proposed to do this, either solving the paraboloid independently in each dimension (Débarre et al., 2007a), or solving it globally, using the measurement data optimally. In this work we use a minimisation scheme to minimize the aberration globally using the minimum $N + 1$ measurements to obtain a solution.

We validated the method successfully in the same custom-built microscope on collagen tendons extracted from rat tail. The metric model identification was verified by cross-validating different calibration datasets against each other, where we found good agreement. Furthermore, we investigated the variation of the model across the specimen, which is important for the correction scheme to be used. If a model varies across the sample, either the model must be re-identified, or a more general correction scheme must be used. Although we have not performed an extensive analysis, we found indication that the model varies through the specimen.

For this chapter my contribution was the construction of the experimental setup and collaboration on the experimental validation of the identification and the correction.

Lessons learned

By implementing both direct wavefront sensing and a sensorless correction scheme for correcting aberrations, we have gained insight in best practices for construction of such systems. We briefly present these findings here as recommendations for future work.

When using a 17-actuator deformable mirror in our experiments we were limited by both stroke and number of modes. In our case we were especially constrained by the limited stroke in spherical aberration, often occurring in microscopy settings. Correcting setup aberrations therefore used a significant portion of our available stroke, limiting the efficacy of the DM in subsequent experiments. Therefore, a deformable mirror that can drive the first 40 Zernike modes, and a stroke of several radians rms for at least the first 15 Zernike modes would be preferable.

Although direct wavefront sensing provides a direct measurement for the wavefront, it can be quite noisy, especially for odd modes which have a lower sensitivity. Furthermore we found significant coupling between different Zernike modes when measuring the response of the DM through the sample; when setting a single mode on the DM, the wavefront sensor would measure not only that mode but also other modes. For example we

noticed cross-talk between modes Z_5 and Z_6. This could be due to calibration issues for the deformable mirror, but we believe that there is also crosstalk between different modes due to the back-scattering in the sample. For optimal control, one would need to measure a response matrix which maps the measured modal representation of the wavefront as function of specific input aberrations set by the DM, something also mentioned in Rahman & Booth (2013).

For sensorless wavefront schemes, a simple model is perhaps the best. While theoretically optimal solutions should provide superior performance, in a laboratory-environment these assumptions can break down and do not provide significant improvement. Considering that a model might not be valid over the whole specimen, a correction scheme should be robust against this change.

In our experiments we used a piezo stage to move the sample instead of the beam. Although this is optically easier and better for implementing AO, it is probably more favourable to use scanning mirrors which have far greater bandwidth and can therefore be much faster when testing correction schemes.

The non-linear speed of a piezo stage is not an issue if position feedback is provided. Using the position feedback, the recorded fluorescence intensity can easily be remapped on a Cartesian grid. In spite of this, the limited stroke of the piezo at faster speeds does pose a problem.

Finally, we have learned it is preferable to build your own setup as opposed to using a commercial system. The ease of use these systems bring often goes hand in hand with difficulty to modify them and lack of technical documentation. Building an optical setup from scratch allows much greater control over the design which makes testing and modifying significantly easier.

2 Introduction to transiting exoplanets

The notion that Earth is a planet was not always obvious, and Earth long possessed a unique position as the centre of the universe. With the discovery of other planets in our Solar System and the realisation that the Earth revolves around the Sun, Earth's position was put into perspective, and was no longer unique, nor the centre. Even though there there were speculations about potential life beyond Earth, it took millennia before these ideas – for example advocated by Epicurus (341–270 B.C.) – could be tested (Crowe, 1999). The number of known planets was constant for a long time, before it rose again rapidly with the discovery of planets around other stars.

The detection of the first exoplanets was almost two decades ago (Wolszczan & Frail, 1992; Mayor & Queloz, 1995) and instigated a whole new field of research. However, the first detections were not planets that scientists anticipated to find, and bear little resemblance to Solar System planets. The planets found by Wolszczan & Frail (1992) were orbiting a millisecond-pulsar from which the timing could be detected using radio telescopes. While a suitable technique for detection, a millisecond-pulsar host star is not an obvious planet-hosting candidate considering the violent history these stars have gone through. Mayor & Queloz (1995) determined the radial velocity of stars by measuring Doppler line-shift using a spectrograph, and found a planet of at least half a Jupiter mass orbiting at 0.05 AU, or a period of 4 days. The authors suspect that the companion might be a very-low-mass brown dwarf stripped of its outer atmosphere. Such a close-in massive planet provides a strong radial velocity signal, but again has no comparison in our Solar System.

With the advent of instruments such as HARPS (Mayor et al., 2003),

CoRoT (Baglin et al., 2006) and Kepler (Koch et al., 1998), many planets have since been detected, and with future instruments, many more will be discovered. At the time of writing, two decades after the first detection, there are around a thousand known exoplanets, with thousands of candidates to be confirmed (see Fig. 2.1 for an overview). Most of these planets are big, massive and in close-in orbits, because detection of exoplanets is strongly affected by observation bias, but Earth-sized planets are now also within the detection range (Borucki et al., 2013).

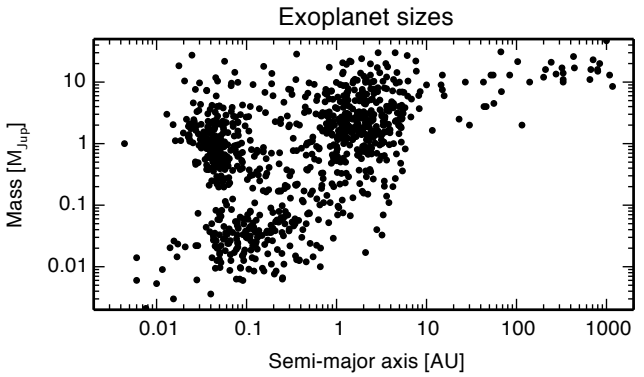


Figure 2.1: Exoplanets discovered to data, plotted as semi-major axis in against planet mass, both on logarithmic scale. Although distribution of planets in this graph are influenced by various observation biases, this plot does show the broad variety of exoplanets that exist, many having no equal in the Solar System. Data obtained from `exoplanet.eu`.

2.1 Detection methods

The reason that exoplanets for a long time remained undiscovered is because of the technical difficulty detecting them and because of the peculiarity of the planets detected so far, such as the objects described in the previous section. Compared to the host star, exoplanets are faint, with a brightness contrast of up to nine orders of magnitude for the case of Jupiter, or ten order of magnitude for Earth for visible light (Des Marais et al., 2002). In addition, exoplanets appear in close proximity to their bright host, making it even more difficult to observe these bodies directly as the telescope's points spread function smears out the starlight (Oppenheimer & Hinkley, 2009). These values are strongly dependent on

Table 2.1: Number of detected exoplanets as of early 2014 grouped by detection method. Source: exoplanet.eu.

Method	Planet systems	Planets
Radial velocity	411	548
Transit	329	435
Imaging	43	47
Microlensing	24	26
Pulsar	13	16
Total	810	1074

wavelength, and choosing this carefully helps significantly in reducing the contrast requirements. However, using techniques such as adaptive optics, coronagraphy, polarimetry, angular differential imaging, etc., this has now been achieved for tens of systems (Oppenheimer & Hinkley, 2009). The first directly imaged exoplanets were under less demanding circumstances, with relatively dim host stars and hot, large, wide-orbiting planets. These circumstances lower the required contrast mentioned before when observing in the infrared; the hot, large planet is intrinsically brighter, the parent star is fainter and the large separation between planet and host star is less demanding on the spatial resolution. Indeed the first direct detection was a roughly 4 Jupiter-mass planet orbiting at 46 AU from its host brown dwarf (Chauvin et al., 2005), and subsequent early detections were under similar circumstances (Marois et al., 2008; Kalas et al., 2008). With the upcoming 40-meter class telescopes and extreme-AO enabled instruments (e.g. GPI (Macintosh et al., 2008), SPHERE (Beuzit et al., 2008)), the observable part of the relevant phase-space will expand, uncovering more planets. In spite of this progress, it is technically less challenging to indirectly detect planets, for example by radial velocity, microlensing, pulsar timing, or by transit surveys Struve (1952). An overview of detected planet systems grouped by method is given in Table 2.1.

2.2 Transit analysis

One technique suitable for routine detection of exoplanets is by monitoring the brightness of stars. In the rare geometrical configuration that an exoplanet passes in between us and its host star, it can be detected as a temporary dimming in brightness of the star. The chance that a randomly oriented circular planetary orbit geometry produces an observable transit

is given by (Sackett, 1999):

$$P_{\text{geom}} = 0.0045 \left(\frac{1\,\text{AU}}{a} \right) \left(\frac{R_\star + R_{\text{pl}}}{\text{R}_\odot} \right) \tag{2.1}$$

with a the semi-major axis, $R_\star + R_{\text{pl}}$ the sum of the stellar and planetary radius and $\text{R}_\odot$ the solar radius. To put this into perspective, for an observer looking at the Sun from a random, distant location, there is a 1 in 200 chance that he or she can observe Earth transit the Sun, and only a 1 in 1000 chance to see Jupiter transit. On the other hand, a planet in an orbit with an 0.03 AU semi-major axis has a geometrical detection probability of roughly 15 per cent, and indeed such close-in planets have been detected abundantly (see Fig. 2.1). In the case that an exoplanet is indeed in such a transiting geometry, the maximum amount of fractional dimming for a non-limb darkened star is equal the ratio of the planet to stellar surface (Sackett, 1999), i.e.

$$\Delta F \approx \left(\frac{R_{\text{pl}}}{R_\star} \right)^2 . \tag{2.2}$$

To confirm that the stellar dimming was not spurious, two or three orbital periods should be observed, i.e. one can only detect planets in transits that have orbital periods a factor two or three less than the observing time, thus favouring close-in planets. Nevertheless, if one observes many stars for a long enough time, some will have planets oriented such that we can observe them transit their host star, and several missions are doing exactly this. The observation bias arising from these constraints therefore favors close-in orbits for large stars with large planetary-to-stellar radius ratios. On top of that, we can only confirm orbits for which the period is smaller than our observation baseline. Finally, there is a significant false-positive rate for planet detection through transit analysis on the order of tens of per cents (Brown, 2003; Collier Cameron et al., 2007) Follow up studies with, for example, radial velocity analysis are necessary to confirm detection of a planet, and not a small star, for example.

Figure 2.2 shows an example of the light curve induced by a transiting exoplanet. While this only gives the intensity as a function of time, a wealth of information can be extracted from the precise shape of the light curve. When multiple transits of the same object are observed, this gives the orbital period of the object, and thus the distance to the host star. The maxim depth of the eclipse (at C in Fig. 2.2) is a proxy for the eclipsed surface of the star and thus the radius of the planet, while a secondary eclipse, where the star eclipses the planet, yields the brightness of the planet. Changes in the

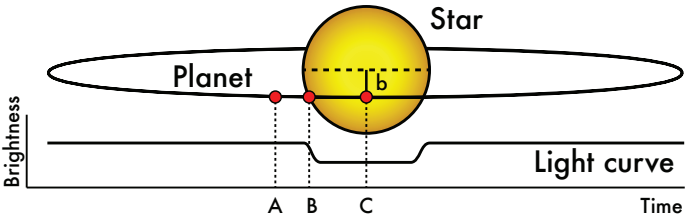


Figure 2.2: Example of a transit light curve. The points A, B and C mark the out-of-transit, ingress and mid-eclipse points, respectively, while b indicates the impact parameter (see text for details). The gradual darkening of the star towards the edge represents limb darkening. *Image credit: Wikipedia, the free encyclopedia, Nikola Smolenski*

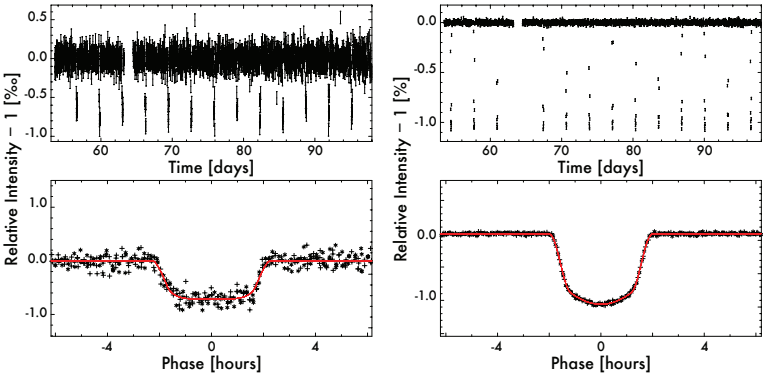


Figure 2.3: Light curves for Kepler 4b (left) and 5b (right), the first exoplanets detected uniquely by Kepler. In both panels, the top graph shows the light curve as a function of time, the bottom graph shows the curve folded with the orbital period. Kepler 4b induces an observed dimming of its host star of 0.07 per cent, and takes about 4 hours to cross the stellar disk while Kepler 5b's dimming is ten times as deep with 0.7 per cent dimming, also lasting 4 hours. Both planet's orbital period are short at 3 d. Note the different vertical scale in the different panels. *Image credit: NASA.*

period can be signs of other bodies perturbing the transiting object, such as other planets. The shape of the ingress (B in Fig. 2.2) and egress gives information on the impact parameter (b in Fig. 2.2), the shape of the object, and the intensity distribution over the stellar surface, i.e. limb darkening. The shape of the minimum of the transit light curve (at C in Fig. 2.2) gives further information on the impact parameter and stellar intensity distri-

bution. If spectra during the transit are available, the difference with the stellar spectrum can even give insight into the composition of the planetary atmosphere, without the need to resolve the planet.

Since transits only require photometry of a star, this can be used to monitor many stars at the same time. In spite of the small chance of transiting geometries, many transiting exoplanets can be detected this way. Such a survey can then be used in a statistical sense, where populations of exoplanets can be characterised. Furthermore, like radial velocity, it is a useful tool for first detection of planets, after which one can investigate the object in more detail.

2.3 Characterising exoplanets

With exoplanets detection now an almost daily routine, the next step is characterisation of the planets found. While parameters such as orbital period, mass and radius are known for many exoplanets, this still leaves room for differences between the objects. For example the Solar System's giant planets all have roughly the same radius and metallicity, but these planets are very different from one another (Oppenheimer & Hinkley, 2009).

There are exoplanets for which the atmosphere has been characterised to some extent, including HD 189733b where Swain et al. (2008) found methane. With future missions such as the Exoplanet Characterisation Observatory (EcHO), this will become more common (Tinetti et al., 2012). The exoplanets treated in this thesis are one of the few exoplanets where we can study the composition of the rocky planet itself in the case of KIC 1255b and exoplanetary rings for J1407.

Although exoplanets were peculiar objects twenty years ago, now they are discovered at rates in the hundreds per year, and we can confirm that exoplanets are indeed common. With exoplanet characterisation being the next step in this field, peculiar objects like J1407 and KIC 1255b can already provide insight into some characteristics. Transit surveys are therefore powerful tools in detecting suitable exoplanets for characterisation, and thus help develop our understanding of planets in general.

2.4 This thesis

In this thesis, we make some first steps towards exoplanet characterisation. We investigate two peculiar transiting exoplanets, *Kepler* KIC 12557548 (KIC 1255 for short) and of *SuperWASP* J140747.93−394542.6 (J1407). The for-

mer shows a 16-hour periodic transit signal with varying eclipse depth and shape from orbit to orbit, with little to no apparent correlation between consecutive orbits. This star was observed with the *Kepler* satellite While the rocky planet itself is too small to observe, it produces large amount of dust that form a cloud trailing the planet, that provides clues on the composition of the planet. The variability of the dust cloud size additionally gives insight on the dynamical processes underlying the dust generation processes. The latter was observed only once with *SuperWASP* during 2007, where the full transit lasted more than 50 days. This light curve shows a wealth of fine-structure while it is roughly symmetric around the transit epoch, consistent with the presence of a complex ring system and potentially exomoons. If this is indeed the case, J1407 is not unlike Saturn in both having complex rings and moons. The presence of a ring structure would also here provide a unique opportunity for investigation of the matter which makes up planets.

Both objects, while being different and peculiar in their own way, give a rare window of opportunity for investigation of the composition of exoplanets. The former exposes its rocky material through evaporation by the host star, the other modulates its host's light, imprinting its signature on the light curve observed by us. It is therefore likely that these planets will be the subject of much scrutiny in the coming years.

Kepler KIC 1255

For planetary transits, the geometry is often symmetrical, and this yields a symmetric transit signal, such as the transits of Kepler 4b and 5 (Fig. 2.3). Deviation from symmetry therefore indicates an asymmetry in the system, which is the case with KIC 1255. The light curve, shown in Fig. 2.4 shows a slight bump just before ingress, a steep ingress, and a gradual egress. Another curiosity of this system is the variable transit depth, ranging from less than 0.1 per cent to more than 1 per cent. The absence of some transits implies that the eclipse signal by the planet is negligible, and that the planet itself must therefore be tiny.

This target was found by Rappaport et al. (2012) and investigated by others in detail (Brogi et al., 2012; Budaj, 2013; Perez-Becker & Chiang, 2013). It appears that the close proximity to the host star of 0.013 AU (roughly 4 stellar radii) is causing the evaporation of the planet. The stellar flux would evaporate the rocky surface of the planet, causing a dust cloud to eclipse the star. According to this scenario, the planet mass would be slightly less than twice the mass of the Moon (0.02 Earth mass), and lose more than 0.1 Earth-mass per gigayear (Perez-Becker & Chiang, 2013), which explains

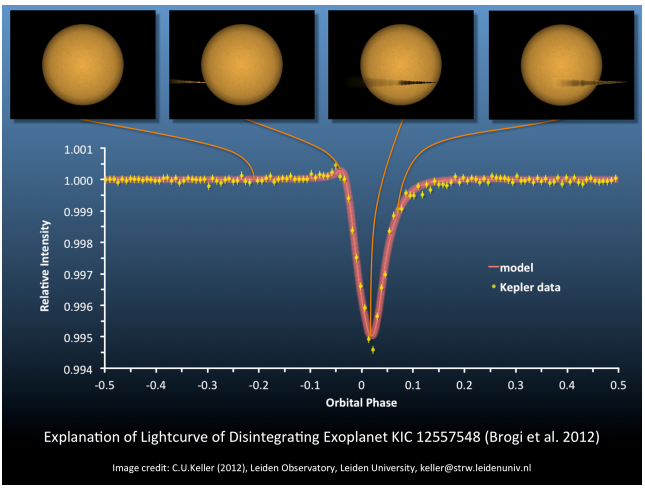


Figure 2.4: Schematic illustration of a KIC1255b model explaining the various features of the light curve.

the absence of some transits: the planet is tiny.

It turns out that a model of an exponential tail fits this data well (Brogi et al., 2012). The bump before ingress can be explained by forward scattering just before ingress due to the dust particles, the deep minimum is when most dust covers the stellar disk and the egress can be interpreted as a long tail of dust trailing the planet. To explain the variable transit depth, the dust evaporation must occur on timescales of less than one orbit, in this case 16 h, and that the dust generation must be highly variable, sometimes ejecting big clouds causing more than 1 per cent dimming, while at other times generating hardly any dust to account for the shallow transits. This is not unlikely considering the planet's proximity to the host star and its influence on the planet.

In Chap. 5 of this thesis, we characterise the dust cloud tail surrounding the planet in more detail, looking at the dust cloud generation and longevity. We investigate the light curve on a per-transit basis, to find correlations between consecutive transit and egress depths, evolution of the transit depth, and differences between different transit depths. We extended the 1-D model developed in Brogi et al. (2012) to a 2-D, 2-component model, where we model the dust cloud by an opaque, spherical core and an exponentially decaying tail attached to it. Using this improved model and the available short cadence data we can investigate the dust cloud

changes from transit to transit. While we note that the model we use is not a unique solution, this analysis does provide constraints for future, more thorough modelling.

SuperWASP J1407

Another peculiar object is J1407, whose observed intensity shows a long transit with maximum dimming of 95 per cent (Mamajek et al., 2012). Although the light curve is roughly symmetric, with the ingress and egress each lasting about 25 days, the fine-structure is different between the two when folding the light curve around the eclipse midpoint. While binary stars or a circumstellar disk can explain parts of the light curve, the only scenario that seems compatible with the observations is a ringed disk around a companion to J1407. Mamajek et al. (2012) proposed and investigated this hypothesis, and found that an initial 4-ring model explains the approximate shape of the light curve.

The intriguing part about this target is the myriad of fine-structure in the light curve, with large light curve gradients of up to tens of per cent per night and inflection points during individual nights, and big intra-night differences in both dimming and light curve gradient. Using a simple face-on ring model, one can compute the light curve gradient for a given orbital speed. By fitting this to the data, one can calculate the minimum speeds implied by these light curve-gradient. We performed this analysis in Chap. 6 and found minimum speeds of up to 33 km s^{-1}, which are incompatible with the maximum orbital speed of 22 km s^{-1} found by Mamajek et al. (2012), even assuming optimal geometric conditions for the ring model.

Using these implied speeds from the light curve and the transit duration of 54 days, we also derived a minimum ring-disk diameter in the range of 0.05 to 0.5 AU, which is 16 to 160 times as large compared to Saturn's 480×10^3 km outer E-ring radius. Furthermore, the maximum semi-major axis for a circular orbit would be between 0.8 to 10 AU.

Since these speeds are incompatible with the constraints presented by Mamajek et al. (2012), this simple model fails to explain the data. One possible solution to this problem could be azimuthal structure in the ring, such as the 'spokes' observed in Saturn's B-rings (Smith et al., 1982). Under favourable orientations, the rotational speed of such spokes could add to the planet speed to account for the observed light curve gradients. Although a possible explanation, it is not clear whether this is indeed the case, as it requires a careful and unlikely alignment of the disk and spokes to match the observations.

Part I

Adaptive optics microscopy

3 Snapshot coherence-gated direct wavefront sensing for multi-photon microscopy

Based on Tim van Werkhoven, Jacopo Antonello, Hoa Truong, Michel Verhaegen, Hans Gerritsen and Christoph Keller *Opt. Express* (accepted)

Deep imaging in turbid media such as biological tissue is challenging due to scattering and optical aberrations. Adaptive optics has the potential to compensate the tissue aberrations. We present a wavefront sensing scheme for multi-photon scanning microscopes using the pulsed, near-infrared light reflected back from the sample utilising coherence gating and a confocal pinhole to isolate the light from a layer of interest. By interfering the back-reflected light with a tilted reference beam, we create a fringe pattern with a known spatial carrier frequency in an image of the back-aperture plane of the microscope objective. The wavefront aberrations distort this fringe pattern and thereby imprint themselves at the carrier frequency, which allows us to separate the aberrations in the Fourier domain from low spatial frequency noise. A Fourier analysis of the modulated fringes combined with a virtual Shack-Hartmann sensor for smoothing yields a modal representation of the wavefront suitable for correction. We show results with this method correcting both DM-induced and sample-induced aberrations in rat tail collagen fibres as well as a Hoechst-stained MCF-7 spheroid of cancer cells.

3.1 Introduction

Microscopy is of fundamental importance in biological and biomedical research and applications. This is to a great extent due to the development of imaging methods that provide 3D imaging in (optically thick) specimens, such as confocal microscopy (Minsky, 1961; Shotton, 1989; Pawley, 2006) and multi-photon excitation microscopy (Denk et al., 1990; Zipfel et al., 2003). However, these techniques suffer from specimen-induced aberrations (Schwertner et al., 2004a, 2007), which results in reduced depth penetration, loss of spatial resolution, and increased phototoxicity (de Grauw et al., 1999) (see Fig. 3.1).

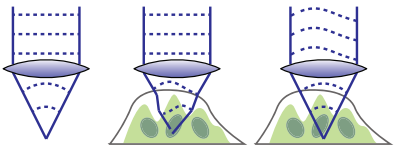


Figure 3.1: *Left:* without sample, a plane wave focusses into a diffraction limited spot. *Centre:* in the presence of a perturbing sample, the focus is distorted leading to loss of signal and resolution. *Right:* using adaptive optics to pre-aberrate the wavefront, the focus can be improved.

Adaptive optics (AO) has been successfully used in scanning microscopy to correct these aberrations, demonstrating the potential for AO in this field (Booth, 2007a,b). One class of methods are sensorless methods (Marsh et al., 2003; Booth, 2007b; Débarre et al., 2009a, 2007a; Facomprez et al., 2012) that iteratively optimise a metric such as the mean intensity or sharpness based on the fluorescence signal, without using a dedicated wavefront sensor. Alternatively, the fluorescence can be used to measure the wavefront directly, either using fluorescent beads (Azucena et al., 2010, 2011), or fluorescent proteins (Tao et al., 2013b, 2012, 2011a,b) as the biological equivalent of astronomical guide stars (Foy & Labeyrie, 1985; Fugate et al., 1991).

Another strategy uses the back-reflected excitation light for wavefront sensing. In this case, selecting only the light from the focal plane and rejecting the light reflected elsewhere along the optical path is a major challenge. Previous attempts solved this with coherence-gated wavefront-sensing methods (Feierabend et al., 2004; Rückel et al., 2006; Tuohy & Podoleanu, 2010) or a pinhole (Cha et al., 2010) to select light from a plane of interest. The former methods use a low-coherence interferometer with either temporal- or spectral modulation, similar to optical coherence tomography (OCT) (Huang et al., 1991), while the latter uses a pinhole to select a volume overlapping with the laser focus.

In this paper we describe an optical system that combines coherence-gating with interferometric wavefront measurements by interfering the back-reflected beam from the sample with a tilted reference wavefront. The coherent mixing of these two signals produces a fringe pattern where the wavefront aberrations are encoded in the fringe deformation, separating it in the Fourier domain from low-frequency noise. A confocal pinhole provides axial sectioning to reject ghost reflections from e.g. the cover glass, while additionally smoothing out speckles due to scattering samples. The wavefront is then recovered using Fourier transform fringe analysis methods (Takeda et al., 1982; Macy, 1983; Bone et al., 1986; Takeda, 1990).

This paper is organised as follows: Section 3.2 describes the optical setup, coherence gating, the fringe analysis and the wavefront reconstruction. Section 3.3 presents results for deformable mirror-induced as well as sample-induced aberrations and corrections. Finally, Sect. 3.4 discusses advantages and potential improvements of this method.

3.2 Method

We use the back-scattered near-infrared light to measure the wavefront aberration induced by (biological) samples in a scanning microscope setup. This light is usually discarded as the specimen is imaged using the fluorescent light. By interfering the back-scattered light with a flat reference wavefront with a known tilt, we measure a fringe pattern in a re-imaged back-aperture plane of the objective. The wavefront aberrations in the sample arm are then encoded as deformations of the fringe pattern (Takeda et al., 1982).

Optical setup

Fig. 3.2 shows the optical schematic for coherence-gated wavefront sensing. We use a Coherent Chameleon Ultra II Ti:Sa 140 fs pulsed, near-infrared laser beam as the source, which is expanded 12 times to a 14 mm beam (L1, L2 in Fig. 3.2). The beam is divided by a non-polarizing 50/50 beamsplitter (BS1, BS017 Thorlabs) into a *reference beam* (long-dashed) and a *sample beam* (solid). We stop down the beam to 9.5 mm before it is reflected by the deformable mirror (DM, Okotech MMDM 19-channel deformable mirror) such that we only use the central, well-controlled region (Paterson et al., 2000). The beam is then re-imaged by lenses L3 and L4 (both f=300 mm, Thorlabs NIR achromats) onto the back aperture of the objective. Since the deformable mirror can only introduce wavefront cor-

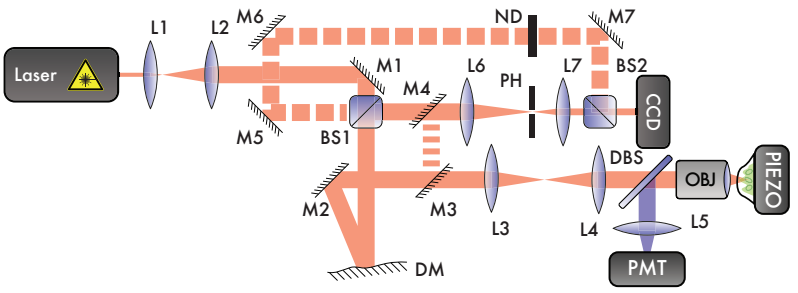


Figure 3.2: Schematic optical setup showing the coherence-gated wavefront sensing principle. The pulsed laser beam is expanded (L1, L2) and then split (BS1) into a reference (long dashed) and a sample (solid) beam. In the sample beam the deformable mirror (DM) is re-imaged onto the objective's back aperture (L3, L4). The sample is mounted on an XYZ piezo stage, and the fluorescence signal is separated by a dichroic beam splitter (DBS) and detected by a photomultiplier tube. The back-reflected light in the sample beam is again reflected by the DM and directed into the wavefront sensor (WFS) by BS1, where the back aperture is re-imaged onto the CCD (L6, L7), after being spatially filtered by the pinhole (PH). The reference beam's optical path length and intensity are matched to that of the sample beam. The beams are combined on the CCD (BS2) where the temporally overlapping laser pulses interfere. By tilting folding mirrors in the reference beam (M6, M7) we tune the carrier frequency of the fringe pattern. Mirrors M3 and M4 are optional and are used to measure the influence of the DM for calibration purposes.

rection in one direction, we bias the mirror to mid-actuation to allow for both positive and negative path-length correction. The collimated beam reflecting off the deformable mirror is therefore converging, which is compensated for by lenses L3 and L4. We use a Nikon 20×/0.75 NA multi-immersion objective with water, which has a back-aperture diameter of 16 mm; the effective numerical aperture is 0.42. Additionally, we use a Nikon 40×/0.90 NA air objective whose 9.5 mm back-aperture is filled.

The sample is mounted on an XYZ piezo stage (PI, Nanocube P-611.3S). Fluorescence emitted from the sample is split off by a 705 nm cut-off dichroic beamsplitter (DBS, Semrock, FF705-Di01-25x36) where we use an additional 720 nm short-pass filter (Semrock FF01-720/SP-25) to block stray laser light. The fluorescence is then imaged onto a 600 µm fibre feeding a photo multiplier tube (PMT, Hamamatsu, GaAsP photocathode H7422P-40).

Laser light reflected back from the sample passes back through the objective and the DM into the wavefront sensor arm via BS1. The sample

beam is de-magnified by a factor 3 by lenses L6 and L7 (f=300 mm, 100 mm, Thorlabs NIR achromats) onto the wavefront-sensing camera (CCD, AVT, Guppy Pro F-033b, 12bit, 656×492 pixels, each 9.9×9.9 µm). We use a 400 µm pinhole (PH, custom made) confocal to the focal plane to reject unwanted light, e.g. from the cover glass, similar to Rahman & Booth (2013). At the position of the pinhole, the diameter at the first minimum of the diffraction-limited Airy disc is 62 µm at a wavelength of 750 nm. The reference beam's optical path length is matched to the sample beam path length by moving mirrors M5 and M6, and the intensity is matched by a neutral density filter (ND, reflective, Thorlabs). The beams are combined by a 50/50 beamsplitter (BS2, BS017 Thorlabs). Mirrors M3 and M4 are optional and used for the deformable mirror calibration.

We use a data acquisition board (National Instruments, PCI-e 6259) for reading out the PMT signal, as well as for controlling and reading out the piezo stage. We drive the piezo stage using a sinusoidal signal for the fast axis with a typical frequency of 20 Hz. Simultaneous to driving the piezo stage, we record the stage position using the strain gauge sensors. With a constant PMT integration time this yields non-square pixels due to the varying piezo speed, which is corrected in software using the position feedback. The camera is connected to a Firewire controller, and the deformable mirror is connected to a stabilised high-voltage amplifier, which is connected to a 32-channel PCI DAQ card. We control the PMT and piezo stage using a Windows computer running LabView, and the camera and deformable mirror are controlled by a Linux computer with Python.

Coherence gating

The sample and reference beams have identical optical path lengths when arriving at the CCD. Since the laser is pulsed, this creates interference *only* between the reference wavefront and light reflected off a thin layer of interest in the sample to which the optical path length of the reference beam is tuned. This principle is shown graphically in Fig. 3.3. Light that is reflected outside of this layer does not arrive at the camera coincidentally with the reference beam pulse and does not interfere, although it does add a background signal. This background signal is reduced by the confocal pinhole.

As the reference beam wavefront is static, the interference pattern only depends on the wavefront aberration in the sample beam. By inducing a wavefront tilt in the reference beam, we generate an interference fringe-pattern on the CCD. The deformation of these fringes contains information on the wavefront aberration, which are due both to the sample being studied and static aberrations in the optical system (Takeda et al., 1982;

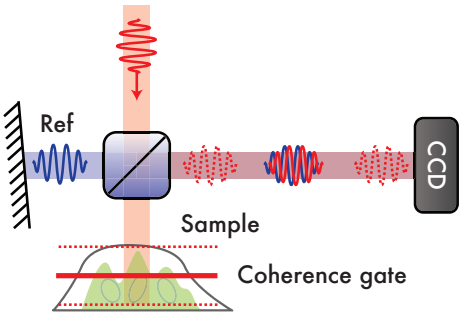


Figure 3.3: Excitation light entering from above is split into a reference and sample beam. In the sample, the light is reflected at different depths, but only backscattered light from a thin layer, the *coherence gate* (solid red), has the same optical path length as the reference beam and will interfere to form fringes on the CCD. Pulses that originate from deeper or shallower in the sample (dashed red lines) arrive later or earlier (dashed red waves) than the reference pulse and do not contribute to the fringe pattern and add a static background.

Takeda, 1990).

We can recover the complex wave using Fourier analysis of the fringe pattern. When bandpass-filtering the modulated fringe pattern in Fourier space, both low- and high-frequency noise such as ghost reflections and other spurious signals will be strongly suppressed.

The coherence length, or the spatial extent over which the pulses of the two beams overlap, is given by the correlation of the reference and sample pulses. With 140 fs hyperbolic secant-squared pulses generated by the laser, the FWHM of the autocorrelation is 65 µm in air. We measured a FWHM coherence length of $l_c = 83.7 \pm 0.3\,\mu\text{m}$ using a mirror as a target, which is 29 % broadened compared to the value given by the manufacturer, most likely by group velocity dispersion due to the optics in the setup. In an aqueous fedium with $n \approx 1.3$ this corresponds to a coherence gate length of $l_c/2n = 32\,\mu\text{m}$ (Wang et al., 2012), similar to values obtained byFeierabend et al. (2004); Feierabend (2004).

Although dispersion compensation could be used to improve the fringe contrast or minimize the coherence gate Feierabend (2004); Rückel et al. (2006), these were not as much a limiting factor here as the speckle-induced amplitude aberrations and the limited deformable mirror stroke and number of actuators. Further exploitation of the method could, however, benefit from such correction.

Confocal pinhole

We use a confocal pinhole in the back-reflected light detection path before the wavefront sensor for two purposes: it rejects out-of-focus light, and the pinhole acts as a low-pass filter that smoothes out speckles in the coherence-gated image. It thereby acts as a low-pass filter for wavefront (and amplitude) aberrations with a cutoff that is commensurate with the number of degrees of freedom of the DM.

The first benefit works in tandem with the coherence gating. The axial depth sectioning of a detection pinhole with diameter d_{ph} in the geometrical regime is given by

$$\Delta z = \sqrt{2}\, d_{\mathrm{ph}} / M \tan \alpha \tag{3.1}$$

where M is the magnification of the system and α the objective acceptance angle of the detection path (Brakenhoff et al., 1990). The 400 µm pinhole in our case gives $\Delta z = 40\,\mu\mathrm{m}$ for the 20× objective in water and $\Delta z = 6.8\,\mu\mathrm{m}$ for the 40× objective in air.

Empirically we determined the axial sectioning of the pinhole by measuring the integrated intensity on the wavefront sensing camera off a cover glass as function of the axial position of the cover glass. We found a full width at half maximum (FWHM) axial resolution of $\Delta z = 46\,\mu\mathrm{m}$ and $\Delta z = 6.8\,\mu\mathrm{m}$ for the 20× and 40× objectives, respectively.

Fourier transform fringe analysis

The fringe analysis method is based on Takeda et al. (1982); Takeda (1990), who describes a heterodyne technique to encode the low-frequency aberrations with a higher frequency carrier signal, separating it from low-frequency noise. Here we apply the method for coherence-gated wavefront sensing in scanning microscopes. Using Fourier filtering we extract the wavefront aberrations from the fringe pattern deformation. In this experiment we use a spatial carrier signal to allow wavefront measurements in a single frame, whereas in time-domain method such as for example in Rückel et al. (2006), four frames are required for a single wavefront measurement.

The fringe pattern is described by Eq. (1) in (Takeda et al., 1982)

$$g(x,y) = a(x,y) + b(x,y)\cos[2\pi f_0 x + \phi(x,y)]\,, \tag{3.2}$$

where $\phi(x,y)$ represents the wavefront and $a(x,y)$ and $b(x,y)$ represent irradiance variations. In this context, $a(x,y)$ is a static background from

e.g. ghost reflections and $b(x, y)$ are amplitude variations in the fringe pattern, such as those due to speckles. f_0 is the known spatial carrier frequency induced by the tilt in the reference wavefront, in this case along the x-axis. $a(x, y)$ can be measured by de-tuning the coherence (such that the second term is zero). $b(x, y)$ can be determined by modulating the fringe pattern by varying the path length difference between the sample and the reference beams. Using the Fourier analysis described in (Takeda et al., 1982), we can recover the complex phase from the real fringe pattern.

Note that temporal variations slower than the exposure time (4 to 8 ms) do not influence the wavefront measurements because the phase is recovered from a single measurement. During the experiments the fringe contrast was sufficient for wavefront sensing in approximately 95 per cent of the frames, while fringes moved significantly from one frame to the next. We did not experience any problems with internal seeing and mechanical vibrations.

Wavefront reconstruction

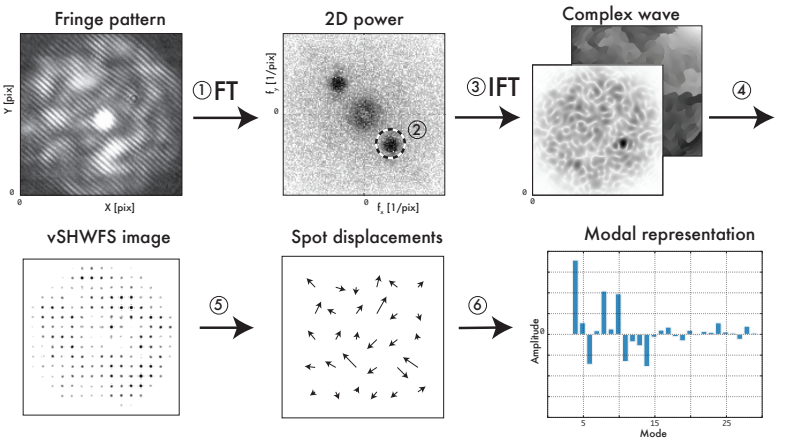


Figure 3.4: Schematic illustration of the fringe analysis method. The fringe pattern recorded with the camera (*top left*) is Fourier transformed (*top center*). One sideband at $\vec{f}_0$ is isolated and inverse transformed, yielding the complex wave (*top right*). The wave is propagated through a virtual Shack Hartman wavefront sensor to obtain a spot pattern (*bottom right*), which is reduced to a vector of spot displacements (*bottom center*). By multiplying this vector with a response matrix of a certain basis (e.g. Zernike), we obtain a modal description of the wavefront (*bottom right*).

Calibration

The wavefront sensor is calibrated before use by removing system aberrations (including the initial non-flatness of our DM), recording a flat-field and the carrier frequency, and measuring the non-common path aberrations between the focal plane and the wavefront sensor.

We first maximise the two-photon fluorescence signal from a fluorescein solution by optimising the deformable mirror shape using a Nelder-Mead simplex algorithm (Nelder & Mead, 1965; Antonello et al., 2012). Since the fluorescence signal is given by the integral over the PSF squared, maximising the intensity minimises the volume of the PSF (So et al., 2000). This system correction removes static aberrations in the setup, including the initial non-flatness of the deformable mirror. We observe a significant improvement of tens of percents of the fluorescence signal after this correction, indicating that we successfully correct static aberrations.

We obtain a flat-field image by de-tuning the coherence gate and averaging over 250 frames. This image is used to divide the raw frames thus correcting for $a(x, y)$, and the inhomogeneous illumination of the CCD. Inhomogeneity in $b(x, y)$ is accounted for by tuning the carrier frequency $\vec{f_0}$ such that the sideband is well-separated from the low-frequency noise peak in the Fourier domain. We verify this by inspecting the power spectrum of the fringe pattern. Through trial and error, we find that approximately 30 fringes across the camera is an optimal value for the fringe analysis method. The radius of the region of interest around $\vec{f_0}$ (dashed circle in Fig. 3.4, panel two) is set to half the length of $\vec{f_0}$.

Finally, we use a scattering-bead solution consisting of fluorescein-containing silica beads (d=106 nm, 2×10^{14} mL^{-1} = 2×10^{2} µm^{-3}) suspended in water to measure the response of the wavefront sensor on the previously defined flat wavefront, where we average over 200 frames to obtain a robust measurement. A scattering-bead solution is more suitable than a mirror since this back-scattering is more similar to what is expected from real samples. This wavefront measurement corresponds to the non-common path errors between the focal plane and the wavefront sensor, and is subtracted from subsequent measurements to obtain the real wavefront errors induced by the specimen.

Wavefront recovery

To obtain the wavefront phase, we record an image using a CCD and process it as described above. This image is apodised and Fourier transformed, and a region around $\vec{f_0}$ is isolated, rejecting both the low-

frequency and high-frequency noise as suggested by Peck (2010) (see Fig. 3.4 upper-centre panel). The crop window size is a tradeoff between the spatial resolution of the wavefront (larger window) and noise rejection (smaller window). Using a radius of approximately $0.5|\vec{f}_0|$ for the circular crop window, we were able to successfully isolate the sideband. We then inverse Fourier transform the isolated sideband to yield the complex wave.

In theory, the phase can be recovered from the complex wave (Eq. 6 from (Takeda et al., 1982)), however this only works robustly if the fringe pattern is well-behaved. In our experience, this straight-forward approach produces unreliable phase maps with discontinuities and singularities. While the former can be solved by unwrapping the phase (Itoh, 1982; Ghiglia & Pritt, 1998), singularities are more difficult to overcome.

We use a virtual Shack-Hartmann wavefront sensor (vSHWFS) to mitigate this, as proposed by Feierabend et al. (2004); Rückel et al. (2006). We take the complex wave and divide it into 16 by 16 sub-apertures of 20 by 20 pixels each, comparable to Rückel et al. (2006). For each of these sub-apertures we compute the power of the Fourier transform, which gives an image of the complex wave at the location of the sub-aperture. This yields an image like one obtained with a Shack-Hartmann wavefront sensor (Fig. 3.4 bottom-left panel), where the displacement of the spot position from the center is proportional to the local phase gradient (Southwell, 1980).

To convert this spot pattern into a modal description of the wavefront $\vec{z}$ we first compute the displacement vector for each spot, $\vec{v}$ (Fig. 3.4 bottom-center panel). We then compute the theoretical vSHWFS response matrix H for a number of Zernike modes (von Zernike, 1934; Noll, 1976) and solve $\vec{v} = H\vec{z}$ in the least-squares sense using the average intensity of each sub-image $\vec{w}$ as weights to obtain $\vec{z}$. This is done for each iteration as both $\vec{v}$ and $\vec{w}$ change. Hence sub-apertures that receive little light are weighted less, preventing those from adding noise.

Scattering and speckles

As biological tissue is generally highly scattering, imaging the back-reflected coherent light onto a detector will produce strong speckle patterns (Goodman, 1976). These speckles arise from coherent superposition of light scattered by the (biological) tissue, which forms an interference pattern observed as speckles. This effect was found to be problematic in phase-step interferometry applications (Feierabend et al., 2004; Rückel et al., 2006) where multiple images have to be combined to recover the complex wave, and the authors solved this by averaging several measurements with different speckle patterns. Wang et al. (2012) investigated the

behaviour of speckles and quantified the influence of a confocal aperture on the speckle sizes, among other things. They found that a smaller aperture will yield larger speckles, which can be explained by the spatial filtering of the aperture.

In our case, the confocal pinhole used for axial sectioning also provides low-pass filtering, smoothing out speckles smaller than the diffraction-limited resolution of the pinhole. This effect is illustrated in Fig. 3.5, where a wavefront sensor image is shown using the scattering beads solution with and without the 400 µm pinhole (left and center panels). When using a confocal pinhole a trade-off needs to be made between speckle reduction (Wang et al., 2012) and wavefront measurement resolution (Rahman & Booth, 2013); a larger pinhole will allow wavefront aberration measurements with higher resolution while a smaller pinhole will reject scattered light better.

In addition to the spatial filtering by the confocal pinhole, the Fourier bandpass filter separates the fringes from the speckles as long as their spatial frequencies do not overlap. The spatial resolution of a plane wave incident on a circular aperture is given by

$$r = 1.22\lambda/2n \sin\theta, \tag{3.3}$$

where θ is the half-opening angle of the beam. In this case, the 400 µm pinhole and the f=100 mm lens (L7) yield a resolution of 230 µm at the CCD, such that any feature smaller than this (i.e. speckles) will be smoothed out. This is in agreement with what was found by Wang et al. (2012), where their smallest field of view corresponds to a ∼3 mm confocal aperture, which has a resolution of ∼30 µm, or ∼3 pixels on their CCD 1. As expected, this ∼3 pixel resolution is a lower limit in their Figure 5 (top row).

Because the reference beam is combined with the sample beam after the confocal pinhole, the spatial frequency of the fringe pattern can be higher than the maximum spatial frequency of the transmitted wave by the confocal pinhole. Since the fringe frequency as well as the pinhole diameter are free parameters, these can be selected such that the speckles and fringes are well-separated in the Fourier domain.

Something the confocal pinhole does not prevent, however, is the amplitude modulation induced by the speckles. As can be seen in Fig. 3.5, there are both saturated and dark spots in the image where the fringe contrast is reduced significantly. In these locations the phase cannot be reconstructed properly, and to mitigate this we measure the wavefront at different locations in the specimen with different speckle patterns, similar to Feierabend et al. (2004); Rückel et al. (2006); Wang et al. (2012). However,

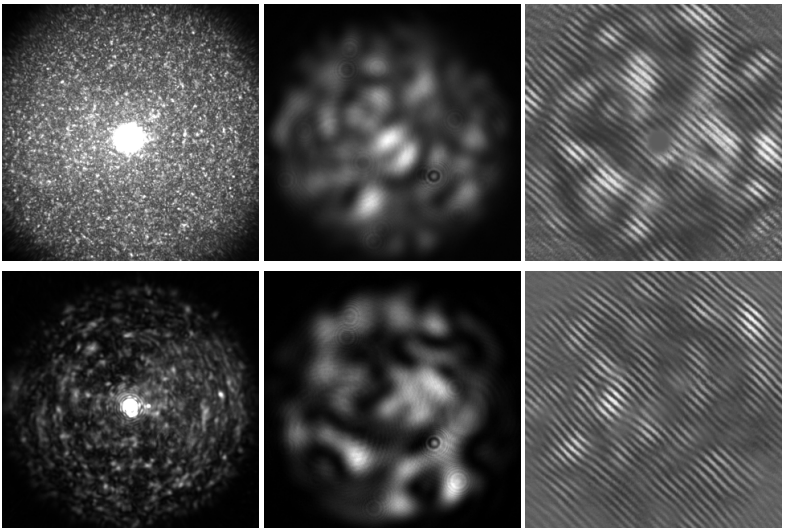


Figure 3.5: Fringe and speckle patterns recorded from back-scattered light off the scattering beads solution (top) and off the spheroid sample (bottom). These images were observed by the WFS camera at the back aperture plane. The left images was obtained without a confocal pinhole, while the center and right images were taken with one in place. Additionally, the right images were flatfield corrected and interfered with the reference beam to produce fringes. The scattering bead solution shows a strong and homogeneous speckle pattern, while the spheroid shows larger speckles, indicating it is less scattering. The speckle pattern is spatially filtered by the confocal pinhole such that the fringe pattern contrast is enhanced, although larger-scale amplitude variations were still present. As can be seen in the center and right columns, the spatial frequency of the filtered speckles is similar for both samples. The fringe pattern has a higher spatial frequency than imposed by the confocal pinhole because the reference beam is introduced after the pinhole. All images were taken with the 20× objective and all fields of view are the same.

compared to those methods we obtain the wavefront from one frame, after which the wavefront measurements are averaged, as opposed to the 4-frame requirement of phase-stepping interferometry.

Double pass

In direct wavefront sensing, the beam passes through the sample twice, once towards the focus and once back-reflected. Because of this, the sensitivity for different wavefront aberration modes can be different. Rahman & Booth (2013) show that there is a difference between even and odd modes

in the case of a pure specular reflection, something which has long been known in retinal imaging (Artal et al., 1995b,a). In the other extreme when a purely scattering sample is used instead of the mirror, there is little to no correlation between the incoming angle and back-reflected angle, such that the scatterers in the specimen all act as point sources and their back-reflections are coherently super-imposed (Rückel & Denk, 2007). In real samples, the sensitivity to modes can differ significantly, but unlike a mirror reflection, odd modes can have non-zero sensitivities and thus can be detected (Rahman & Booth, 2013).

Wavefront correction

To verify the wavefront-sensing method, we control the deformable mirror in closed loop, minimising the residual aberrations measured by the wavefront sensor.

Before correction, we characterise the response of the deformable mirror using mirrors M3 and M4 to direct the laser beam directly onto the wavefront sensor after reflecting off the deformable mirror only once. To interfere the reference beam with the shorter calibration beam, we use the alignment laser, which is a continuous wave and as such also interferes with a beam of unequal optical path length. Because this gives a well-behaved fringe pattern, we can directly obtain the phase from the complex wave using Eq. 6 from (Takeda et al., 1982). This phase is unwrapped using Itoh's weighted phase unwrapping scheme (Peck, 2010) with the amplitude of $c(x, y)$ as weight.

The result of the above calibration produces an influence matrix of the deformable mirror, mapping the input voltages to mirror shape. We use the pseudo-inverse of this matrix to control the mirror, where we typically use all but one singular value in the inversion. Using a basis set of Zernike modes we can construct a matrix that maps a vector of Zernike amplitudes to mirror voltages, which we use for our control. Because of the finite number of actuators, we limit the control of the DM to Zernike modes 5 to 15 (ordering from Noll (1976)), excluding mode 12, which the DM cannot produce due to the actuator configuration. When driving only even modes, we use Zernike modes 5, 6, 11, 13, 14 and 15.

When running the system there are occasionally frames with blurred out fringes, probably due to lab seeing or vibrations, for which the wavefront is poorly defined. To mitigate this, we reject frames with a low fringe contrast, as defined by the power in the extracted side band. Images with less than 75 per cent of the average power of the last 100 frames are discarded. For accepted frames, we take sets of 4 images, for each of which we com-

pute the vSHWFS spot pattern. We take the median image of this set, from which we compute the spot displacement vector. We compute the modal phase representation by multiplying the displacement vector with the response matrix, weighing it with the average intensity in each sub-image, as described above.

Although the speckle pattern induced by the scattering samples is smoothed by the spatial filtering of the confocal pinhole, it nonetheless induces strong amplitude variations. As described above, we mitigate this by averaging speckle patterns recorded at different positions. This is achieved by applying the phase corrections to the deformable mirror with a proportional gain of 0.05 while scanning the sample over a variable-sized area.

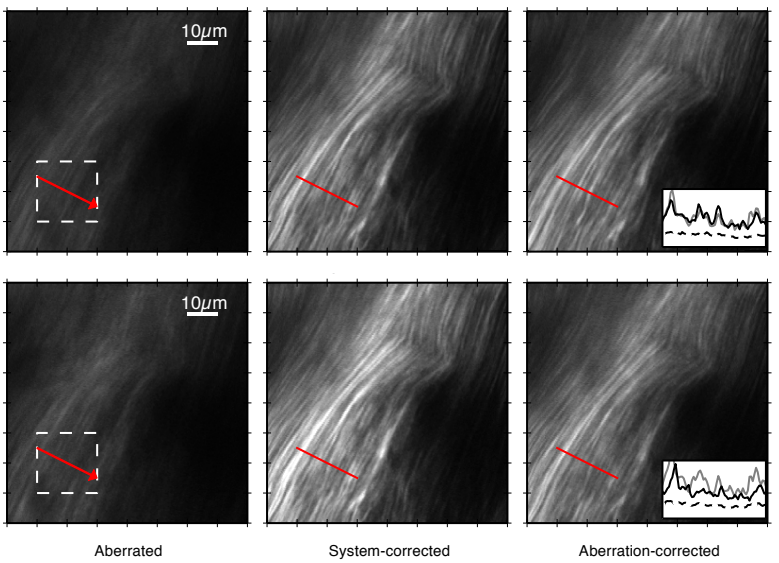


Figure 3.6: Correction of DM-induced aberrations Z_5, Z_6, Z_{11} (top) and Z_5–Z_8 (bottom) with the 20× objective in rat tail collagen fibres using the 20× objective. The panels show an image with the DM in aberrated shape (left), the system-corrected image (centre) and an image after aberration correction (right). For both cases, modes Z_5–Z_{15} were controlled with the DM when correcting, and the sample was scanned over the area indicated by the dashed box. After correction we obtain images similar to the system-corrected DM case, although the even-only aberration correction (top) performed better in this case. The grayscale is identical in each row. The cross-section profile is taken along the arrow: dashed for aberrated, gray for system-corrected, and black for aberration-corrected data.

3.3 Results

To verify our wavefront-sensing method, we test it in different biologically relevant scenarios using both a 20×/0.75 NA and a 40×/0.90 NA objective. We perform DM-induced aberration correction and correction of sample aberrations, which are compared with a flat DM with respect to the fluorescent optimum as described in Sect. 3.2. When running the correction, scanning over an area to average out the amplitude modulation effect of speckles also allows us to record the fluorescence intensity during the correction. We stop the correction when the average intensity of the scanned area stops to increase, which is typically after approximately 100 iterations, or roughly one minute. After correction, a 256 by 256 pixel image is recorded at 0.5 ms pix^{-1}, which takes 30 s.

Collagen tendons

Collagen fibres were extracted from a rat tail and washed 4 times in distilled water. Following fixation in 4 per cent paraformaldehyde, the collagen fibres were washed in PBS and then embedded in 3 per cent agarose

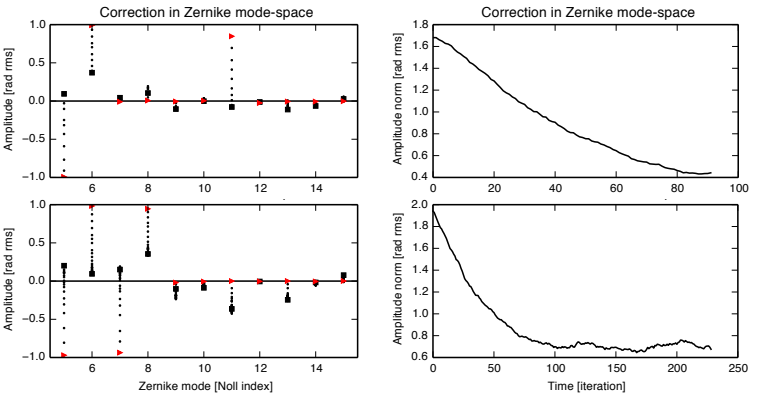


Figure 3.7: Correction loop diagnostics associated with Fig. 3.6: the top and bottom rows correspond to the respective rows in Fig. 3.6. *Left:* modal decomposition of deformable mirror shape during the correction experiment. The triangles note the initial DM shape and the squares indicate the final shape with respect to a system-corrected DM, the dots in between denote intermediate corrections. *Right:* norm of the DM shape during the experiment. The decay in the norm of the DM shape indicates that the correction is successful in recovering the system-corrected setting. See text for details.

(Sigma-Aldrich chemie GmbH) in a 35 mm glass bottom dish (MatTek Corporation).

We image the fibres using 900 nm excitation light to produce the second harmonic signal approximately 50 µm deep inside the rat tail fibre, using the underfilled Nikon 20×/0.75 NA objective. The DM is set to an aberrated shape using mostly even modes (Z_5, Z_6, Z_{11}, Figs. 3.6–3.7 top) as well as both even and odd modes (Z_5, Z_6, Z_7, Z_8, Figs. 3.6–3.7 bottom) and run the control loop as described above.

During correction of the aberrations, we scan the sample over a sub-region of 20 by 20 µm, indicated by a dashed box in the figures. We run the loop until the fluorescence intensity in the scanned area stabilises. Using this corrected DM shape, we image the original region again for comparison, shown on the right in Fig. 3.6, while the centre column shows images obtained with only correction of system aberrations for comparison.

Correction of only even aberrations yields better results than correction of even-and-odd aberrations, as can be seen from the increased brightness and contrast in Fig. 3.6. This is possibly due to the double-pass effect, which makes the correction of odd modes more difficult, although as can be seen in Fig. 3.7, these modes can, to a large extent, be corrected. In both cases, we are able to correct the induced aberrations as shown Fig. 3.7 and recover the image quality as obtained after only correcting system aberrations. While the amplitude of the induced aberrations should decrease, it does not necessarily have to go to zero in the case that there are aberrations in addition to the system-correction. In any case the corrected shape should be the same for both correction runs, as they were obtained at the same location. As can be seen in the left panels of Fig. 3.7, the correction generally obtains the correct value for each mode, although some modes show a discrepancy (such as Z_{11}).

MCF-7 Spheroids

We also measure and correct wavefront aberrations inside an MCF-7 breast cancer spheroid (Soule et al., 1973). The spheroids were generated from cultured MCF-7 using liquid-overlay technique as described in Friedrich et al. (2009). After aggregating for 7 days, spheroids of 0.5 to 1.2 mm diameter were harvested for measurement. These spheroids were fixed in 4 per cent formaldehyde and permeablised in 0.2 per cent Triton x-100 (Sigma-Aldrich). The samples were washed 3 times in PBS. Following blocking with 10 per cent FBS, the spheroids were incubated with Hoechst 33342 (Molecular Probes) and finally washed with PBS. The samples were then embedded in 3 per cent agarose in concave microscope

slides (Marienfeld-superior). We excite the Hoechst using 750 nm excitation light, and during the experiments we note no significant bleaching (<5 per cent).

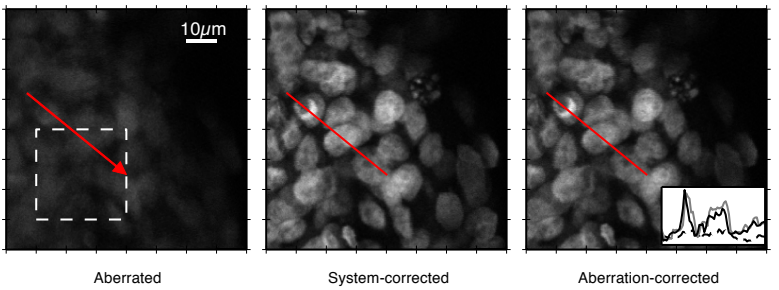


Figure 3.8: DM-induced aberrations Z_5–Z_8 correction in a Hoechst-stained MCF7-spheroid sample using the 20× objective. The averaging area for the correction is indicated by the dashed box. Note that there is no observable bleaching in this area. The grayscale is identical, the cross-section profile is taken along the arrow.

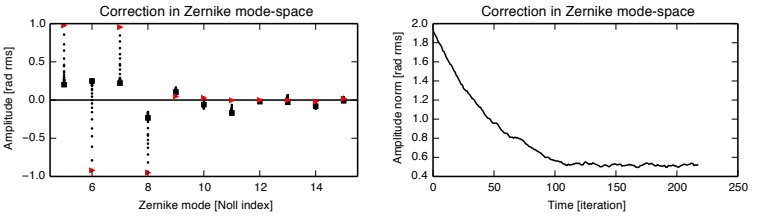


Figure 3.9: Correction loop diagnostics associated with Fig. 3.8, plots as in Fig. 3.7. *Left:* modal decomposition of the DM shape during the experiment. *Right:* norm of the DM shape during the experiment.

As with the collagen fibres, we introduce even and odd aberrations with Zernike modes 5 through 8 (Figs. 3.8–3.9). We image a region of 80 µm by 80 µm at 35 µm depth inside the spheroid using the underfilled 20×/0.75 NA objective with the DM set to the aberrated shape. We run the loop until the fluorescence intensity in the small scanned area stabilises, after which we image the larger region. These results, comparing the aberrated, flat and corrected DM shapes, are shown in Fig. 3.8.

As with the collagen fibres, also here we are able to correct the induced aberrations and significantly improve the image quality as given by a system-aberration corrected setup. Odd aberrations are also corrected

here, indicating that we indeed also obtain odd aberration information from the wavefront sensing (Fig. 3.9).

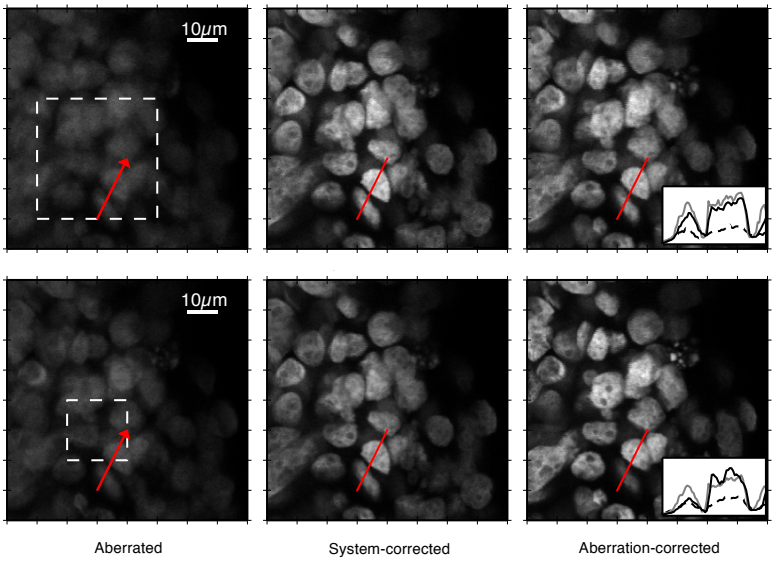


Figure 3.10: Correction of DM-induced aberrations Z_5, Z_6, Z_{11} (top) and Z_5–Z_8 (bottom) in a Hoechst-stained MCF7-spheroid sample using the 40× objective. The corrected image (right) is shown with an aberrated (left) and a system-corrected reference (center) image. During correcting, the sample was scanned over the region indicated by the dashed box. The grayscale is identical in each row. The cross-section profile is taken along the arrow: dashed for aberrated, gray for system-corrected, and black for aberration-corrected data.

Subsequently, we imaged the same sample with a 40×/0.90 NA air objective. We image a region 20 µm deep inside the spheroid, and use the correction collar to correct for depth-induced spherical aberration, setting the collar to correct for 17 µm glass ≈ 20 µm water (assuming $n_{\text{sample}} \approx 1.3$). We add even-and-odd and even aberrations on the DM with the same modes and amplitudes as before, and correct these.

The results are shown in Fig. 3.10. We observe a slightly different image location, which is possibly due to correction of spherical aberration, which reduces the depth of focus, or accidental introduction of (de-)focus due to inaccuracies in the calibration. In both cases we are able to correct the DM-induced aberrations. In the second case when correcting even-and-odd aberrations (bottom panels), the control loop improves the image

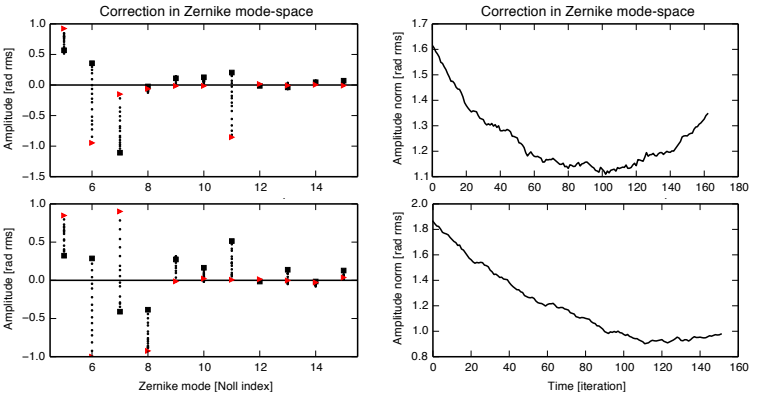


Figure 3.11: Correction loop diagnostics associated with Fig. 3.10, plots as in Fig. 3.7. *Left:* modal decomposition of the DM shape during the experiment. *Right:* norm of the DM shape during the experiment. See text for details.

compared to only system-aberration correction, showing increased brightness but also contrast. As can be seen in Fig. 3.11, the correction does not converge to zero amplitude for the different modes, however the two experiments do both converge to similar values (e.g. Z_5, Z_6 and Z_7), indicating that we are correcting sample-induced aberrations as well.

Finally, we image a deeper region at 35 µm without further adjustment of the correction-collar ring on the objective. This yields depth-dependent sample aberrations due to the refractive index mismatch of air and the spheroid. In this case we had to stop the correction prematurely due to saturation of the DM actuators. These results are shown in Fig. 3.12, showing both higher intensity and contrast as compared to only system correction, and smaller features in the cell nuclei become visible. As can be seen in Fig. 3.13, there is significant spherical aberration (Z_{11}) as expected because of the uncorrected refractive index mismatch, but we additionally find strong coma (Z_7). Since these images were obtained at the edge of the spheroid, the spherical surface might induce such coma aberration. Additionally, we find <0.5 rad amplitudes for modes Z_5, Z_6, Z_8 and Z_9. Correcting depth-dependent sample aberrations is important as there is never a perfect match between immersion fluid and sample, which results in these types of aberrations (Débarre et al., 2009a). We find that our method is capable of correcting such aberrations, but during the experiments we were constrained by the limited stroke of the deformable mirror. A mirror with large stroke, especially for spherical aberration, is therefore desirable.

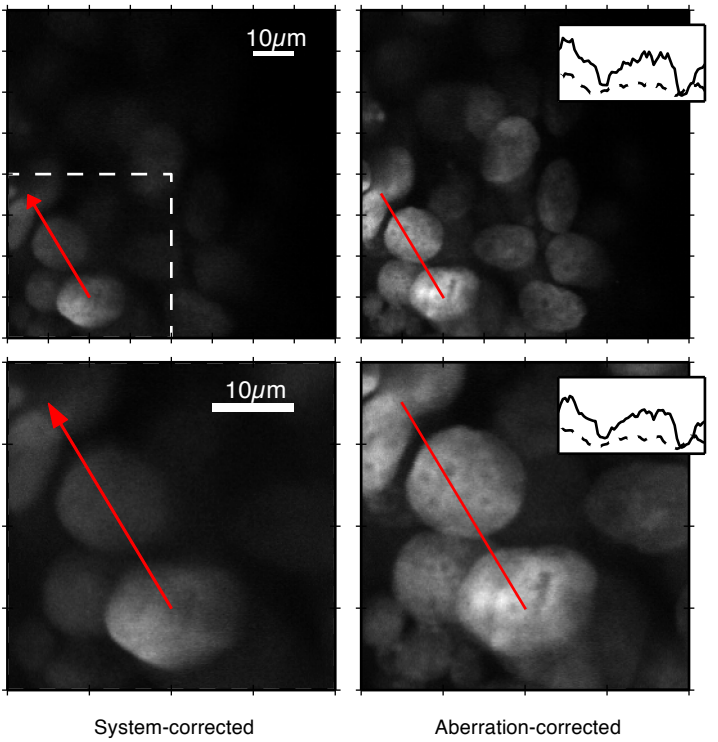


Figure 3.12: Correction of sample-induced depth-dependent aberrations at 35 µm inside the spheroid imaged with the 40× objective. The bottom panels show a zoomed-in region of the top images, and were taken with the same corrected DM settings. Left is the flattened DM setting, right shows the image after correction. The dashed box indicates both the region shown in the bottom panels, as well as the area used to scan over during correction for averaging. The grayscale is identical per row. The cross-section profile is taken along the arrow: dashed indicates the system-corrected data, black the aberration-corrected data.

3.4 Discussion

We developed a direct wavefront-sensing method that is capable of deriving the phase from a single image instead of four. By using a confocal pinhole, we strongly reduce ghost reflections and smooth out speckles in our camera. The Fourier analysis of the fringe pattern additionally allows us to filter out low and high spatial frequency noise.

Since we retrieve the complete wavefront information from a single mea-

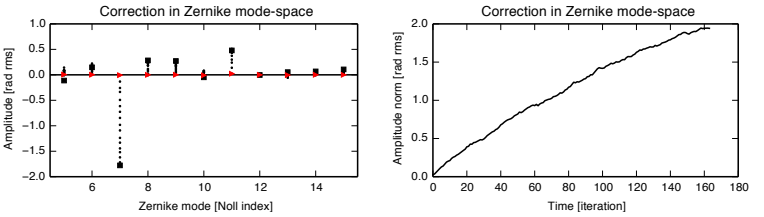


Figure 3.13: Correction loop diagnostics associated with Fig. 3.12, plots as in Fig. 3.7. *Left:* modal decomposition of the DM shape during the experiment. *Right:* norm of the DM shape during the experiment.

surement, as compared with time-domain methods (Feierabend et al., 2004; Rückel et al., 2006; Fiolka et al., 2012; Tuohy & Podoleanu, 2010; Wang et al., 2012), this method is faster and does not require moving parts to generate the interferogram. Because of this, vibrations and internal seeing are less problematic. Although in strongly scattering samples averaging is required to overcome speckle-induced amplitude variations, this is similar to other methods. In addition, most low spatial frequency noise can easily be separated in the Fourier domain, as long as there is sufficient fringe pattern contrast. Finally, this method accommodates varying back aperture sizes and aberration characteristics without changing optical elements, and the trade-off between dynamic range and sensitivity can be adjusted by simply changing the reference-wavefront tilt.

Like Rückel et al. (2006), we find that the phase extracted from the fringes is ill-defined and contains singularities and discontinuities. This is likely due to the back-scattering occurring in a volume, such that the measured image is a coherent super-position of fringe patterns from different scatterers in the focal volume. We observe this fact in, for example, the scattering-bead solution, where the fringes bifurcate, indicative of a phase singularity. We find that using a virtual Shack-Hartmann wavefront sensor overcomes these problems.

The double-pass through the sample changes the sensitivity for different wavefront modes, where even modes are more easily sensed than odd modes (Rahman & Booth, 2013). While odd modes are measured by direct wavefront sensing, these measurements have a lower SNR. One way to solve this is to combine a direct wavefront-sensing method with intensity optimisation. Since the methods measure the wavefront differently, the reconstruction becomes better defined, alleviating the even-odd asymmetry problem. Combining these two feedbacks for the correction is therefore a logical next step.

Since speed is important in scanning microscopy, so is the speed of the adaptive optics method to be used. In this context, we estimate that the method is, in principle, fast enough to keep up with scanning microscopes (Rückel et al., 2006). The wavefront sensor camera exposure time used for these experiments is 4 to 8 ms. Using a more sensitive and faster camera should be sufficient to run the control loop at up to 1 kHz. Although this does not include computational load, the core of our wavefront-sensing analysis comprises Fourier transforms, which for 10^5 pixels can be done at 1 kHz. Currently we need ~100 iterations per correction, in part due to the low 0.05 gain, which was chosen for robustness, and we are confident that this can be decreased if the control is improved. One improvement is using quadratic programming as done by Rückel et al. (2006), which additionally prevents saturation of the DM.

Although per-point correction is not feasible in this way, considering that the wavefront changes gradually over the sample (Zeng et al., 2012), we can measure and correct the wavefront on a coarser grid than what we are scanning, such that 6^2 to 8^2 wavefront measurements are needed per sample. We conclude that our method is, in principle, fast enough to be applied to real-time systems.

3.5 Conclusions

We present an aberration correction method that directly measures the wavefront of the back-reflected light in multi-photon microscopes. By measuring the interference pattern between a reference beam and a sample beam, we are able to retrieve the wavefront and correct aberrations. Using coherence gating and a confocal pinhole, only light from a shallow layer of interest of the order of 10 µm is selected. Furthermore, the pinhole smoothes out speckles produced by the scattering samples. Coherence-gated fringe analysis is insensitive to temporal variations slower than the exposure time (internal seeing, vibrations) as it obtains the wavefront information from a single measurement as opposed to four measurements in phase-step interferometry.

We use the back-reflected laser light for wavefront sensing, by using both a confocal pinhole for smoothing speckles and to reduce the background signal from e.g. the cover slip, as well as coherence gating to extract wavefront information. Although coherence gating could, in principle, be used exclusively for depth selection, this adds a strong background signal, degrading the fringe contrast and making it difficult to extract the wavefront. Furthermore, using a confocal pinhole is advantageous in the presence of speckles, which are generally present in scattering specimen. While we

have not investigated this in these experiments, the size of the confocal pinhole is a free parameter and allows for a trade-off between speckle rejection (Wang et al., 2012) versus wavefront sensing resolution (Rahman & Booth, 2013). Since the reference beam is combined with the sample beam after the confocal pinhole, this does not constrain the fringe frequency, which can also be freely chosen. These degrees of freedom could allow for deeper measurements of the wavefront, and requires further investigation.

The maximum depth at which direct wavefront sensing in scattering samples works was investigated by Wang et al. (2012). They find that at a certain depth the wavefront cannot be recovered anymore, which they attribute to multiply scattered photons from random locations in the sample accidentally falling in the coherence gate. Because of this, there is not a single unique wavefront that can describe the phase of those photons anymore. Since a pinhole rejects scattered light, we suspect that using this allows measuring deeper in the sample. The depth gains that a pinhole brings need to be investigated in more detail.

We show the feasibility of the method in rat tail collagen fibres as well as in a spheroid of Hoechst-stained MCF-7 cells, indicating that it also works in the presence of scattering-induced speckles. We correct both only even and even-and-odd DM-induced aberrations, and are able to improve the image compared to only system-aberration correction. Additionally we are able to correct depth-dependent sample aberrations present in any microscopy system due to refractive index mismatches between sample and immersion fluid.

Acknowledgments

The authors are grateful for insightful discussions with A. V. Agronskaia and G. A. Blab, and thank D. J. van den Heuvel for the scattering beads sample preparation. The authors also thank G. Vdovin at Okotech for supporting this project. This research is supported by the Technology Foundation STW under project number 10433, which is part of the Smart Optics Systems perspective program.

4 Optimisation-based wavefront sensorless adaptive optics for non-linear microscopy

Based on Jacopo Antonello, Tim van Werkhoven, Michel Verhaegen, Hoa Truong, Christoph Keller and Hans Gerritsen *J. Opt. Soc. Am. A* (to be submitted)

Optical aberrations have detrimental effects in multiphoton microscopy. These effects can be curtailed by implementing model-based wavefront sensorless adaptive optics, which only requires the addition of a wavefront shaping device such as a deformable mirror to an existing microscope. The aberration correction is achieved by maximising a suitable image quality metric. We implemented a model-based aberration correction algorithm in a second-harmonic microscope. We remove the tip, tilt and defocus aberrations from the basis functions used for the control of the deformable mirror, as these aberrations induce distortions in the acquired images. We compute the parameters of a quadratic polynomial that is used to model the image quality metric directly from experimental input-output measurements. Finally we apply the aberration correction by maximising the image quality metric using the least-squares estimate of the unknown aberration.

4.1 Introduction

Multiphoton microscopy techniques such as two-photon fluorescence microscopy Denk et al. (1990) and second-harmonic microscopy Campagnola et al. (2001) are commonly employed to image biological specimens. Exploiting the image sectioning properties of these processes, one can create high-resolution 3D reconstructions that are invaluable for biomedical research. One limiting factor is the presence of specimen-induced aberrations. Because the index of refraction is not homogeneous within the specimen, aberrations affect both the resolution and the maximum depth of penetration Booth (2007b). Using adaptive optics Tyson (2010), these detrimental effects can be minimised by reducing the phase aberrations. A phase aberration can be introduced in the excitation beam by means of a deformable mirror, for example. Chosen correctly, such a phase aberration can suppress some amount of the specimen-induced aberrations.

Direct measurement of the specimen-induced aberrations is challenging. One solution involves measuring the aberrations of the excitation light that is back-scattered from the specimen Rahman & Booth (2013); van Werkhoven et al. (forthcoming, 2012); Rückel et al. (2006); Cha et al. (2010); Bueno et al. (2010). In this case, the difficulty arises in excluding the light that is reflected from the out-of-focus layers of the specimen Rahman & Booth (2013); van Werkhoven et al. (forthcoming). Additionally, these measurements are weakly sensitive to odd aberrations Rahman & Booth (2013), due to the double pass effect Artal et al. (1995b). In another solution, instead, the emission from a point source inside the specimen is used to perform Shack-Hartmann wavefront sensing Azucena et al. (2010); Aviles-Espinosa et al. (2011); Shaw et al. (2013); Tao et al. (2013b,a). Here the difficulty stems from the lack of such reference point sources within the specimen and from the limited number of photons available in the emission signal.

An alternative, indirect approach involves deducing the specimen-induced aberrations solely by examining the emission signal. This approach only requires the addition of a deformable mirror into the excitation path of an existing microscope. A solution that is based on the segmentation of the pupil has been proposed Ji et al. (2009); Milkie et al. (2011). Other solutions are based on the optimisation of an image quality metric, which attains its global maximum when the residual aberration is maximally suppressed. In practice, different trial aberrations are sequentially applied with the deformable mirror until the image quality metric attains its maximum.

General optimisation algorithms can be used to maximise the image qual-

ity metric Vdovin (1998); Albert et al. (2000); Vorontsov (2002); Marsh et al. (2003); Murray et al. (2005); Wright et al. (2005). However, because these algorithms have no prior knowledge about the metric, a large number of trial aberrations must be evaluated before the metric is maximised Booth (2007b, 2006); Antonello et al. (2012). Reducing the number of trial aberrations is critical in achieving short image acquisition times and in limiting side effects such as photobleaching and phototoxicity. For small aberrations, the response of the image quality metric can be approximated using a simple model, such as a quadratic polynomial Neil et al. (2000); Booth et al. (2002); Booth (2006); Débarre et al. (2007b); Booth (2007c); Débarre et al. (2008b, 2009b); Jesacher et al. (2009); Facomprez et al. (2012); Zeng et al. (2012); Antonello et al. (2012). Model-based aberration correction algorithms Neil et al. (2000); Booth et al. (2002); Booth (2006); Débarre et al. (2007b); Booth (2007c); Débarre et al. (2008b, 2009b); Jesacher et al. (2009); Song et al. (2010); Linhai & Rao (2011); Yoo et al. (2012); Facomprez et al. (2012); Zeng et al. (2012); Antonello et al. (2012) exploit the knowledge about the model of the metric to quicken the aberration correction, thus curtailing the aforementioned side effects.

In this paper, we investigate applying a model-based wavefront sensorless aberration correction algorithm to a second-harmonic microscope.

In scanning microscopy, the aberration correction should not introduce x-tilt, y-tilt and defocus Zernike aberrations Noll (1976). These aberrations do not affect the image quality. Instead, they affect the position of the focal point within the specimen and they induce translations or distortions in the acquired images Débarre et al. (2008b, 2009b); Thayil & Booth (2011). We control the deformable mirror using linear combinations of the influence functions Paterson et al. (2000); Vdovin et al. (2013) instead of using the Zernike modes. By applying the Singular Value Decomposition (SVD), we ensure that minimal amounts of x-tilt, y-tilt and defocus are induced by the deformable mirror.

We also consider the aberration correction problem, where we compute the least-squares estimate of the unknown aberration. This involves solving a possibly non-convex optimisation that may exhibit local optima. By exploiting results already applied in the solution of localisation problems Beck et al. (2008), we can efficiently compute the least-squares estimate. Once the estimate is computed, the aberration correction can be applied by maximising the image quality metric.

We are able to compute the parameters of the quadratic polynomial used to model the image quality metric directly from input-output measurements, using a previously proposed method Antonello et al. (2012). Here,

we validate this method using a microscope and a biological specimen for the first time. We assess the accuracy of this method with a number of validation and cross-validation experiments. Finally we validate the aberration correction algorithm proposed in this paper.

The paper is organised as follows. In section 4.2 we discuss the definition of the basis functions for the control of the deformable mirror. In section 4.3 we outline our proposed aberration correction algorithm. In section 4.4 we report the experimental results. The conclusions are drawn in section 4.5.

4.2 Definition of the basis functions for the control of the deformable mirror

The basis functions should satisfy two different requirements. The first requirement is that the basis functions be orthogonal to the x-tilt, y-tilt and defocus aberrations. As mentioned in the introduction, these aberrations induce unwanted displacements or distortions to the region of the specimen that is acquired by scanning the focal point Débarre et al. (2008b, 2009b); Thayil & Booth (2011). When using Zernike modes as the basis functions, the unwanted aberrations can be set to zero.

A second requirement is that the basis functions express the capabilities of the deformable mirror in an accurate and concise form, by taking into account the mechanical limitations of the mirror and the misalignment in the optical system as much as possible. This requirement is not satisfied when using Zernike modes as the basis functions since a deformable mirror with N_a actuations cannot accurately induce a set of N_a Zernike modes.

We now discuss a simple procedure to define a new set of basis functions that satisfy the two requirements discussed so far. This procedure is based on the decomposition of a matrix H that approximately describes a linear relationship between the control signals of the deformable mirror and a set of Zernike coefficients. For completeness, we first report how H can be computed from input-output measurements.

Computation of matrix H from input-output measurements

Let N_a be the number of actuators of the deformable mirror. Assuming that the deformable mirror is a linear device, the phase aberration $\Phi(\boldsymbol{\xi})$ is given by the superposition of the influence functions $\psi_i(\boldsymbol{\xi})$ of each actuator, where $\boldsymbol{\xi}$ is the spatial coordinate in the pupil and u_i is the control

signal of the i-th actuator

$$\Phi(\boldsymbol{\xi}) = \sum_{i=1}^{N_a} u_i \psi_i(\boldsymbol{\xi}). \tag{4.1}$$

For a suitable number N_z of Zernike modes $\mathcal{Z}_j(\boldsymbol{\xi})$, we have that

$$\Phi(\boldsymbol{\xi}) \approx \sum_{j=2}^{1+N_z} z_j \mathcal{Z}_j(\boldsymbol{\xi}), \tag{4.2}$$

where z_j is the j-th Zernike coefficient. We neglect the piston mode $\mathcal{Z}_1$ since this does not affect the image and assume that $\Phi(\boldsymbol{\xi})$ and $\psi_i(\boldsymbol{\xi})$ have zero mean value over the pupil.

We collect the coefficients u_i and z_j respectively into vectors $\mathbf{u} \in \mathbb{R}^{N_a}$ and $\mathbf{z} \in \mathbb{R}^{N_z}$. By considering a grid defined in the pupil, N_c samples of $\Phi(\boldsymbol{\xi})$ are collected into a vector $\boldsymbol{\phi} \in \mathbb{R}^{N_c}$. Similarly, we evaluate $\psi_i(\boldsymbol{\xi})$ and $\mathcal{Z}_j(\boldsymbol{\xi})$ over the grid and define two matrices $\Psi \in \mathbb{R}^{N_c \times N_a}$ and $Z \in \mathbb{R}^{N_c \times N_z}$. From (4.1) and (4.2) we have that $\boldsymbol{\phi} = \Psi \mathbf{u}$ and $\boldsymbol{\phi} \approx Z\mathbf{z}$.

We would like to recover a matrix H which maps an actuation vector $\mathbf{u}$ into the corresponding vector of Zernike coefficients $\mathbf{z}$, i.e., $\mathbf{z} \approx H\mathbf{u}$. H can be computed using input-output measurements, so that the misalignment in the optical system is accounted for. Using a Shack-Hartmann wavefront sensor or interferometric methods Wang & Booth (2009); Fernandez & Artal (2003); Vdovin et al. (2013); Booth et al. (2005) one can collect a set of measurements of the phase $\boldsymbol{\phi}_1, \dots, \boldsymbol{\phi}_D$ corresponding to different settings of the deformable mirror $\mathbf{u}_1, \dots, \mathbf{u}_D$.

We compute H by minimising the following criterion,

$$\min_H \sum_{i=1}^{D} \|\boldsymbol{\phi}_i - ZH\mathbf{u}_i\|^2. \tag{4.3}$$

Setting the derivative with respect to H to zero, we have the following normal equation

$$Z^T Z H \Big(\sum_{i=1}^{D} \mathbf{u}_i \mathbf{u}_i^T\Big) = Z^T \Big(\sum_{i=1}^{D} \boldsymbol{\phi}_i \mathbf{u}_i^T\Big), \tag{4.4}$$

which can be solved by multiplying from the left and from the right by the inverse matrices of $Z^T Z$ and $\sum_{i=1}^{D} \mathbf{u}_i \mathbf{u}_i^T$. For a properly defined grid, the

inverse of $Z^T Z$ exists, since Z is full column rank due to the orthogonality property of the Zernike modes. Additionally, vectors $\mathbf{u}_i$ can be selected so that $\sum_{i=1}^{D} \mathbf{u}_i \mathbf{u}_i^T$ is full rank.

In our system we have that $N_a = 17$ and $N_c = 75912$. We performed $D = 4N_a$ measurements of the phase $\boldsymbol{\phi}_1, \ldots, \boldsymbol{\phi}_{4N_a}$. In each measurement, a single actuator is poked while the other actuators are at rest. We empirically chose $D = 4N_a$, other choices are possible provided $D \geq N_a$. The choice of N_z is more critical. With a poor choice of N_z, the accuracy requirement discussed at the beginning of section 4.2 may not be fulfilled and the approximation $\mathbf{z} \approx H\mathbf{u}$ may be too rough. We chose $N_z = 35$ by evaluating the error in approximating the phase measurements $\boldsymbol{\phi}_i$ using an increasing number of Zernike modes.

SVD-based removal of the x-tilt, y-tilt and defocus aberrations

From the previous section, we conclude that, in our system, the $N_a = 17$ influence functions approximately span a subspace of the space spanned by the first $N_z = 35$ Zernike modes. Because $\text{rank}(H) < N_z$, there exist non zero vectors $\mathbf{z}$ that do not belong to the range of H and the Zernike modes should not be used as the basis functions for the control of the deformable mirror.

We can split vector $\mathbf{z}$ and matrix H so that $\mathbf{z} \approx H\mathbf{u}$ is partitioned as

$$\begin{bmatrix} \mathbf{z}_l \\ \mathbf{z}_h \end{bmatrix} \approx \begin{bmatrix} H_l \\ H_h \end{bmatrix} \mathbf{u}, \tag{4.5}$$

where the x-tilt, y-tilt and defocus coefficients are collected into $\mathbf{z}_l = [z_2, z_3, z_4]^T$. The SVD of H_l is

$$H_l = U_l \begin{bmatrix} \Sigma_l & \mathbf{0} \end{bmatrix} \begin{bmatrix} V_{l1}^T \\ V_{l2}^T \end{bmatrix}, \tag{4.6}$$

where Σ_l has dimensions 3×3, V_{l1} has dimensions $N_a \times 3$ and V_{l2} has dimensions $N_a \times N_p$ with $N_p = N_a - 3$. The required constraint that $\mathbf{z}_l \approx \mathbf{0}$ is enforced if we choose $\mathbf{u}$ such that $H_l\mathbf{u} = 0$. This is achieved by parametrising $\mathbf{u}$ using the columns of V_{l2}, i.e., letting $\mathbf{u} = V_{l2}\mathbf{p}$, where $\mathbf{p} \in \mathbb{R}^{N_p}$.

In summary, we have that

$$\Phi(\boldsymbol{\xi}) = \sum_{i=1}^{N_a} u_i \psi_i(\boldsymbol{\xi}), \quad \text{s.t.} \quad \mathbf{u} = V_{l2}\mathbf{p}, \tag{4.7}$$

or equivalently

$$\Phi(\boldsymbol{\xi}) = \sum_{i=1}^{N_p} p_i \omega_i(\boldsymbol{\xi}), \tag{4.8}$$

where $\omega_i(\boldsymbol{\xi})$ are the new basis functions. We have that $\omega_i(\boldsymbol{\xi}) = \sum_{j=1}^{N_a} (V_{l2})_{ji} \psi_j(\boldsymbol{\xi})$, where $(V_{l2})_{ji}$ denotes element (j, i) of V_{l2}. For a given vector $\mathbf{p}$, we can compute the control signals of the actuators with $\mathbf{u} = V_{l2}\mathbf{p}$. Similarly, for a given $\mathbf{p}$, the Zernike analysis of the induced wavefront aberration is given by $\mathbf{z} \approx HV_{l2}\mathbf{p}$.

We also applied regularisation Hansen (2010) by truncating the SVD of HV_{l2} to $U_1\Sigma_1 V_1^T$. Using no more than 80% of the sum of the singular values, Σ_1 is a 7×7 matrix. In our experiments the deformable mirror was controlled using a vector $\mathbf{r} \in \mathbb{R}^N$ where $N = 7$. For a given $\mathbf{r}$, the control signals of the actuators of the deformable mirror are computed using $\mathbf{u} = V_{l2}V_1\mathbf{r}$. The Zernike analysis of the induced wavefront is computed using $\mathbf{z}_h \approx H_h V_{l2} V_1 \mathbf{r}$ and the rms of the phase profile is given by computing the 2-norm, i.e., $\|\mathbf{z}_h\|$. This is equivalent to applying another parametrisation to (4.8). We remark that in this way, no pseudo-inverse is ever computed or used to control the deformable mirror, differently from what is done in Paterson et al. (2000), for example.

4.3 Least-squares estimation of the unknown aberration

In this section we discuss the aberration correction algorithm. In Débarre et al. (2009b); Facomprez et al. (2012); Zeng et al. (2012), the authors show that, for small aberrations, the image quality metric can be modelled using a quadratic polynomial. We denote a measurement of the image quality metric at time instant k with $\tilde{y}_k$. We have that

$$\tilde{y}_k = c_0 - (\mathbf{x} + \mathbf{r}_k)^T Q (\mathbf{x} + \mathbf{r}_k) + \epsilon_k, \tag{4.9}$$

where c_0 and Q are the parameters of the quadratic polynomial. Matrix Q is a positive semidefinite matrix, i.e., $Q \succeq 0$ Antonello et al. (2012). Vector $\mathbf{x}$ represents the unknown aberration whereas $\mathbf{r}_k$ accounts for the aberration induced by the deformable mirror. The term ϵ_k is a placeholder that collects both the uncertainty in modelling the image quality metric and the measurement noise, and as such it cannot be measured by definition. By including this term, a measurement $\tilde{y}_k$ can be set equal to the right hand side of (4.9). Excluding ϵ_k, the right hand side of (4.9) is referred to as the approximate image quality metric in Antonello et al. (2012). The parameters c_0 and Q can be computed using the input-output measurements

recorded in a calibration experiment described in section 4.4 or using the methods described in Débarre et al. (2008b, 2012).

Definition of the least-squares problem

The aberration correction is achieved by maximising the image quality metric, i.e., by letting $\mathbf{r}_k = -\mathbf{x}$ in (4.9). For this reason, we must first estimate the unknown vector $\mathbf{x}$. This can be done by applying $m \geq N + 1$ trial aberrations $\mathbf{r}_1, \ldots, \mathbf{r}_m$ with the deformable mirror and by taking the corresponding measurements $\tilde{y}_1, \ldots, \tilde{y}_m$.

Collect $\epsilon_1, \ldots, \epsilon_m$ into a vector $\boldsymbol{\epsilon}$ and $\tilde{y}_1, \ldots, \tilde{y}_m$ into a vector $\tilde{\mathbf{y}}$. By stacking m instances of (4.9), we can define a vector-valued function $\mathbf{g}$

$$\mathbf{g}(\mathbf{x}) = \begin{bmatrix} c_0 - (\mathbf{x} + \mathbf{r}_1)^T Q (\mathbf{x} + \mathbf{r}_1) \\ \vdots \\ c_0 - (\mathbf{x} + \mathbf{r}_m)^T Q (\mathbf{x} + \mathbf{r}_m) \end{bmatrix}, \tag{4.10}$$

such that

$$\tilde{\mathbf{y}} = \mathbf{g}(\mathbf{x}) + \boldsymbol{\epsilon}. \tag{4.11}$$

The least-squares estimate of $\mathbf{x}$ can be obtained by minimising $\|\boldsymbol{\epsilon}\|^2$, i.e., by solving

$$\min_{\mathbf{x}} f(\mathbf{x}), \tag{4.12}$$

where

$$f(\mathbf{x}) = \|\tilde{\mathbf{y}} - \mathbf{g}(\mathbf{x})\|^2. \tag{4.13}$$

We note that if vector $\boldsymbol{\epsilon}$ follows a multivariate normal distribution with zero mean and covariance proportional to the identity matrix, solving (4.12) provides the maximum likelihood Torrieri (1984) estimate of $\mathbf{x}$.

Analysis of the least-squares problem

Finding the global minimum of $f(\mathbf{x})$ appears to be non trivial as $f(\mathbf{x})$ may be non-convex. This is illustrated with a two-dimensional example in Fig. 4.1. Here, the contour plot of (4.13) is shown, when $m > N + 1$ measurements of $\tilde{y}$ are taken. The measurement noise is zero, i.e., $\boldsymbol{\epsilon} = \mathbf{0}$. Nevertheless, $f(\mathbf{x})$ is not convex and exhibits two critical points. In addition to the least-squares solution $\mathbf{x}^{\text{ls}}$ of (4.12), which is the global minimum and for which $f(\mathbf{x}^{\text{ls}}) = 0$, a local minimum $\mathbf{x}^{\text{loc}}$ is present. In case one uses $\mathbf{r}_k = -\mathbf{x}^{\text{loc}}$ to perform the aberration correction, then the residual aberration is not zero and the image quality metric is not maximised.

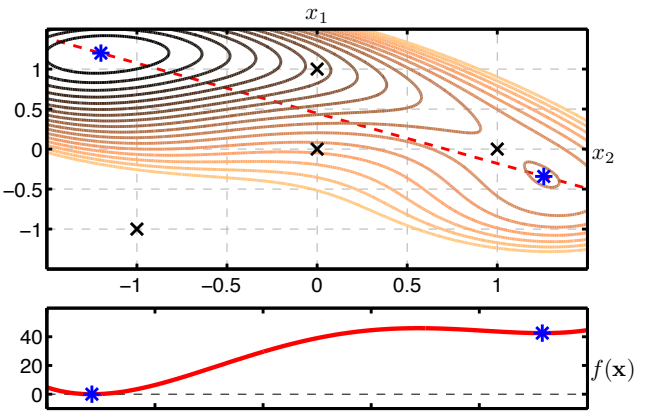


Figure 4.1: (Color online) Contour plot of (4.13). In this example, $f(\mathbf{x})$ is not convex and exhibits a local minimum. The parameters are $c_0 = 100$ and $Q = \begin{bmatrix} 1.25 & 0.433 \\ 0.433 & 1.25 \end{bmatrix}$. Four measurements of $\tilde{y}$, taken at $\mathbf{r}_1 = (0,0)^T$, $\mathbf{r}_2 = (1,0)^T$, $\mathbf{r}_3 = (0,1)^T$ and $\mathbf{r}_4 = (0,1)^T$ are marked with $\times$ symbols. The global minimum $\mathbf{x}^{\text{ls}} = (-1.2, 1.2)^T$ and the local minimum $\mathbf{x}^{\text{loc}} \approx (1.2582, -0.3421)^T$ are indicated with $*$ symbols. Isolines with an elevation greater than 70 have been removed for clarity. A cross section along the dashed line is reported in the plot in the bottom.

Because the convexity property is not satisfied in general, it is unclear how $\mathbf{x}^{\text{ls}}$ can be computed. For example, a gradient based method applied to solve (4.12) may fail to compute $\mathbf{x}^{\text{ls}}$. Alternatively, more sophisticated algorithms may be unsuitable to meet the requirements of a real time implementation. Nevertheless, the global solution of (4.12) can be computed efficiently even when $f(\mathbf{x})$ is not convex, as is outlined in the following section.

Efficient computation of $\mathbf{x}^{\text{ls}}$

In Beck et al. (2008), an efficient algorithm to find the global solution of a possibly non-convex optimisation similar to (4.12) was developed, in the context of localisation problems. In this section we show how the solution proposed in Beck et al. (2008) can be applied to our problem. For the remaining part of the paper we assume that Q is strictly positive definite. This assumption is reasonable since if there are aberrations that do not affect the image quality metric then these cannot be corrected and they should be neglected during the aberration correction Antonello et al.

(2012).

Introducing an additional scalar variable α, we can reformulate (4.12) into the following equivalent constrained optimisation

$$\begin{aligned} \min_{\mathbf{x},\alpha} & \sum_{k=1}^{m} \left(-\alpha + 2\mathbf{r}_k^T Q\mathbf{x} + c_0 - \mathbf{r}_k^T Q\mathbf{r}_k - \tilde{y}_k\right)^2 \\ \text{s.t.} \quad & \alpha = \mathbf{x}^T Q\mathbf{x}. \end{aligned} \tag{4.14}$$

Problem (4.14) is written concisely in matrix form as

$$\min_{\mathbf{w}} \|A\mathbf{w} - \mathbf{b}\|^2 \quad \text{s.t.} \quad \mathbf{w}^T D\mathbf{w} + 2\mathbf{f}^T\mathbf{w} = 0, \tag{4.15}$$

where

$$\begin{gathered} \mathbf{w}^T = \begin{bmatrix}\mathbf{x}^T & \alpha\end{bmatrix}, R = \begin{bmatrix}\mathbf{r}_1^T & \dots & \mathbf{r}_m^T\end{bmatrix}^T, \\ A = \begin{bmatrix}-2R^T Q & -\mathbf{1}\end{bmatrix}, \mathbf{b} = \begin{bmatrix}\mathbf{r}_1^T Q\mathbf{r}_1 + \tilde{y}_1 - c_0 \\ \dots \\ \mathbf{r}_m^T Q\mathbf{r}_m + \tilde{y}_m - c_0\end{bmatrix}, \\ D = \begin{bmatrix}Q & \mathbf{0} \\ \mathbf{0} & 0\end{bmatrix}, \mathbf{f}^T = \begin{bmatrix}\mathbf{0} & -1/2\end{bmatrix}^T \end{gathered} \tag{4.16}$$

and $\mathbf{1}$ and $\mathbf{0}$ denote vectors of appropriate dimensions where all components are respectively ones and zeros. The authors in Beck et al. (2008) note that (4.15) is a Generalised Trust Region Subproblem Moré (1993). Such problems, although non-convex in general, have necessary and sufficient optimality conditions Moré (1993). In particular, from Beck et al. (2008); Moré (1993), we have that $\mathbf{w}$ is a global minimiser of (4.15) if and only if there exist a Lagrange multiplier ν such that

$$\begin{aligned} (A^T A + \nu D)\mathbf{w} &= A^T\mathbf{b} - \nu\mathbf{f} \\ \mathbf{w}^T D\mathbf{w} + 2\mathbf{f}^T\mathbf{w} &= 0 \\ A^T A + \nu D &\succeq \mathbf{0}. \end{aligned} \tag{4.17}$$

We assume that matrix A is full column rank, which in turn implies that $m \geq N + 1$. This assumption on A is by no means restrictive. Because $Q \succ 0$, it can be factored $Q = V\Delta V^T$, where Δ is diagonal and full rank. Choose $R = [V, \mathbf{0}]$, where $\mathbf{0}$ is a vector of zeros, then A is full column rank. We further assume that the optimal Lagrange multiplier ν^* is such that $A^T A + \nu^* D$ is strictly positive definite. The authors in Beck et al. (2008) point out that this more restrictive assumption could be removed with a more refined analysis. However the case where ν^* is such that $A^T A + \nu^* D$

is not strictly positive definite is unlikely to occur both in theory and in practice Beck et al. (2008).

Under these assumptions, one can compute

$$\mathbf{w}(\nu) = (A^T A + \nu D)^{-1}(A^T \mathbf{b} - \nu \mathbf{f}) \tag{4.18}$$

for a fixed value of ν. By replacing $\mathbf{w}$ in the second equation in (4.17) with the right hand side of (4.18), one finds a univariate rational polynomial equation in ν

$$\mathbf{w}(\nu)^T D \mathbf{w}(\nu) + 2\mathbf{f}^T \mathbf{w}(\nu) = 0. \tag{4.19}$$

The optimal Lagrange multiplier ν^* can be found examining the solutions of (4.19). From the assumption $A^T A + \nu D \succ 0$, it can be derived Moré (1993); Beck et al. (2008) that ν must be in the interval $(I_l, +\infty)$, where $I_l = -1/\lambda_{\max}(\Delta_A^{-1/2} V_A^T D V_A \Delta_A^{-1/2})$ and we used the factorisation $A^T A = V_A \Delta_A V_A^T$. In addition, it is known Moré (1993); Beck et al. (2008) that (4.19) is strictly decreasing in ν within the considered interval. Therefore the desired root ν^* of (4.19) can be found efficiently, for example via a bisection algorithm Beck et al. (2008). Once ν^* is found, the estimate of $\mathbf{x}$ is extracted from the first N components of $\mathbf{w}(\nu^*)$.

The aberration correction algorithm is therefore applied in the following manner. First the data collection step takes place, whereby the $m \geq N + 1$ trial aberrations are applied and the corresponding measurements $\tilde{y}_1, \dots, \tilde{y}_m$ are taken. Then, ν^* is computed by finding the root of (4.19) within $(I_l, +\infty)$. The estimate $\mathbf{x}^{\text{ls}}$ of the aberration is found in the first N components of $\mathbf{w}(\nu^*)$. The aberration correction step involves applying the correction with the deformable mirror, letting $\mathbf{r} = -\mathbf{x}^{\text{ls}}$. These steps can be repeated in the following time instants by including more than m measurements to achieve a refined correction. We note that in Antonello et al. (2012), the least-squares estimate of $\mathbf{x}$ was not computed since the quadratic constraint in (4.14) was neglected to obtain a linear least-squares problem.

4.4 Experimental results

We implemented the model-based wavefront sensorless algorithm and report the experimental results in this section. Following Jesacher et al. (2009), we employ the mean image intensity as a metric to correct aberrations in our second-harmonic microscope. Our experiments show successful aberration correction using this metric (see section 4.4).

Our first purpose is to validate a previously proposed method Antonello et al. (2012) to compute the parameters c_0 and Q of (4.9) using input-output

measurements. This validation has not been previously done in a realistic setting, since in Antonello et al. (2012) no microscope and no specimen were involved. Additionally, we intend to validate the aberration correction algorithm described in section 4.3. We report our results in the following sections.

Description of the experimental setup

A schematic of the experimental setup is shown in Fig. 4.2. The source is a Coherent Chameleon Ultra II Ti:Sa 140 fs pulsed, near-infrared laser, with a beam diameter of 1.2 mm. This beam is expanded to a 14 mm wide beam by lenses L1 and L2.

The beam is stopped down to 9.5 mm (AP) before it is reflected under an angle of about 10° by the deformable mirror (DM; Okotech, MMDM 19-channel deformable mirror). The deformable mirror is re-imaged one-to-one onto the objective back-aperture by lenses L3 and L4 (focal lengths, 300 mm). Because the deformable mirror can only introduce negative deflection, we bias the mirror so that we can apply positive and negative deflections to correct the wavefront (see Vdovin et al. (2013)). In addition, the relationship between the control signal u_i of each actuator of the deformable mirror and the voltage applied to the corresponding electrode is quadratic, so that a linear displacement of the membrane is expected Vdovin et al. (2013). Due to this bias, the collimated beam coming from L2 is converging after being reflected by the deformable mirror. We corrected this by using lenses L4 and L3, so that a collimated beam is fed into the objective.

The sample is mounted on an xyz-piezo stage (XYZ; PI, Nanocube P-611.3S). The second-harmonic signal from the specimen is collected by the objective and split off by a 705 nm cut-off dichroic beam splitter (DBS; Semrock, FF705-Di01-25x36). This light is focused onto a 600 µm multimode fibre which is connected to a photo multiplier tube (PMT; Hamamatsu, GaAsP photocathode H7422P-40). The objective used (OBJ) is a 40×/0.9 NA Nikon air objective with spherical correction collar. We manually adjusted the collar to correct for the spherical aberration due to the cover glass and the specimen at the selected depth.

For characterisation of the deformable mirror, we interfere a tilted reference beam with a sample beam deflected off the deformable mirror to create fringes that encode the wavefront deformation. To allow this, a 50/50 beam splitter (BS1; Thorlabs, BS016) splits off part of the light into a reference arm beam, which is relayed onto the camera (CCD; AVT, Guppy Pro F-033b) by mirrors M2, M3, and M4. The sample arm beam is deflected by

the deformable mirror once before flip mirror FM1 directs the light into the calibration arm. Lenses L6 and L7 re-image the deformable mirror onto a camera. For this calibration we use the alignment laser, which is a continuous wave. We used the method described in van Werkhoven et al. (forthcoming) to decode the wavefront from the fringe patterns.

The piezo stage is controlled with a data acquisition board (National Instruments, PCI-e 6259) on a Windows computer running LabView. The deformable mirror is controlled through a PCI DAC card on a Linux computer running MATLAB and custom C code.

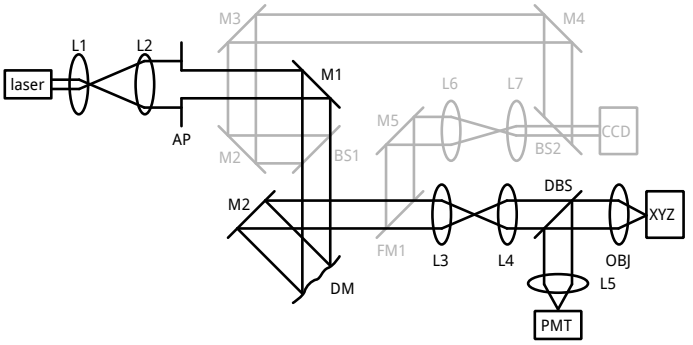


Figure 4.2: Illustration of the optical setup. The components in black are used throughout the aberration correction experiments. The components in grey are used only for the initial characterisation of the deformable mirror (DM). A pulsed laser beam is expanded with lenses L1 and L2, clipped by aperture AP and reflected by flat mirror M1 onto the DM. The DM is in an image of the back aperture of the microscope objective (OBJ), using lenses L3 and L4. The DM is illuminated under an angle of about 10° using the flat mirrors M1 and M2. The microscope objective (OBJ) focuses the light onto the specimen, which is supported by an xyz stage (XYZ). The second-harmonic signal emitted from the focal point inside the specimen is collected with the objective and separated from the illumination beam using a dichroic beam splitter (DBS). The emitted signal is focused by lens L5 onto a photo multiplier tube (PMT). For characterising the DM, the surface of the DM is re-imaged onto a CCD camera (CCD) using the flip mirror FM1, flat mirror M5 and lenses L6 and L7. A reference arm is created using beam splitter BS1, flat mirrors M2, M3, M4 and beam splitter BS2. A coherence-gated fringe analysis method described elsewhere van Werkhoven et al. (2012, forthcoming) is applied to the fringe pattern generated onto CCD.

In the aberration correction experiments we imaged collagen fibre extracted from rat tail washed 4 times in distilled water. Following fixation in 4% paraformaldehyde, the fibre was washed in phosphate buffer saline

and then embedded in 3% agarose (Sigma-Aldrich chemie GmbH) in a 35 mm glass bottom dish (MatTek Corporation). We used 900 nm excitation light to generate the second-harmonic signal.

Preparation of the experiments

We first imaged a $20 \times 20\ \mu m^2$ region, approximately 33 µm deep into the collagen fibre. The region is labelled with A in Fig. 4.3. The influence of the size of the region used for the aberration correction has been studied elsewhere Zeng et al. (2012). For a certain setting of the deformable mirror $\mathbf{r}$, the region is scanned using the xyz stage. The corresponding value of the image quality metric $\tilde{y}$ is measured as the mean image intensity Jesacher et al. (2009); Débarre et al. (2009b); Facomprez et al. (2012); Zeng et al. (2012), i.e., the mean pixel value recorded over the region. The pixel dwell time is 0.5 ms and the sampling is 24×24 pixels^2. With these settings, the xyz stage does not reach the full 20 µm distance in the x scanning direction, which is the fast axis. This was not an issue since such a coarse sampling was only used to perform the aberration correction in a short time Facomprez et al. (2012). The final images taken after the aberration correction were recorded with a higher sampling. The image deformation due to both the non-linearity and non-uniform speed of the xyz stage were removed from the final images, using interpolation and the signals recorded with the position sensors of the xyz stage.

First, the static aberrations in the system due to misalignment and imperfections in the optical components were corrected. We used the non-zero initial aberration that was found during the calibration of the deformable mirror in section 4.2 (about 0.79 rms rad at 900 nm, mostly astigmatism). At this point the Nelder–Mead algorithm Lagarias et al. (1998) was executed four times to find a value $\mathbf{r}$ that maximises $\tilde{y}$. Unfortunately, this led to the saturation of two actuators, indicating that the stroke of the deformable mirror may be insufficient for completely suppressing the aberration in this region. We selected a slightly sub-optimal vector $\mathbf{r}$ from the vectors generated by the Nelder–Mead algorithm. For the selected vector, the maximum normalised voltage of the actuators was 0.72, i.e., $\|\mathbf{u}\|_\infty \leq 0.72$, $\tilde{y}$ improved by 3% and a total aberration of about 0.18 rms rad was suppressed. We used this state as the new initial condition for the rest of the experiments, i.e., $\mathbf{r} = \mathbf{0}$ is mapped to this setting of the deformable mirror. In the following sections, all the units in rad are referenced to the 900 nm excitation laser light.

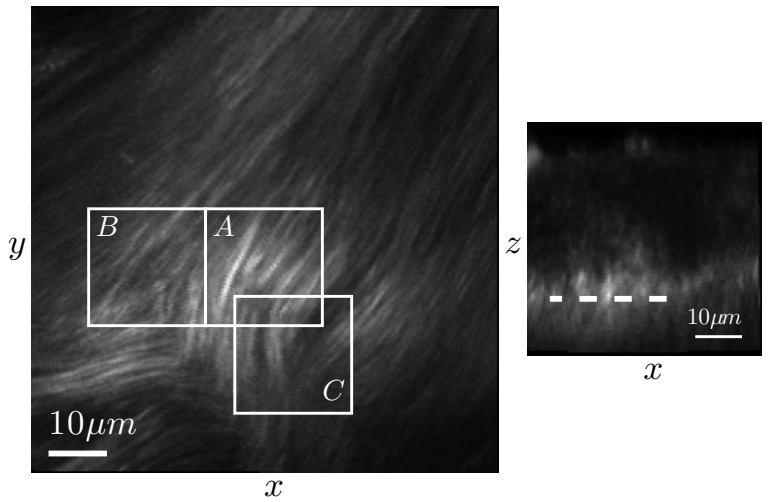


Figure 4.3: Cross sections of rat tail collagen fibre used in our experiments. The smaller image on the right hand side is an xz cross section (50×50 μm^2, 128×128 pixels^2). The dashed line denotes an xy cross section (80×80 μm^2, 256×256 pixels^2) approximately $33\,\mu$m deep, which is shown on the left hand side. Three different $20m \times 20$ μm^2 regions are marked.

Computation of the parameters of the quadratic polynomial using input-output measurements

We executed the computation of the parameters of the quadratic polynomial used for modelling the image quality metric multiple times. Each time, the sequence of input vectors consisted of two sub-sequences. The first subsequence contained random vectors $\mathbf{r}_1$, ..., $\mathbf{r}_{250}$ and was used for the validation and the cross-validation. The second subsequence contained 70 fixed vectors (each vector having a single non zero component). The second subsequence was included because the 250 random vectors may be insufficient to uniformly sample the N-dimensional space of the inputs. In Antonello et al. (2012) 3750 random vectors were used, but this was impractical here, due to the time necessary to move the xyz stage. The maximum rms rad of the input aberrations did not exceed 0.81. This value was empirically tuned by examining the goodness of fit as a function of the the maximum rms Antonello et al. (2012).

For each input vector in the sequence, we measured the corresponding output of $\tilde{y}$. The resulting input-output data, i.e., collections $\tilde{y}_1$, ..., $\tilde{y}_{320}$ and $\mathbf{r}_1$, ..., $\mathbf{r}_{320}$ were used to formulate the following optimisation prob-

lem Antonello et al. (2012)

$$
\begin{aligned}
&\min_{\tilde{\mathbf{x}}} \|\tilde{A}\tilde{\mathbf{x}} - \tilde{b}\| \quad \text{s.t.} \\
&\tilde{A} = \begin{bmatrix} 1 & \mathbf{r}_1^T & -\mathbf{r}_1^T \otimes \mathbf{r}_1^T \\ \vdots & \vdots & \vdots \\ 1 & \mathbf{r}_{320}^T & -\mathbf{r}_{320}^T \otimes \mathbf{r}_{320}^T \end{bmatrix}, \\
&\tilde{b} = \begin{bmatrix} \tilde{y}_1 & \dots & \tilde{y}_{320} \end{bmatrix}^T, \\
&\tilde{\mathbf{x}} = \begin{bmatrix} c_0 & \mathbf{c}_1^T & \mathrm{vec}(Q)^T \end{bmatrix}^T, \\
&\quad Q \succeq 0,
\end{aligned}
\tag{4.20}
$$

where vec(·) denotes the vectorisation operation and × the Kronecker product. This programme was solved using CVXOPT Andersen et al. (2013) (see Andersen et al. (2011) for further details).

Validation and cross-validation of the computed parameters

The results of applying (4.20) in region A marked in Fig. 4.3 are shown in Fig. 4.4. The computation of the parameters was repeated six times. Each time, a new input-output data set was acquired, $\mathcal{D}_1$, ..., $\mathcal{D}_6$. For each input-output data set, optimisation (4.20) was solved generating six sets of parameters, each set comprising c_0, $\mathbf{c}_1$ and Q. The sets are denoted as $\mathcal{M}_1$, ..., $\mathcal{M}_6$. We quantified the goodness of fit for all combinations of $\mathcal{D}$ and $\mathcal{M}$ by means of the $\mathcal{R}^2$ indicator. Using the random input subsequence of $\mathcal{D}_i$ and $\mathcal{M}_j$, we computed the predicted output $\hat{\mathbf{o}} \in \mathbb{R}^{250}$ of the image quality metric. The input-output data points obtained from the deterministic input subsequence of 70 vectors were discarded and not included in the computation of the $\mathcal{R}^2$ indicator. The following equations were employed for the $\mathcal{R}^2$ indicator,

$$
\begin{aligned}
&\mathcal{R}^2 = 1 - S_r/S_t, \qquad S_r = \|\mathbf{o} - \hat{\mathbf{o}}\|^2, \\
&S_t = \|\mathbf{o} - \bar{o}\mathbf{1}\|^2, \quad \bar{o} = (1/250)\mathbf{1}^T\mathbf{o},
\end{aligned}
\tag{4.21}
$$

where $\mathbf{o}$ is the measured output of $\mathcal{D}_i$. An $\mathcal{R}^2 = 1$ implies a perfect fit of the experimental data.

Fig. 4.4 reports $\mathcal{R}^2$ indicators that are close to one, implying a good fit of the experimental data. A good fit is also found for the combinations that are off the main diagonal. Here the parameters c_0, $\mathbf{c}_1$ and Q allow to accurately predict cross validation output data. Similar results were found for the other two regions marked in Fig. 4.3.

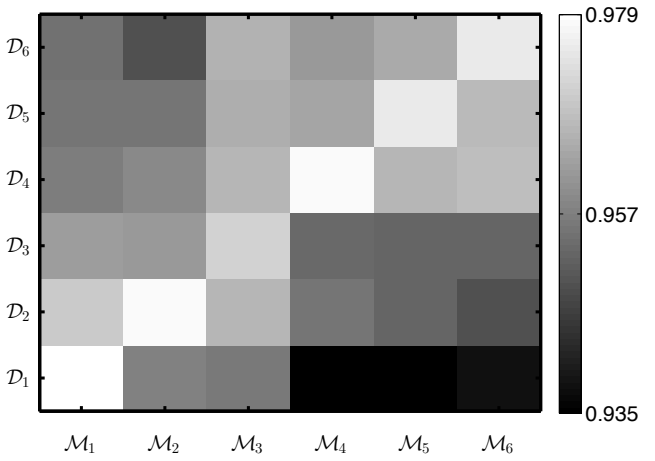


Figure 4.4: Validations and cross validations of the computation of c_0, $\mathbf{c}_1$ and Q using (4.20). The computation has been performed 6 times in region A in Fig. 4.3. $\mathcal{D}_i$ denotes the input-output data taken during the i-th time. $\mathcal{M}_i$ denotes the set of parameters (c_0, $\mathbf{c}_1$ and Q in (4.20)) computed from $\mathcal{D}_i$. For each combination $\mathcal{M}_i$ and $\mathcal{D}_j$, the i-th random input subsequence and $\mathcal{M}_j$ are used to compute the predicted output $\hat{\mathbf{o}} \in \mathbb{R}^{250}$. Each rectangle reports the goodness of fit ($\mathcal{R}^2$, see (4.21)) computed comparing $\hat{\mathbf{o}}$ with the corresponding measured output $\mathbf{o} \in \mathbb{R}^{250}$ of $\mathcal{D}_j$. A value of one for the goodness of fit indicates that the model fits the data without error. High values of the goodness of fit are reported in all combinations showing that (4.20) is a robust method to compute the parameters.

Correction of the residual aberration

In this section we apply the aberration correction algorithm described in section 4.3. First, we attempt to further reduce the residual aberration in region A, which is marked in Fig. 4.3. Some aberration may not have been completely suppressed by the Nelder–Mead algorithm, which was applied to region A in section 4.4. We therefore expect no improvement or a small improvement in region A. Second, we also apply the aberration correction to regions B and C, where the Nelder–Mead algorithm was not applied. Here we expect some improvement, provided that the aberrations found in regions B and in C are different from the aberration found in region A.

We take the parameters c_0, $\mathbf{c}_1$ and Q that were computed using $\mathcal{D}_6$ in section 4.4. In order to apply the algorithm, the following modified parameters need to be used, i.e., $c_0' = c_0 + (1/4)\mathbf{c}_1^T Q^{-1}\mathbf{c}_1$, $\mathbf{c}_1' = \mathbf{0}$ and $Q' = Q$. This modification is necessary since, for simplicity, in section 4.3 we neglected the linear term $\mathbf{c}_1$. This term corresponds to the aberration that

is present when computing the parameters of the quadratic polynomial, (see Antonello et al. (2012) fore more details). The aberration correction experiment is applied in the three regions using the corresponding parameters for each region. A summary of the results is given in Fig. 4.5.

In Fig. 4.5a, the normalised measurements of the image quality metric are reported for region A (curve with $\circ$ markers), region B (curve with $\square$ markers) and region C (curve with $\star$ markers). The measurements are normalised using the corresponding maximum recorded measurement $\tilde{y}_{\max}$ in each region. The initial value of $\tilde{y}$ is reported at sample time $k = 0$. This measurement is not supplied to the aberration correction algorithm. The data collection step is performed between time $k = 1$ and $k = 8$ inclusive, where $N + 1 = 8$ trial aberrations are applied. From time $k = 9$ onwards, the aberration correction step is applied.

Marginal improvement is found in region A (curve with $\circ$ markers) and B (curve with $\square$ markers), where aberrations of about 0.38 and 0.37 rms rad are corrected. The rms of each aberration is estimated using $\|H_h V_{l2} V_1 \mathbf{x}^{\text{ls}}\|$, adjusted for the 900 nm excitation light. In region C, (curve with $\star$ markers) an estimated 1.27 rms rad aberration is corrected, leading to an improvement of 20% of the image quality metric. Nevertheless, two actuators of the deformable mirror were saturated.

Two 256×256 pixels^2 images of region C are reported in Fig. 4.5b and Fig. 4.5c. These images are recorded before ($k = 0$, Fig. 4.5b) and after ($k = 24$, Fig. 4.5c) the aberration correction. The cross sections marked in the images are reported in a single graph in Fig. 4.5d. The image taken at time $k = 24$ is 18% brighter and shows finer detail in the bottom and right parts. Here some structure of the fibre was not visible at time $k = 0$. The improvement is less clear when examining the left and top parts of the region. One possible reason for the variability of the improvement is that the aberration is not spatially invariant over the considered field of view. Finally, we note that the applied correction was not optimal since two actuators of the deformable mirror were saturated. We conclude by observing that this improvement after the aberration correction is compatible with what was achieved by running four iterations of the Nelder–Mead algorithm in section 4.4.

Validation of the aberration correction algorithm

To assess whether the aberration correction algorithm is effectively removing aberration we performed a different kind of experiment. First we introduce a known amount of aberration using the deformable mirror. We then apply the aberration correction algorithm to suppress this

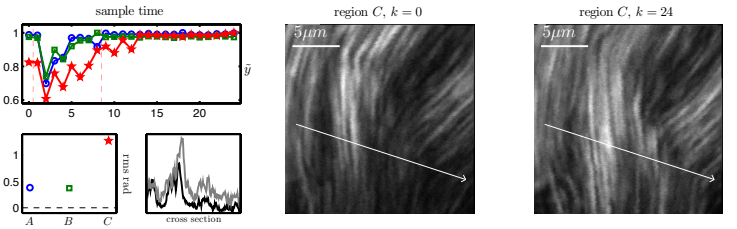


Figure 4.5: Summary of three aberration correction experiments. (a) evolution of the normalised image quality metric. The experiments were performed in region A (curve with $\circ$ markers), B (curve with $\square$ markers) and C (curve with $\star$ markers), which are marked in Fig. 4.3. For each region, the corresponding parameters computed by solving (4.20) were used. $\tilde{y}_{\max}$ is the maximum measurement of $\tilde{y}$ in each region. The estimated rms rad of each aberration is 0.38 for region A, 0.37 for region B and 1.27 for region C. (b) 256×256 pixels^2 image of region C at sample time $k = 0$. (c) 256×256 pixels^2 image of region C at sample time $k = 24$. (d) cross sections taken along the arrows marked in (b) and (c), black for (b) and grey for (c).

aberration. The algorithm is not supplied with any information about the known aberration. Finally we evaluate the residual aberration by comparing the estimate of the aberration provided by the algorithm with the known aberration. This kind of experiment is commonly employed in the literature to assess the effectiveness of the aberration correction Débarre et al. (2007b); Facomprez et al. (2012); Zeng et al. (2012); Antonello et al. (2012).

Fig. 4.6 reports a summary of the correction of 20 random aberrations introduced with the deformable mirror in region A. The upper plot in Fig. 4.6 shows some statistical indicators of the normalised measurements of the image quality metric. The measurements have been normalised using the maximum measurement of the metric $\tilde{y}_{\max}$ that is recorded throughout the 20 experiments. The median, 25th and 75th percentiles are computed in each time instant, see the caption of Fig. 4.6 for a detailed legend. The same analysis has been made for the residual aberration and is reported in the lower plot in Fig. 4.6. The rms of the residual aberration is computed as the rms of the difference between the known aberration introduced by the deformable mirror and the respective estimate provided by the algorithm. From this figure, we conclude that the image quality metric is consistently maximised, as the median is close to 1 after the aberration correction is applied from sample time $k = 9$ onwards. This is consistent with the reduction in the residual aberration reported in the lower

plot.

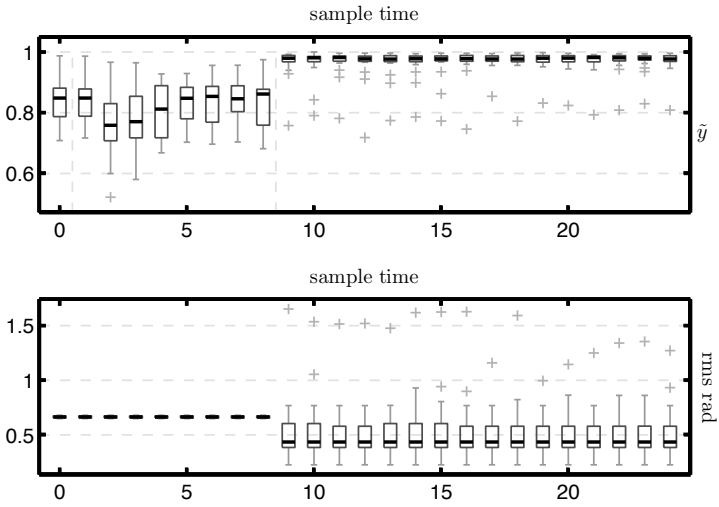


Figure 4.6: Summary of the correction of 20 random aberrations induced by the deformable mirror in region A, which is marked in Fig. 4.3. The upper plot reports the normalised measurements of the image quality metric. The measurements are normalised using the maximum measurement $\tilde{y}_{\max}$ that is recorded throughout the 20 experiments. At time $k = 0$ the initial value of $\tilde{y}$ is reported, this data point is not supplied to the aberration correction algorithm. Between time $k = 1$ and $k = 8$, the data collection step is executed. From time $k = 9$ onwards, the aberration correction step is applied. A statistical analysis is made at each time instant using the function `boxplot` from MATLAB. The tops and bottoms of the rectangles denote the 25th and 75th percentiles, the horizontal lines in the middle of the rectangles denote the medians, the whiskers extend to the furthest measurements not considered as outliers. The + symbols denote single outliers. The same statistical analysis is performed for the residual aberration and the results are shown in the lower plot.

The same experiments were performed in region B and C. In both cases we used the corresponding modified parameters, computed using $\mathcal{D}_6$ in section 4.4. The results are reported in Fig. 4.7 and Fig. 4.8. Whereas the results for region C are similar to the results obtained in region A, the results in region B do not show a good performance, since the medians of the residual aberration are comparable with the initial aberration before the correction.

Out of the 20 trials in region B, we report respectively the ones that resulted in the maximum and in the minimum improvement of $\tilde{y}$ in Fig. 4.9.

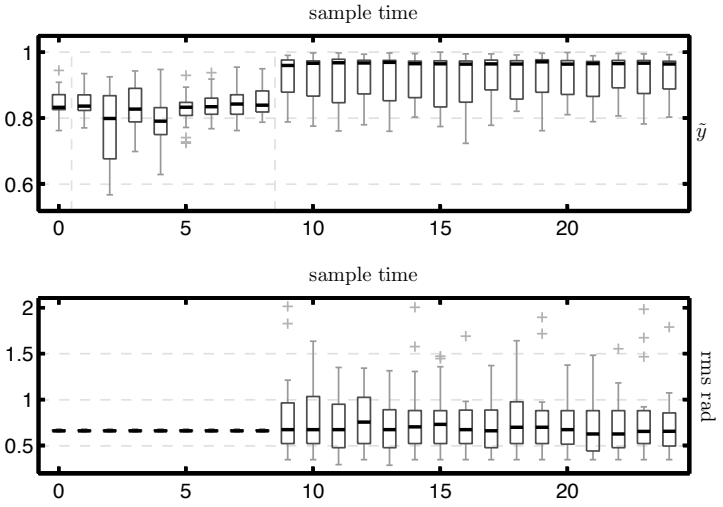


Figure 4.7: Summary of the correction of 20 random aberrations induced by the deformable mirror in region B. See the caption of Fig. 4.6 for a legend of the plots.

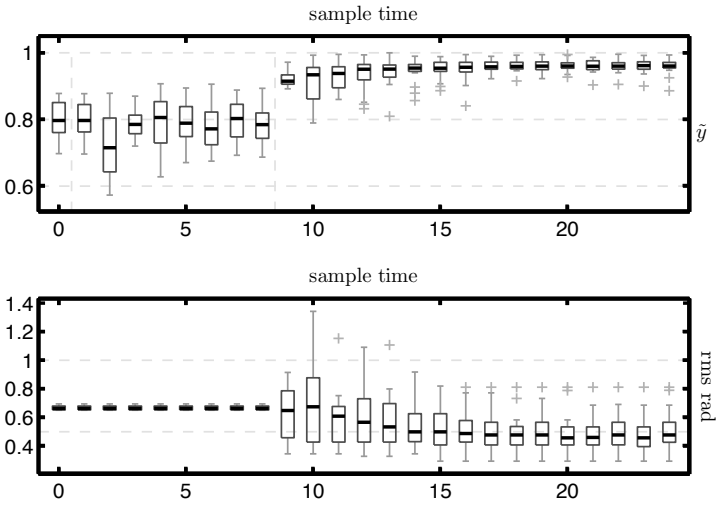


Figure 4.8: Summary of the correction of 20 random aberrations induced by the deformable mirror in region C. See the caption of Fig. 4.6 for a legend of the plots.

In Fig. 4.9a, some fine structure of the fibre is more visible after the correction, which is compatible with a successful aberration correction. On the other hand, in Fig. 4.9b, the aberration correction failed, as both the image after the correction is visually worse and the intensity is slightly decreased.

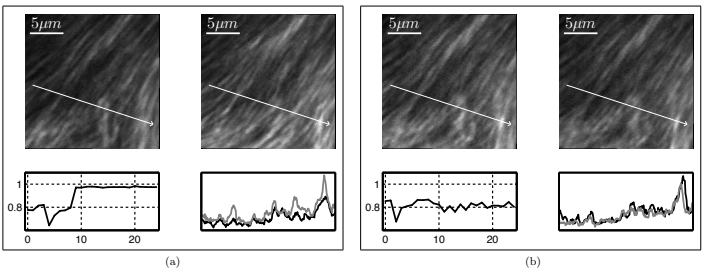


Figure 4.9: Two aberration correction experiments from the set of experiments reported in Fig. 4.7. These two experiments resulted respectively in the maximum (a) and the minimum (b) improvement of $\tilde{y}$. Both in (a) and (b), a 256×256 pixels^2 image is taken before (on the left, $k = 0$) and after (on the right, $k = 24$) the aberration correction. The graphs in the bottom of (a) and (b) show respectively the evolution of the normalised metric (on the left) and the cross sections indicated by the arrows in the images (on the right). In the cross section graphs, the dark and the light lines correspond respectively to $k = 0$ and $k = 24$.

The experiments resulting in the maximum and in the minimum improvement of $\tilde{y}$ in region C are also reported in Fig 4.10. In Fig 4.10a, a successful aberration correction is shown, with a clear maximisation of $\tilde{y}$ and a noticeable improvement in the contrast of the image after the correction. In Fig 4.10b, the improvement is more marginal.

Analysis of the experimental results

We computed the correlation among different quantities to concisely assess the results of the 60 aberration correction experiments reported in Fig. 4.6, Fig. 4.7 and Fig. 4.8. Considering the last time instant $k = 24$, we set up a saturation indicator variable $s_1, ..., s_{60}$ that is 1 if saturation of some actuators of the deformable mirror occurred and -1 otherwise.

We computed a correlation of -0.2477 between the normalised measurement of the image quality metric and the saturation indicator variable. These two quantities are slightly inversely correlated, meaning that saturation of some actuators negatively affected the final value of the image

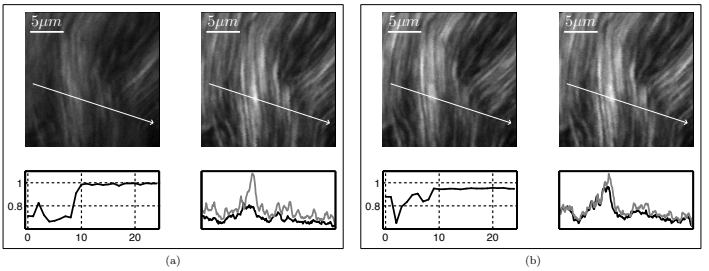


Figure 4.10: Two aberration correction experiments from the set of experiments reported in Fig. 4.8. These two experiments resulted respectively in the maximum (a) and the minimum (b) improvement of $\tilde{y}$. Refer to the caption of Fig. 4.9 for a detailed legend.

quality metric. The correlation between the final rms of the residual aberration and the saturation indicator variable was 0.3457. This positive correlation shows that a larger amount of residual aberration was found when the stroke of the deformable mirror was exhausted. Finally we computed a correlation of -0.7388 between the normalised measurement of the image quality metric and the rms of the residual aberration, which confirms that some aberration is removed by maximising the image quality metric. We conclude that saturation of the deformable mirror was an issue that hampered the results in our experimental validation.

Variations of the parameters over the field of view

We report spatial variations in the parameters c_0, $\mathbf{c}_1$ and Q. The parameters differed when computed respectively in region A, B and C in Fig. 4.3. Due to the variations, we were not able to apply the aberration correction algorithm using a single set of parameters, e.g., by correcting aberrations in region B and C using the parameters computed from region A. Variations in the parameters represent a challenge for model-based aberration correction algorithms, since the parameters are computed once only with a calibration experiment Neil et al. (2000); Booth et al. (2002); Booth (2006); Débarre et al. (2007b); Booth (2007c); Débarre et al. (2008b, 2009b); Jesacher et al. (2009); Song et al. (2010); Linhai & Rao (2011); Facomprez et al. (2012); Zeng et al. (2012); Antonello et al. (2012).

Parameter $\mathbf{c}_1$ depends on the non-zero aberration that is present when collecting the input-output measurements used in (4.20). This parameter can be removed by applying the aberration correction algorithm as done in

section 4.4. Instead, c_0 is dependent on the maximum value of the image quality metric, which differed in the three regions. We found variations in Q. For example, the largest eigenvalue of Q varied by about 30% in region B and by about 22% in region C with respect to its value in region A. The eigenvectors of Q were also affected. For instance, the eigenvector corresponding to the second largest eigenvalue of Q was rotated by about 7° in region B and by about 21° degrees in region C with respect to its orientation in region A.

We note that when using the algorithms discussed in Facomprez et al. (2012), the eigenvalues of Q are computed from scratch each time the aberration correction is applied (see section 4 in Antonello et al. (2012)). This is possible by applying at least $2N + 1$ trial aberrations and by approximating the solution of (4.12). Nevertheless, also these algorithms do not account for possible variations in the orientation of the eigenvectors, as found in our experiments. These variations affect the performance of the aberration correction, as investigated in section 3 in Facomprez et al. (2012). Therefore updating the parameters of the model as different regions of the specimen are acquired is an open research challenge.

4.5 Conclusion

In this paper we present a procedure to define a new set of basis functions for the control of the deformable mirror. The new basis functions can be made approximately orthogonal to a set of Zernike modes. This is necessary for applying aberration correction in scanning microscopy applications, where the deformable mirror must not induce the x-tilt, y-tilt and defocus aberrations.

The second contribution concerns the algorithm used for the aberration correction. We consider computing the least-squares estimate of the unknown aberration. Although this problem is non-convex in general, the least-squares estimate can be computed efficiently by exploiting results already applied in the solution of localisation problems Beck et al. (2008). Once the least-squares estimate is computed, the aberration correction is applied by maximising the image quality metric.

We implement the aberration correction algorithm in a second-harmonic microscope. First we validate a previously proposed Antonello et al. (2012) method to compute the parameters of the quadratic polynomial used to model the image quality metric directly from input-output measurements. Second, we validate the aberration correction algorithm discussed in this paper. We also report the measurement of variations in the parameters of

the quadratic polynomial over the field of view.

Acknowledgements

This research is supported by the Technology Foundation STW under project number 10433 that is part of the Smart Optics Systems perspective programme directed by Prof. M. Verhaegen.

Part II

Peculiar transiting exoplanets

5 Analysis and interpretation of 15 quarters of Kepler data of the disintegrating planet KIC 12557548 b

Based on Tim van Werkhoven, Matteo Brogi, Ignas Snellen and Christoph Keller 2013, *Astron. Astrophys.* 561, A3

The *Kepler* object KIC 12557548 shows irregular eclipsing behaviour with a constant 15.685 h period, but strongly varying transit depth. The object responsible for this is believed to be a disintegrating planet forming a trailing dust cloud transiting the star. A 1-D model of an exponentially decaying dust tail was found to reproduce the average eclipse in intricate detail. Based on radiative hydrodynamic modelling, the upper limit for the planet mass was found to be twice the mass of the Moon. In this paper we fit individual eclipses, in addition to fitting binned light curves, to learn more about the process underlying the eclipse depth variation. Additionally, we put forward observational constraints that any model of this planet-star system will have to match. We manually de-correlated and de-trended 15 quarters of *Kepler* data, three of which were observed in short cadence mode. We determined the transit depth, egress depth, and stellar intensity for each orbit and search for dependencies between these parameters. We investigated the full orbit by comparing the flux distribution of a moving phase window of interest versus the out-of-eclipse flux distribution. We fit short cadence data on a per-orbit basis using a two-parameter tail model, allowing us to investigate potential dust tail property variations. We find

two quiescent spells of ~30 orbital periods each where the transit depth is <0.1 per cent, followed by relatively deep transits. Additionally, we find periods of on-off behaviour where >0.5 per cent deep transits are followed by apparently no transit at all. Apart from these isolated events we find neither significant correlation between consecutive transit depths nor a correlation between transit depth and stellar intensity. We find a three-sigma upper limit for the secondary eclipse of 4.9×10^{-5}, consistent with a planet candidate with a radius of less than 4600 km. Using the short cadence data we find that a 1-D exponential dust tail model is insufficient to explain the data. We improved our model to a 2-D, two-component dust model with an opaque core and an exponential tail. Using this model we fit individual eclipses observed in short cadence mode. We find an improved fit of the data, quantifying earlier suggestions by Budaj (2013) of the necessity of at least two components. We find that deep transits have most absorption in the tail, and not in a disk-shaped, opaque coma, but the transit depth and the total absorption show no correlation with the tail length.

5.1 Introduction

Rappaport et al. (2012) discovered the peculiar target KIC 12557548 in the *Kepler* database, which shows dips in the light curve with a period of about 15.7 h (constant to within 10^{-5}), but a depth varying from less than 0.2 per cent to up to 1.3 per cent. The phase-folded light curve shows no signs of ellipsoidal light variations, which limits the mass of planet candidate KIC 12557548 b[1] to $<3\,M_J$.

Rappaport et al. (2012) exclude several other scenarios for this target, including a dual-planet system and a low-mass eclipsing stellar binary. The authors argue for a disintegrating planet as the most likely scenario in which the close proximity to the host star causes parts of the planet's surface to evaporate. The evaporated gas drags dust along, which subsequently eclipses parts of the star. Because of the stochastic nature of this process, transits have variable depth. The authors also find evidence of forward scattering due to this dust cloud, which creates a slight increase in intensity just before ingress. This scenario puts an upper limit on the planet's escape velocity, such that a Mercury-mass planet is the most likely candidate. The authors qualitatively investigate the likeliness of this scenario and find it to be consistent with observations.

[1] Henceforth denoted as KIC 1255b.

Subsequently, Brogi et al. (2012) quantitatively investigated the planet hypothesis using a 1-D model to constrain the transit parameters, the shape of the dust cloud, and the average particle size. They find that a dust cloud with ∼0.1 µm-sized particles best matches the observed, average eclipse light curve. The authors also find different system parameters for subsets of the transits consisting of relatively shallow (0.2 per cent to 0.5 per cent) and deep (>0.8 per cent) eclipses.

Perez-Becker & Chiang (2013) investigate the evaporation dynamics of this planet candidate and, through radiative hydrodynamic modelling, argue that it is losing mass at a rate of $\dot{M} \gtrsim 0.1\, M_{\oplus}$ Gyr^{-1}. They conclude that the planet candidate has a mass of less than 0.02 $M_{\oplus}$, or twice the mass of the Moon. According to the authors, KIC 1255b may have lost up to 70 per cent of its initial mass, with only the inner iron core left. Budaj (2013) investigates the dust tail properties in more detail and argues that the particle size changes along the tail, where micron-sized particles best explain ingress while 0.1 to 0.01-micron sized particles fit egress best.

Here we investigate the shape of individual eclipses and provide statistical constraints on the system. We extend the previous model from Brogi et al. (2012) to a pseudo 2-D variant where the vertical extent of the cloud of dust is not neglected, and an opaque core is included as a disk centred on the planet candidate. The short cadence data for quarters 13 through 15 allow us to fit the model on a per-transit basis to compare individual transits. Using the 15-quarter coverage of the target we investigate correlations between the transit depth, the depth at egress, and the stellar activity as well as variations of these parameters. Additionally, we derive a three-sigma upper limit for the secondary eclipse of 4.9×10^{-5}, which is consistent with an object radius smaller than 4600 km for an albedo of 1.

The data reduction is explained in Sect. 5.2, and the light curve analysis is presented in Sect. 5.3. Section 5.4 describes the improved model with which we investigate per-orbit properties for the short cadence data as well as an analysis and interpretation of those results.

5.2 Data reduction

Kepler is a 0.95 m-aperture Schmidt telescope with a 16° diameter field of view, observing 156 453 stars with a 95 megapixel, 42 science CCD focal plane (Koch et al., 2010). The telescope was launched on 6 March 2009 and started observing on 2 May 2009. Unfortunately, in May 2013 *Kepler* went into safe mode due to a second reaction wheel failing.

At the basic level, frames are recorded by integrating for 6.02 s, followed

by a 0.52 s readout. Data is then stored in short cadence (SC) mode by co-adding 9 frames, giving a cadence of 58.86 s with a 54.18 s exposure time (Gilliland et al., 2010) or in long cadence (LC) mode by co-adding 270 frames resulting in a cadence of 29.4244 min and an exposure time of 27.1 min (Jenkins et al., 2010a). Because of the limited telemetry bandwidth, *Kepler* observes no more than 512 targets – about 0.3 per cent of the total – in SC mode.

At the time of writing, 15 quarters of *Kepler* data are publicly available at the Multi-Mission Archive (MAST[2]) at the Space Telescope Science Institute. KIC 12557548 was observed in LC mode for quarters 1 to 12, and in SC mode during the last quarters (13 to 15). The latter data have a 30 times shorter exposure time, resulting in a factor of $\sqrt{30} \approx 5.5$ lower signal-to-noise ratio, assuming pure Poisson noise.

The *Kepler* pipeline (Jenkins et al., 2010b) delivers the light curve as *Simple Aperture Photometry* flux (`SAP_FLUX`), as well as an automatically reduced *Presearch Data Conditioning* (PDC) flux (`PDCSAP_FLUX`). The PDC aims to remove systematic errors from the raw flux time series. Since KIC 1255 exhibits a very peculiar light curve, the pipeline may have difficulty in automatically reducing the data. Indeed, we find that for the LC data, deep transit data are flagged as outliers. The automatically reduced SC data has a noise level approximately 1.5 times higher than our manually reduced data. Because of this, we started the data reduction from the raw `SAP_FLUX` and manually de-correlate the data.

The data reduction is explained in more detail below. First we selected usable data and filtered out bad data. For `SAP_FLUX` only, we manually decorrelated the signal using linear co-trending basis vectors supplied by the *Kepler* pipeline to remove systematics from the signal. Finally, we removed the stellar signal by de-trending the flux.

Data selection

We used all LC data from *Kepler* quarters 1 through 15, and all SC data, which were taken during quarters 13 to 15. We ignored non-finite data and most non-zero `SAP_QUALITY` flagged data. The *Kepler* pipeline erroneously marks some data with flag 2048 ("Impulsive outlier removed after cotrending", (Fraquelli & Thompson, 2012, p. 20) for transits deeper than ~0.8 per cent. Considering that this is a highly variable target, it is not unlikely that the automated *Kepler* pipeline has some difficulties, and indeed the *Kepler* archive manual warns of these cases (`kepcotrend` docu-

[2] http://archive.stsci.edu/kepler/data_search/search.php

mentation[3]). Ignoring these data points would not correctly sample deep transits, and we therefore included these flagged data. We have not found any other SAP_QUALITY flags that correlate with the orbital phase, and we therefore removed all other flagged data from our analysis.

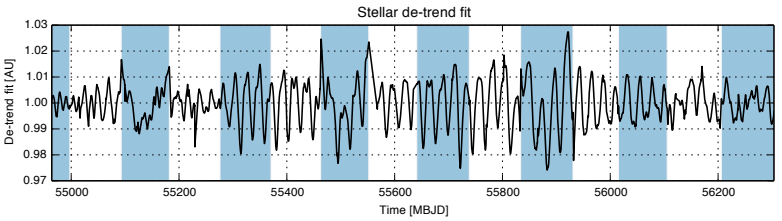


Figure 5.1: De-trending fit for the de-correlated, long cadence SAP_FLUX, showing the stellar variability during quarters 1 through 15. Odd quarters are shaded.

De-correlating systematics

To remove systematics from the light curve, the *Kepler* archive provides a set of 16 linear co-trending basis vectors (LCBV) for each detector, which are derived from a subset of highly correlated and quiet stars and are meant to remove satellite systematics from the data (Fraquelli & Thompson, 2012, p. 21). The pipeline automatically de-correlates the light curve for all targets against these LCBVs, but the *Kepler* archive manual recommends to manually de-correlate signals that are highly variable.

Linear co-trending basis vectors are only available for LC data and cannot be constructed manually for the SC data because not enough targets are observed in SC to generate a set of LCBVs. To de-correlate the SC data, we linearly interpolated the LCBVs on the SC time points before de-correlation. We successfully used this technique to minimise the out-of-eclipse residual variance of the SC data of quarters 13 to 15. We note that this approach is unable to correct for systematics occurring on timescales significantly shorter than 29 min (i.e. that of the long cadence).

We de-correlated the flux as follows. First we interpolated the *Kepler* LCBVs on SAP_FLUX exposure times for the SC data. Then we excluded the primary eclipse at orbital phase $\varphi \in [-0.15, 0.20]$ from the fitting process. We least-squares fit all 16 vectors to the remaining data and selected the first n vectors that reduce the out-of-eclipse variance significantly (see Fig. 5.2). In our case, we used $n = 2$ vectors for the SC data and $n = 2$ to 5 for the each of the LC data quarters.

[3] http://keplergo.arc.nasa.gov/ContributedSoftwareKepcotrend.shtml.

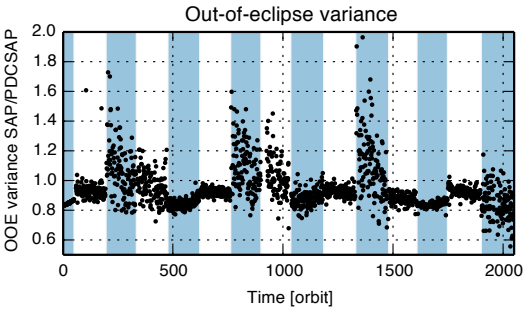


Figure 5.2: Ratio of long cadence `SAP_FLUX` over `PDCSAP_FLUX` standard deviation during the out-of-eclipse part of the orbit (phase $\varphi \in [-0.4, -0.1]; [0.2, 0.4]$). Odd quarters are shaded. The flux variance differs per quarter.

Removing stellar variability

When the de-correlation is performed correctly, it must preserve all astrophysical signal. This includes stellar variability, which has to be removed to analyse the eclipse behaviour. Since the data has jumps in both time and flux due to gaps in the data, we de-trended the data piece-wise per block, where a block is delimited by either a jump in time or flux. A time jump is defined as a gap in data of more than twice the data cadence, a flux jump is defined as a change of flux of more than 0.5 per cent (LC) and 2 per cent (SC) difference between two consecutive data points, and if the difference in the median of the 20 data points around the jump differs by more than 3.5 times the standard deviation of those points. These parameters were chosen on trial and error basis.

For each block, we masked out the transits during orbital phase $\varphi \in [-0.15, 0.20]$ when fitting. For blocks up to 10 data points or spanning a single orbit, we normalised the data by the median. For blocks up to 200 data points we normalised using a second-order polynomial fit, and for larger blocks we fit a cubic spline to the data with a knot at each orbit at $\varphi = 0.5$, exactly opposite to the transit. We flagged data that are near the edge of a block: within one orbital period of the edge, or data outside the outer knots of a spline fit. Both of these flags are excluded in subsequent analyses since these data are poorly fitted and may introduce unwanted errors in subsequent steps. By removing the stellar variation with only one degree of freedom per orbit (i.e. a spline knot), we keep the transit signal intact.

Table 5.1: Comparison between the literature values of the orbital period of KIC 1255b and the value determined through our analysis.

Source	Period (d)
Rappaport et al. (2012)	0.653 56(1)
Budaj (2013)	0.653 552 1(15)
This analysis	0.653 553 8(1)

Table 5.2: Best-fit 1-D model parameters and their one-sigma uncertainties, as derived from the Markov-chain Monte Carlo (MCMC) analysis. From top to bottom: impact parameter b, mid-transit phase offset $\Delta\varphi_0$, decay factor λ, total extinction cross-section (in units of stellar area) c_e, asymmetry parameter g, and single-scattering albedo ϖ.

Parameter	Average	Deep	Shallow
b	$0.63^{+0.022}_{-0.023}$	$0.51^{+0.031}_{-0.022}$	$0.62^{+0.030}_{-0.029}$
$\Delta\varphi_0 \times 10^3$	$-1.35^{+0.43}_{-0.44}$	$-1.01^{+0.55}_{-0.51}$	$0.27^{+0.66}_{-0.80}$
λ	$5.83^{+0.21}_{-0.19}$	$5.84^{+0.23}_{-0.24}$	$4.80^{+0.24}_{-0.20}$
c_e	$0.0227^{+0.0020}_{-0.0013}$	$0.0415^{+0.0048}_{-0.0035}$	$0.0139^{+0.0009}_{-0.0007}$
g	$0.809^{+0.033}_{-0.045}$	$0.810^{+0.026}_{-0.039}$	0.809 (fixed)
ϖ	$0.49^{+0.090}_{-0.088}$	$0.96^{+0.12}_{-0.11}$	0.49 (fixed)

Figure 5.1 shows the de-trending fit by which we divided the de-correlated SAP_FLUX to remove the stellar signal. The variability in this plot is caused by star spots coming in and out of view due to the stellar rotation. We find a period for the stellar rotation of 22.65 ± 0.05 d by auto-correlating the signal, which is consistent with the findings of Kawahara et al. (2013).

Orbital parameters

Once the signal is de-correlated and the stellar variability is removed, we computed the orbital period of the planet candidate by minimising the difference between the phase-folded LC data and the best-fitting model from Brogi et al. (2012) in a least-squares sense. This method is similar to phase dispersion minimisation (Stellingwerf, 1978) with the exception that we use a model instead of smoothed data. For these analyses, we used barycentric Julian days expressed in barycentric dynamical time, as given by the *Kepler* pipeline (Eastman et al., 2010). The values for the period with one-sigma uncertainties are listed in Table 5.1. We do not find significantly different values for the period for PDCSAP_FLUX and SAP_FLUX

reduced data. Additionally, we fitted the 1-D model parameters described in Brogi et al. (2012) using 15 quarters of data (see Sect. 5.4 and Sect. 5.4 for more details). Likewise, we calculated these parameters for deep (more than 0.8 per cent) and shallow transits (0.2 per cent to 0.5 per cent) as well as all transits. The results are shown in Table 5.2.

5.3 KIC 1255b analysis

After reducing all available data, we investigated them on a per-orbit basis. We numbered the orbits sequentially; orbit 1 is the first orbit observed by *Kepler* and 2050 is the last orbit in these data. We used light curves for 1773 orbits in total, excluding bad data as identified by the *Kepler*pipeline (i.e. `SAP_QUALITY`) as well as data that was poorly de-trended. As an example orbit 1700 is shown in Fig. 5.3 for both the long cadence (LC) and short cadence (SC) data.

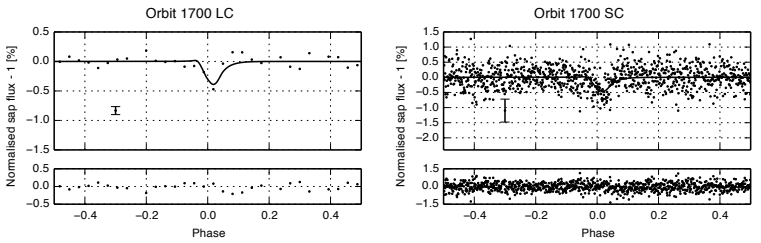


Figure 5.3: Orbit 1700 of KIC 1255b in long- (*top*) and short cadence (*bottom*) data. The short cadence data has a 30× higher temporal resolution, but correspondingly higher noise. We overplot the best-fitting 1-D model in the upper panels.

Primary eclipse

Since the LC data has a cadence of 29.4 min, we have an average of 31.984 data points per 15.6854 h orbit and are limited to about $31.984 \times 0.1 \approx 3$ data points per individual transit (see for example Fig. 5.3, top panel). Hence we only fit the depth of the transits using the 1-D model with the best-fitting parameters of Brogi et al. (2012). To measure the depth, we performed a least-squares fit using a scaled best-fitting model to the data at $\varphi \in [-0.1, 0.2]$.

For LC data, we convolved the model data with the *Kepler* exposure time before fitting, as explained in Brogi et al. (2012). For the SC data we did not convolve the data with the *Kepler* exposure time, and there is sufficient

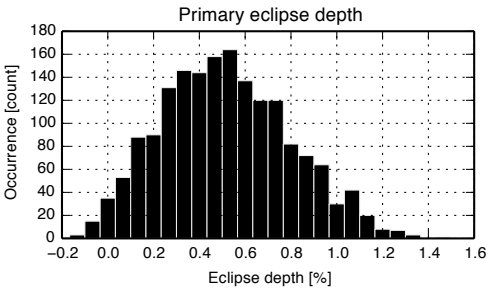


Figure 5.4: All long cadence transit depths obtained from scaling the best-fitting model to SAP_FLUX data. Here we plot the depth at the minimum of those scaled models. The few negative eclipse depths are due to noise in the data.

temporal resolution to additionally fit the onset of the transit. The transit onset is determined by shifting the best-fitting model in time. A least-squares fit yields both onset and depth simultaneously for each transit.

A histogram of all LC transit depths is shown in Fig. 5.4. Figure 5.5 shows the normalised, binned, phase-folded light curve for short cadence and long cadence data of transits of depth 0.8 to 1.0 per cent, with the residuals between the model and the data shown below. The model deviates from the data during egress, showing that the simple, exponential tail model cannot explain the light curve in full detail.

In Fig. 5.6 we plot the transit depth as a function of orbit, with a 30-orbit (approximately one stellar rotation) moving average plotted as a solid black line. There are two quiet regions around orbit 50 and 1950 (MBJD 55 000 and 56 250) during which the average transit depth is on the order of 0.1 per cent, the former is plotted in the left panel of Fig. 5.7. The quiet periods are followed by periods during which the moving-average depths are 0.8 per cent and 0.7 per cent, approximately 0.1 to 0.2 per cent deeper than the average. Additionally, we find times at which the transits appear in an on-off-pattern where >0.5 per cent deep transits are followed by apparently no transit signal at all for up to 11 orbits. This occurs for example around orbit 1076 and to a lesser extent around orbit 940 (MBJD 55 667 and 55 578). In the right panel of Fig. 5.7 we show the period around orbit 1076.

Transit depth correlation

Rappaport et al. (2012) argue that the dust has a sublimation lifetime of 3×10^4 s, or 8 h, such that it does not survive one orbit. To test this, we

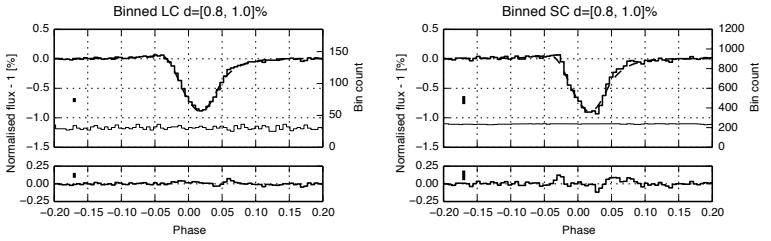


Figure 5.5: Phase-folded, normalised flux binned in 192 bins for long cadence (*top*) and short cadence (*bottom*) data for transits with depths from 0.8 to 1.0 per cent. The number of data points in each bin is indicated with horizontal bars according to the right y-axis. The model is based on the parameters from the 'deep' column in Table 5.2. The best-fitting model deviates from the short cadence data during both ingress and egress. The vertical bar indicates the median three-sigma error of the binned flux. The residual RMS for these data in the phase displayed are 2.0×10^{-4} and 3.7×10^{-4} respectively, showing a greater mismatch between the short cadence data and the 1-D model.

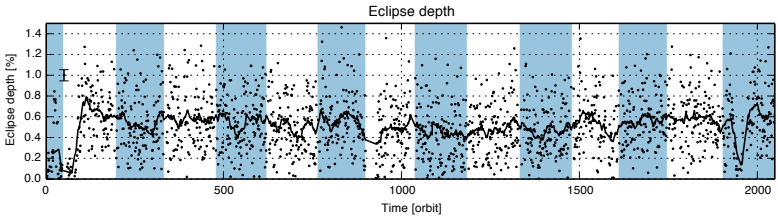


Figure 5.6: Transit depth as a function of orbit for long cadence data. The solid line is a 30-orbit moving average. There are two quiet regions around orbit 50 and 1950. Odd quarters are shaded.

investigated the correlation between consecutive transit depths, as well as transit depth and consecutive egress depth, defined as the depth during $\varphi \in [0.055, 0.15]$. In the former case we expect a correlation if the dust generation lasts longer than one orbit, such that subsequent transits are correlated. When investigating the transit and consecutive egress depths, we quantified the longevity of a dust cloud, such that deep transit clouds survive as an elongated tail in the next transit, which would lead to a particularly long egress signature. In this scenario, we would observe a deep transit as caused by a recent outburst where the dust is close to the planet candidate, eclipsing a large part of the star. Under the influence of gravity and the stellar radiation pressure this cloud would deform into a comet-

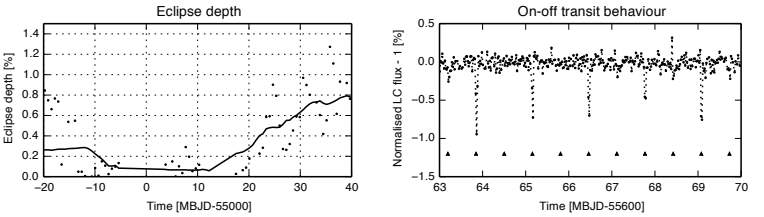


Figure 5.7: *Left:* transit depth as a function of orbit, showing the first quiet period around orbit 50 in Fig. 5.6 in more detail. *Right:* long cadence flux as a function of time where KIC 1255b shows on-off-like behaviour in the transit depth. The triangles indicate the midpoint of the transits.

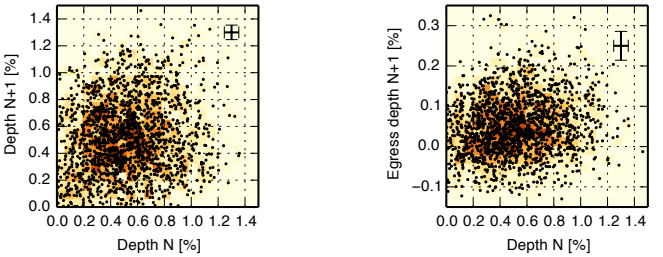


Figure 5.8: *Left:* transit depth plotted against the following transit depth, showing no correlation ($R^2 = 0.026$). This constrains the dynamic processes that underlie the dust cloud generation. The lower-left corner is slightly overpopulated due to two quiet periods of KIC 1255b as explained in Sect. 5.3. The cross indicates the median error for the depth estimate. *Right:* transit depth versus egress intensity for consecutive orbits for the long cadence data. We observe no significant correlation ($R^2 = 0.0019$).

like tail during the orbit, such that a more tenuous dust tail would eclipse the star during the following orbit.

For this analysis we selected all pairs of sequential orbits. We plot the depth versus the depth and egress depth for consecutive orbits in the LC data in the left and right panel of Fig. 5.8, respectively.

There is no obvious correlation between the depth of consecutive transits (Pearson's $R^2 = 0.026$) nor between depth and consecutive egress depth ($R^2 = 0.0019$). The absence of a correlation between consecutive transit depths indicates that the process underlying the dust generation is erratic and occurs on time scales shorter than one orbit. Furthermore, the lack of correlation between transit and consecutive egress depths is consistent

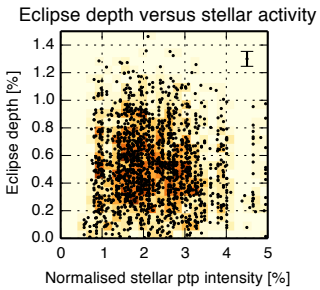


Figure 5.9: Transit depth versus stellar activity, as determined by the peak-to-peak value of the stellar intensity in a 24 d moving window centred at the time of the transit.

with earlier findings by Rappaport et al. (2012) and puts an upper limit on the dynamical time scale of the dust tail dissipation at one orbit, i.e. 15.7 h. This is consistent with Perez-Becker & Chiang (2013), who calculated the dynamical timescale of the dust tail and found it to be approximately 14 h.

We also investigated the correlation between the stellar intensity in the *Kepler* band and the transit depth. The intensity is a proxy for the stellar activity, and a correlation might reveal the influence of stellar activity on the transit depth and thus gas and dust generation, as observed for Mercury (Potter & Morgan, 1990). As we observe the rotational modulation of the stellar intensity due to star spots, and not the absolute intensity, the amplitude of the cyclical intensity variations is also influenced by the spatial and size distributions of the star spots. As is observed on the Sun, we assume that there is a positive correlation between the amplitude of the cyclical intensity variations and the magnetospheric effects affecting the planet and its dust tail.

We measured the stellar variability as a by-product of our data detrending, as described in Sect. 5.2. From the stellar variability we computed the peak-to-peak value in a moving 24 d window (about one stellar rotation period) to obtain a proxy for the stellar activity, where we assume that a higher amplitude corresponds to a more active star. We show the transit depth versus the stellar activity at each orbit in Fig. 5.9. We observe no significant correlation between transit depth and our stellar activity proxy (R^2 = 0.0011). Kawahara et al. (2013) performed a time-series analysis of the transit depth evolution and found a periodicity close to the rotation period of the star, which they interpret as an influence of stellar

activity on the atmospheric escape of the planet candidate. We have not found any evidence for this using our method.

Secondary eclipse

We divided the LC flux during the expected secondary eclipse ($\varphi \in [0.45, 0.55]$) by the out-of-eclipse flux. Using this ratio ensures that any residual deviations from unity in the flux due to inaccuracies in the detrending will not be mistaken for a real signal at the time of the secondary eclipse. This ratio is plotted in Fig. 5.10, with the mean error of the data points shown on the left. The weighted average and the error of all data give a three-sigma upper limit for the secondary eclipse of 4.9×10^{-5}. Using simple geometry we find that a planet candidate with radius smaller than 4600 km, or an Earth-sized object with an albedo of ~0.5 is consistent with this finding. Furthermore, we find no correlation ($R^2 = -0.012$)

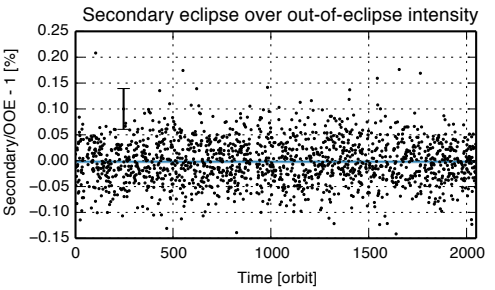


Figure 5.10: Long cadence flux during the secondary eclipse divided by flux during out-of-eclipse periods as a function of orbit. The weighted-average three-sigma upper limit is 4.9×10^{-5}. The black error bar is the mean error of the individual data points.

between the depth of the primary eclipse and the secondary eclipse, indicating that even deep primary eclipses do not leave any significant, back-scattering dust after half an orbit.

Full orbit

To investigate features at phases other than the transit, we compared the flux distribution at different phases against the distribution of the out-of-eclipse flux. Since we expect a flat light curve during the out-of-eclipse phase, significant deviations from zero in the differences between these distributions are indications of potential features due to the dust. We plot

the difference between the out-of-eclipse distribution and a distribution of the flux in a moving phase bin as histograms in Fig. 5.11. We used 100 intensity bins, and we oversampled the phase bin direction four times, resulting in 320 overlapping phase bins each 0.0125 wide in phase. There is

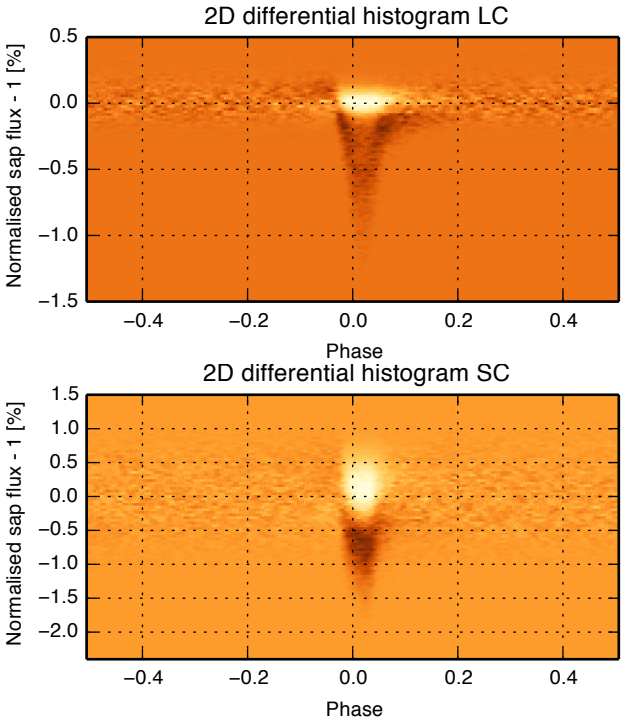


Figure 5.11: 2-D differential histogram comparing the flux distribution at a certain phase against the out-of-eclipse flux distribution for all long- (*top*) and short cadence (*bottom*) data. Excess flux compared to the out-of-eclipse histogram shows up as bright areas. We use 100 intensity bins and 320 phase bins of width 0.0125 such that the phase bins are 4× oversampled. Note the excess flux around phase $\varphi = -0.1$ in the long cadence data indicating the forward scattering, which is not visible in the SC data.

larger spread in the SC data, and both the forward scattering peak as well as the egress are less visible than in the LC data. As expected from the analysis in Sect. 5.3, there are no signs of a secondary eclipse in either plot. Besides features already investigated above, we see no features apart from those explained by the 1-D model.

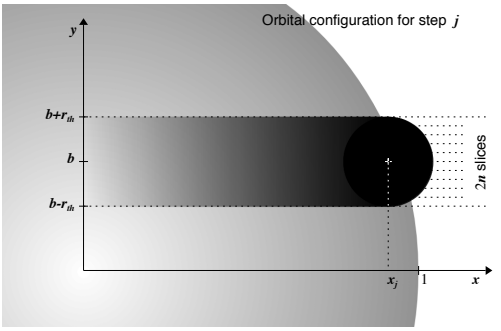


Figure 5.12: Geometry of the two-component, 2-D model. x is along the planet candidate orbit, y is perpendicular to that. $r_{\rm th}$ is the radius of the opaque core.

5.4 Cloud model

Model for the LC data

The *Kepler* long cadence (LC) data are fitted by employing the one-dimensional model of Brogi et al. (2012). Since the LC data have an insufficient number of points per transit, it is unrealistic to fit individual eclipses; we therefore only fit phase-folded and binned LC data with this model, realising that we may average events that differ by more than just the transit depth.

Because of the improved data reduction and the much larger dataset available, we also compared the best-fit values and uncertainties derived here with those of our previous work.

Model for the SC data

The analysis of the morphology of the short cadencedata (SC) (see Sect. 5.3 and Fig. 5.5) reveals residual structures after subtracting a properly scaled 1-D model. This suggests that this model is not sufficient to describe the full morphology of the KIC 1255 light curve, when observed with a higher time resolution.

In an attempt to better fit the *Kepler* SC data, we developed an improved cloud model by accounting for the vertical extent of the cloud, and by splitting the cloud structure into an opaque component for the dust around the planet, and an optically thin, exponentially-decaying cloud of dust following it. When not accounting for the vertical extent of the cloud, the impact parameter derived from fitting the data is an *effective* impact parameter,

meaning that it is averaged over the non-uniform brightness of the stellar disk occulted by the cloud of dust. This could partially explain why the impact parameters for *deep* and *shallow* data differ (see Table 5.2). In addition, it is plausible that – at least for the strongest outbursts – the amount of dust ejected from the planet is sufficient to create an opaque envelope of material, which could also explain why the maximum transit depth seems to be limited to 1.2 to 1.4 per cent. Indeed, it has been pointed out that the cloud should consist of at least two component with different properties for the dust (Budaj, 2013).

As a trade-off for the higher number of free parameters, we neglected the scattering component in this new model. This is justified by the much lower photometric precision (by a factor of $\sim$ 45) of the *Kepler* SC, unfolded and unbinned data, with respect to the phase-folded and binned LC data. The scattering properties are almost exclusively constrained by the small peak just before ingress, which is completely buried in the noise in these data, as shown in Fig. 5.11. The egress part of the light curve is also affected by scattering, but it is degenerate with the properties of the tail, meaning that by changing the exponential decay of the tail we could mimic the effects of a forward-scattering component. Therefore, although this model may miss some of the physics involved, our aim is not to derive physical parameters or to compare them to the previous study, but to better understand the basic geometry of the cloud, and therefore we only focus on its structure.

The 2-D model shares the mathematical background of the previous 1-D code, i.e. it generates the light curve by convolving the profiles of the dust cloud and the stellar disk. The vertical extent of the cloud is described by $2n$ slices centred on the position of the planet candidate. The model is parametrised via the radius of the opaque, circular cloud of dust ($r_{\rm th}$), the scale-length of the exponentially-decaying, optically thin tail of dust ($\lambda_{\rm tail}$), and the impact parameter of the transit (b). All these quantities are expressed in units of the stellar radius (R_S), and both vertical and horizontal directions are quantised with the same step size Δr. This is defined by subdividing the total length of the orbit into m steps. For a semi-major axis of $a = 4.31$, we have $\Delta r = 2\pi a/m \approx 27.1/m$. We denote the vertical direction with y (i.e. perpendicular to the orbit) and the horizontal direction with x (i.e. along the orbit). The curved path of the planet across the stellar disk for $b \neq 0$ is approximated with a straight line.

Figure 5.12 shows the geometry and the main quantities involved in this new 2-D model. The position along the orbit is traced through the vector $\vec{x} = [x_1, x_2, \cdots, x_m]$, in the interval $(-\pi a, \pi a)$, which is related to the or-

bital phase φ via $\vec{\varphi} = \vec{x}/a$. The zero point for the orbital phase coincides with the centre of the star and – for $j = 1$ in Eq. (5.6) – with the position of KIC 1255b. At any time, the centre of the planet candidate, which also coincides with the centre of the opaque part of the cloud, is placed at $x = x_j$ and $y = b$. The vertical position of the slices is defined by the vector $\vec{y}_{\rm cloud} = [y_1, y_2, \cdots, y_{2n}]$, which is centred on $y = b$ and in the interval $(b - r_{\rm th}, b + r_{\rm th})$. For each slice i, the absorption properties of the cloud are defined by

$$\mathcal{C}_i(\vec{x}) = \begin{cases} 1 & \text{for } |\, x \,| \leq x_{{\rm th},i} \\ \mathrm{e}^{-\vec{x}/\lambda_{\rm tail}} & \text{elsewhere} \end{cases}, \tag{5.1}$$

where $\vec{x}_{\rm th} = \sqrt{r_{\rm th}^2 - (\vec{y}_{\rm cloud} - b)^2}$ denotes the intersections of the orbital vector $\vec{x}$ with the opaque disk of dust, for each of the $2n$ slices.

The stellar profile is computed by assuming a quadratic limb-darkening law, meaning that the intensity of the stellar disk is expressed as

$$I_{\rm S}(\mu) = 1 - u_1(1-\mu) - u_2(1-\mu)^2, \tag{5.2}$$

where u_1, u_2 are the quadratic limb-darkening coefficients for a star of similar properties as KIC 12557548 (Claret & Bloemen, 2011a), while μ is the cosine of the angle between the line of sight of the observer and the normal to the stellar surface. Therefore, μ is a function of two variables (the x, y coordinates), and it is meaningful only for points inside the stellar disk. For a given slice i, the intersection between the orbital vector $\vec{x}$ and the stellar disk (i.e. the edge of the star) is given by

$$x_{{\rm star},i} = \sqrt{1 - y_{{\rm cloud},i}^2}\,. \tag{5.3}$$

The stellar brightness profile is therefore

$$\mathcal{S}_i(\vec{x}) = \begin{cases} I_{\rm S}(\vec{x}, y_{{\rm cloud},i}) & \text{for } |\, x \,| \leq x_{{\rm star},i} \\ 0 & \text{elsewhere} \end{cases}. \tag{5.4}$$

The previous relation can be expressed explicitly by substituting

$$\mu = \sqrt{1 - \vec{x}^2 - y_{{\rm cloud},i}^2} \tag{5.5}$$

in Eq. (5.2). The total light curve is finally given by convolving the stellar and the cloud profiles for each slice, which is

$$F(\vec{x}) = 1 - \sum_{i=1}^{2n} \sum_{j=1}^{m} \frac{\mathcal{S}_i(\vec{x})\, \mathcal{C}_i(\vec{x} - x_j)}{\mathcal{S}_{\rm tot}}. \tag{5.6}$$

The normalisation factor $\mathcal{S}_{\rm tot}$ (the total flux from the star) is precomputed by discretising the full limb-darkened stellar disk with the same step size as for the model, and summing over all pixels of the matrix. It is therefore a much more straightforward normalisation than in our previous 1-D model, where the stellar profile for $b = 0$ had to be normalised such that the sum of its elements was equal to unity.

In this two-dimensional model, fractional pixel coverage is also taken into account via a linear approximation. This is particularly important for very small $r_{\rm th}$, or when b approaches $(1 + r_{\rm th})$, and prevents us from using a too-small (and computationally demanding) Δr. We validated our model with the Mandel & Agol (2002) transit code by choosing a very small value for $\lambda_{\rm tail}$, which is equivalent to neglecting the optically thin part of the cloud. For an optimal value of $m = 3000$, found via trial and error, the two models differ by less than 10^{-5}, which is at least two orders of magnitude better than the accuracy of the *Kepler* short cadence data.

Fitting of LC data

The updated parameters of KIC 1255b, as derived from the fitting of 15 quarters of *Kepler* LC data, are listed in Table 5.2. For deep and shallow transits, we observe a better mixing of the Markov-chain Monte Carlo (MCMC) chains (quantified following Gelman & Rubin (1992)) than in the previous work, which is expected from the much larger amount of data used. This also results in more symmetric posterior distributions and less correlation between the parameters.

Fitting of SC data

To compare our 2-D model with the 1-D model, we fit it to the subset of deep transits (see Table 5.2), and perform a similar analysis as in Sect. 5.3. Even though we excluded the forward scattering from the 2-D model, the residuals are reduced as compared to the 1-D model, as shown in Fig. 5.13, indicating a better match to the data.

Subsequently, we selected 213 light curves sampled in SC mode with transit depths greater than 0.5 per cent, no discontinuities in the data and sufficient photometric precision to fit the 2-D model to the individual transits to be used for a per-orbit analysis. Each individual orbit contains 447 points.

We fit the geometry of the dust cloud individually for each light curve by first performing a least-squares fit with a grid of widely spaced parameter values. The set of parameters corresponding to the minimum χ^2 is then

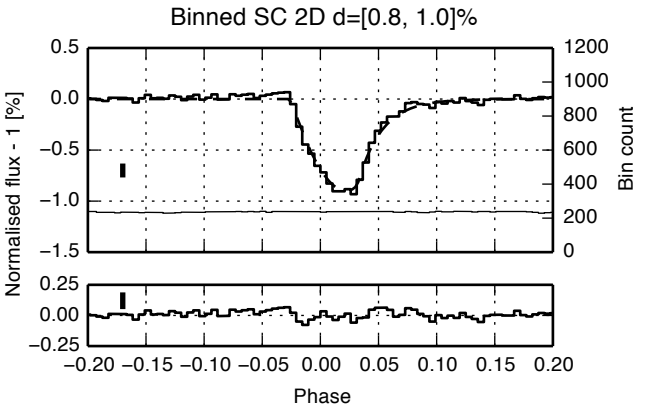


Figure 5.13: Phase-folded, normalised flux binned in 192 bins for short cadence data for transits with depths from 0.8 to 1.0 per cent, as in Fig. 5.5, but instead using the 2-D model to fit the data. Even though the 2-D model does not include forward scattering, we observe an improved fit and reduced residuals, with a 16 per cent lower residual RMS of 3.1×10^{-4}.

used as input for a single MCMC chain of 50 000 steps. This assures that the chain is already started in the proximity of the global χ^2 minimum, and does not get stuck in a local minimum.

By performing the analysis with all free parameters, we notice that a large fraction of the MCMC chains fail to converge. However, by investigating those chains that do converge, we observe that all transits are consistent with the same impact parameter ($b = 0.6 \pm 0.1$). This suggests that our two-component, two-dimensional model is a better approximation for the cloud of dust, and does measure a *true* impact parameter, as opposed to the previous one-dimensional model. Therefore, we fixed the impact parameter to 0.6 and recomputed the MCMC chains.

To verify that the model parameters are not degenerate, we computed artificial light curves by:

1. Assuming a small tail ($\lambda = 0.3$) and only varying the opaque part;
2. Fixing $\tau_{\rm th} = 0.01$ and only varying the exponential tail.

We added noise based on the photometric precision of the *Kepler* SC data, and fit these simulated data with the method described above. In both cases the retrieved parameters are compatible with the single-parameter distributions given as input, and we do not obtain a mixing between the

two model components, as is observed in the real data. This suggests that the two components are indeed present in the data.

In our model the size of the opaque part also drives the vertical extent of the cloud. This means that, if the tail covers a certain effective area (integrated over the x and y direction), its specific length λ has to change as a function of the size of the opaque core. Larger $r_{\rm th}$ means a wider cloud, which requires a faster decay in order to maintain approximately the same effective area. Simply plotting $r_{\rm th}$ versus λ therefore shows a correlation due to the model, and tells little about the system itself (see Fig. 5.14). For this reason we plot the absorption by the core as $\pi\, r_{\rm th}^2$ and of the tail as $2\, r_{\rm th}\, \lambda_{\rm tail}$.

The results of the short cadence transit analysis are shown in Fig. 5.14. Besides plotting the model parameters, we also show the absorption of the core versus the absorption of the tail and the tail length versus the total absorption since these better describe the physics of the system. The total absorption is highly correlated with transit depth (bottom panel; as determined in Sect. 5.3) which is expected: more material will yield a deeper minimum. We find no relation between the tail length and the transit depth or total absorption (bottom panel). Although shallow transits are evenly distributed between the core and tail absorption, deeper eclipses appear to have most absorption in the tail, and not in a disk-shaped, opaque coma (bottom-right corner of middle panel).

Our model is certainly not the only possible geometric description of KIC 1255b, and it is unclear whether the two components we propose are an accurate physical description of the system. One could, for example, fix the vertical extent of the cloud to allow for a variable optical depth for the core with an exponentially decreasing tail. The main point is that the data suggest the necessity of at least two independent components for the cloud, in agreement with Budaj (2013).

5.5 Conclusions

We have manually de-correlated and de-trended 15 quarters of *Kepler* data, of which three quarters were observed in short cadence mode for KIC 12557548, and investigated statistical constraints on system dynamics as well as a per-orbit analysis using three quarters of SC data.

We find two quiescent spells of KIC 1255b of ~30 orbits each around orbit 50 and 1950 where the average transit depth is 0.1 per cent. These two periods appear to be followed by periods where the transit depth is 0.1 to 0.2 per cent deeper than the average transit depth. Additionally,

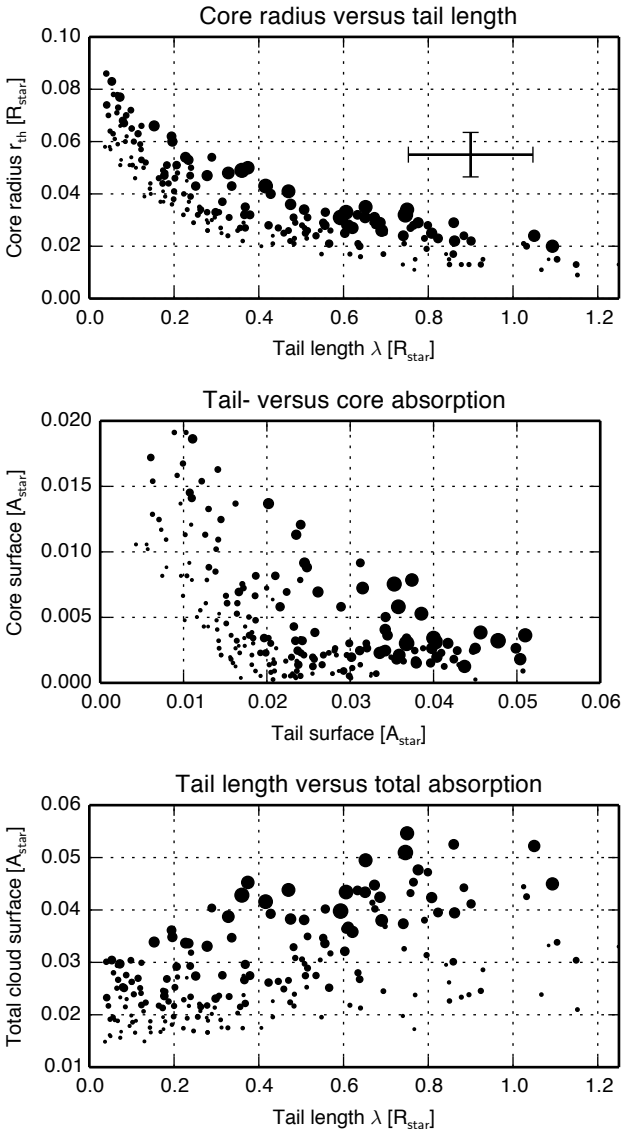


Figure 5.14: Deep (>0.5 per cent) short cadence transit *Kepler* data fitted with the two-component 2-D model. The point surface scales with the transit depth squared (for clarity) as determined in Sect. 5.3. The top plot shows the best-fitting model parameters $r_{\rm th}$ and λ for each individual SC transit. The correlation in this plot is largely due to the model (see text). We show the effective area of the tail versus the size of the opaque core (*middle*), and the tail length versus total absorption (*bottom*). We observe no clear relation between the tail or the core absorption, while the transit depth scales with the total absorption. Very deep transits seem to require a strong tail component, as is evident by the cluster of points in the bottom-right corner of the middle panel. The void in the triangular lower-left corner in the middle panel is due to the exclusion of shallow transits, where iso-transit depth lines run diagonally.

we find times at which the transits show on-off-like behaviour, where >0.5 per cent deep eclipses are followed by hardly any eclipse at all. Aside from these isolated events, we find no significant overall correlation between consecutive transit depths, nor between transit depth and consecutive egress depth. This implies that the majority of the dust does not survive a single orbit, and that the process underlying the dust generation occurs erratically. The independence of transit depth and consecutive egress depth implies that an opaque dust cloud yielding a deep transit does not survive to form part of a more tenuous and stretched-out dust cloud during the next eclipse, in agreement with Perez-Becker & Chiang (2013).

Furthermore, we put an upper limit of 4.9×10^{-5} on the secondary eclipse. This constrains the radius of a planet candidate to less than 4600 km, or one Earth radius for an albedo of ~0.5.

We find a significant discrepancy when fitting our previous 1-D model to the SC data around egress. Our improved model adds a second dimension and represents the dust tail with two components, which better fits individual SC transits than the old 1-D model. We verified that the two components are not degenerate in the model and are data-driven.

Our results suggest that a 1-D, exponentially decaying dust tail is insufficient to represent the data. We find that deep transits have most absorption in the tail, and not in a disk-shaped, opaque coma, but the transit depth or total absorption show no correlation with the tail length. Although our model is only one possible interpretation of the cloud structure, there is also qualitative evidence (Budaj, 2013) of the need of at least two components. We anticipate that models including realistic physics and geometry are required to fully understand this peculiar system.

Acknowledgements

The authors would like to thank Ernst de Mooij for his constructive discussion concerning the modelling of this object.

This paper includes data collected by the *Kepler* mission. Funding for the *Kepler* mission is provided by the NASA Science Mission directorate.

All of the data presented in this paper were obtained from the *Mikulski Archive for Space Telescopes* (MAST). STScI is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS5-26555. Support for MAST for non-HST data is provided by the NASA Office of Space Science via grant NNX09AF08G and by other grants and contracts.

6 Analysis of *SuperWASP* J140747.93−394542.6 transit fine-structure: hints of exomoons

Based on Tim van Werkhoven, Matthew Kenworthy and Eric Mamajek *Mon. Notices Royal Astron. Soc.* (submitted)

A recently discovered $V = 12.3$ mag K5 star in the *SuperWASP* database shows a peculiar light curve with a highly structured eclipse pattern covering a timespan of at least 54 d with maximum dimming of at least 3.3 mag (Mamajek et al., 2012). The central eclipse is surrounded by two 1 mag eclipses at ± 12 d and ± 26 d. The authors speculate that the star is eclipsed by a substellar companion with an extended and highly structured ring system. To investigate the nightly light curve structure and to confirm the multiple ring hypothesis, we have carried out a calibrated reduction of the *SuperWASP* data, removing both systematic errors and periodic stellar variability. We count at least 24 inflection points on ingress and 16 on egress, consistent with the presence of at least 24 rings in this disk. By measuring the light curve gradient, we find implied speeds for the eclipsing object that are incompatible with a closed Kepler orbit with $P = 2.3$ yr. We propose several scenarios that could give rise to such light curve gradients and find that azimuthal ring structure (analogous to "spokes" seen in Saturn's rings) can account for the observed light curve. The highly structured ring system also implies the presence of exomoons orbiting the secondary companion.

6.1 Introduction

Mamajek et al. (2012) discovered the peculiar light curve of a young ~16 Myr, 0.9 $M_\odot$, $V = 12.3$ mag K5 star (1SWASP J140747.93−394542.6 = ASAS J140748−3945.7, hereafter J1407) in *The Super Wide Angle Search for Planets* (henceforth *SuperWASP*) database. The light curve as observed by *SuperWASP* shows a deep >3.3 mag eclipse with two pairs of eclipses occurring symmetrically 12 d and 26 d before and after the eclipse midpoint, parts of which are confirmed by the lower-cadence *All Sky Automated Survey* data. Light curve variability induced by extended disk eclipses is not unique, as seen in the stars EE Cep (Mikolajewski & Graczyk, 1999) and ϵ Aurigae (Guinan & Dewarf, 2002) which show related complex asymmetric eclipses caused by discs orbiting the primary star.

The authors propose and reject several explanations for these observations. An eclipse by an isolated (sub)stellar companion alone would not yield a >95 per cent dimming, nor does it explain the eclipse fine-structure seen in the nightly photometry. An eclipsing binary where a red giant eclipses a fainter, bluer star is rejected on spectroscopic grounds; there is no evidence for a giant in J1407, nor would it give rise to the structure observed in the eclipse weeks before and after the primary eclipse. A disk orbiting an old stellar remnant companion obscuring the primary is unlikely because the system is very young. A circumbinary or circumstellar disk obscuring the primary is unlikely because there is no excess infrared observed at J1407 indicative of such a thick disk. Additionally, the period is incompatible with this scenario. A disk orbiting J1407 obscuring the star light due to the relative motion between us and J1407 would yield a one-time eclipsing event, but is disregarded because it requires a very thin disk (aspect ratio height over radius $\sim 10^{-8}$) and it does not explain the symmetric dimming around the central minimum. Finally, a more massive companion obscured by a disc orbiting it is unlikely again due to lack of infrared excess.

The scenario Mamajek et al. (2012) propose involves a circumsecondary or circumplanetary disc (similar to ϵ Aurigae; Guinan & Dewarf, 2002), which obscures the primary star, J1407. Based on the absence of a second transit, the authors constrain the period $P > 2.33$ yr, yielding an orbital radius greater than 1.7 AU and a circular orbital velocity less than 21.7 km s^{-1} (assuming $m_1 = 0.9\ M_\odot$). They propose a preliminary 4-ring model with a central opaque disk accompanied by three discrete rings with opacities $\tau = 0.5, 0.2, 0.05$. To match the observations, this model has an orbital inclination of 89.955° and the ring plane is tilted by 13° to the orbital plane.

In this paper, we carry out a new reduction of the *SuperWASP* data, tak-

ing care of systematic inter-CCD differences and additionally remove stellar variability due to the rotation of the star (described in Sect. 6.2). We then investigate the nightly fine structure in the photometric signal, focusing specifically on the light curve gradient and its implications for the nature of the eclipsing structure (Sect. 6.4 and Sect. 6.5). We construct a toy-model for the ring system consisting of N discrete rings with variable opacity which we fit to the cleaned up photometric signal in Sect. 6.6. We propose several possible scenarios for these findings and conclude with suggestions to better constrain the system in Sect. 6.7.

6.2 Data reduction

The Super Wide Angle Search for Planets (*SuperWASP*; Pollacco et al., 2006) observatories are located at the Observatorio del Roque de los Muchachos on La Palma, and at the Sutherland Station of the South African Astronomical Observatory. Both facilities consist of a telescope with eight lenses (Canon 200 mm $f/1.8$), each mounted on a 2048 × 2048-pixel back-illuminated CCD (e2v). Each lens has a field of view of ∼64 deg², for a total of ∼482 deg² per telescope. The data acquisition cadence is 9 to 12 min for each camera (Butters et al., 2010).

The *SuperWASP* database has coverage of J1407 during the seasons of 2006 through 2008 (depicted in the upper panel of Fig. 6.1). The star was observed on CCDs 221, 227, and 228. CCDs 227 and 228 have coverage of J1407 during 2006, 2007 and 2008, while CCD 221 only has coverage during the 2007 season. The *SuperWASP* pipeline provides simple aperture photometry (Kane et al., 2004), as well as automatically reduced flux according to the `TAMUZ` correction (Cameron et al., 2007; Tamuz et al., 2005). We find that the automatic correction does not produce optimal photometry for this specific target, which can be expected considering the peculiar behaviour of the light curve (see Fig. 6.3) and the star's location at the corner of the field of view of the three cameras.

Because of this variability, we began our reduction with the simple aperture photometry flux instead of using the automatically reduced flux. First we de-correlated the systematic errors from the light curve by using a selected ensemble of nearby stars to act as a photometric standard, allowing direct comparison of the data from different CCDs (Sect. 6.2). After decorrelation, we removed the stellar variability from the light curve through phase dispersion minimisation (Sect. 6.2).

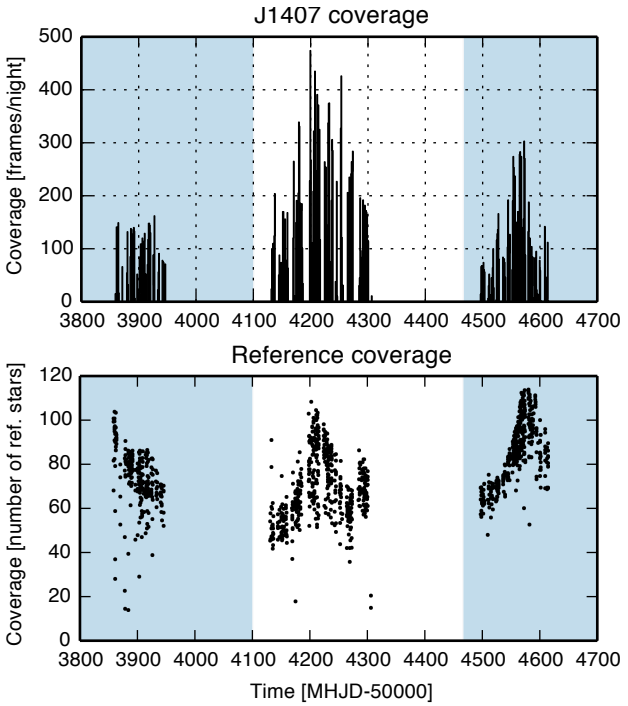


Figure 6.1: *Top:* Number of *SuperWASP* observations of J1407, summed for all CCDs. *Bottom:* Number of reference stars co-temporally observed with J1407 for all CCDs, averaged per day. Years 2006 and 2008 are shaded blue. The transit occurred in 2007 around MHJD 54 222.

Systematics de-correlation

The multiple CCDs of *SuperWASP* provide significantly different light curves, as observed by Mamajek et al. (2012) and Norton et al. (2011). To mitigate this effect, we de-correlated systematic errors between the CCDs based on Tamuz et al. (2005) and de-correlation techniques used by *Kepler* (Fraquelli & Thompson, 2012, p. 21)[1].

First we select a set of 388 reference stars in a cone of 20 arcmin centered on J1407 that are also observed by *SuperWASP*, R_{all}. The bottom panel in Fig. 6.1 shows the number of reference stars co-temporally observed with J1407, summed for all CCDs, averaged per day. Because of temporal gaps in both the target and reference data, we cannot perform a global analy-

[1] http://keplergo.arc.nasa.gov/ContributedSoftwareKepcotrend.shtml.

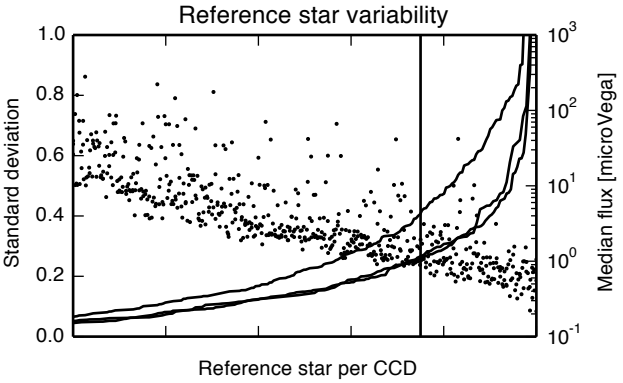


Figure 6.2: Flux variability (lines, left axis) and median flux (right axis, dots) for reference stars observed on the three different CCDs. The vertical line denotes the 75 per cent cut-off we chose for selecting quiet reference stars, only stars to the left are used in consequent de-correlation. While CCD 221 and 227 have similar variability roll-off, the upper solid line for CCD 228 indicates this CCD has noisier data. See text for details.

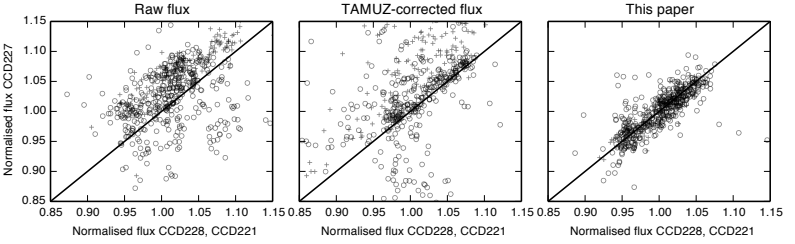


Figure 6.3: Hourly-binned flux from CCDs 221, 227 and 228 plotted against each other. The raw flux (FLUX) shows an offset between the CCDs as well as a significant scatter (Pearson's correlation coefficient $\rho = 0.49$), while the automatically reduced flux (TAMFLUX) is more consistent between different CCDs but still shows a large spread between the data ($\rho = 0.54$). The manually de-correlated flux shows no offset and significantly reduced scatter, indicating consistency between the three different CCDs ($\rho = 0.81$).

sis of a set of reference stars for our target star, and adapted the methods mentioned above to work on a point-by-point basis.

For each CCD separately, we used only reference stars for which at least 25 per cent of the data points are co-temporally observed with J1407, $\mathrm{R_{cotemp}}$.

$$\mathrm{R_{cotemp}} = \{x \in \mathrm{R_{all}} | \#(x_{\mathrm{exp}} \cup \mathrm{tgt_{exp}}) > 0.25\#(\mathrm{tgt})\}. \tag{6.1}$$

Of this subset we selected 75 per cent that have a low variance σ_{ref} (i.e. quiet stars), defined as the standard deviation of the median-normalised flux

$$\sigma_{\mathrm{ref}} = \mathrm{stddev}(F_{\mathrm{ref}})/\mathrm{median}(F_{\mathrm{ref}}). \tag{6.2}$$

We require photometrically quiet stars to measure the systematics of the cameras. This criterion was chosen to reject the most variable reference stars (see Fig. 6.2). In this way we obtained approximately 60 to 100 suitable reference stars, $\mathrm{R_{quiet}}$, per exposure of our target.

$$\begin{aligned}\mathrm{R_{quiet}} = \{x \in \mathrm{R_{cotemp}} | \\ \sigma(x/\mathrm{median}(x)) < \sigma(x/\mathrm{median}(x))_{75\%}\}.\end{aligned} \tag{6.3}$$

For this subset of reference stars, we calculated the mean of the median-normalised flux for all exposures co-temporal with our target.

$$F_{\mathrm{decorr}}(t) = \sum^{i \in R_{\mathrm{quiet}}} F_i(t)/\#(R_{\mathrm{quiet}}(t)) \quad \text{for } t \in t_{\mathrm{obs}}, \tag{6.4}$$

where t_{obs} are times at which J1407 was observed. We then divide the flux of J1407 by the thus obtained de-correlation vector F_{decorr} for each CCD.

Using the average of these reference stars suppresses any astronomical signal that might be present in an individual reference star, and any residual variation is then dominated by systematics. Since these stars are close to the target star, they are along nearly identical lines of sight through the atmosphere and camera optics, and so sample similar optical conditions. Additionally, these stars are close together on the CCD surface, such that they sample similar vignetting and electronic noise in the data.

Finally, to normalise the flux and allow inter-CCD comparison, we normalised the J1407 light curve by a first order polynomial fit for each CCD individually, excluding the transit window during MHJD 54 180 to 54 250. Through this de-correlation process, the photometry from the different CCDs now show significantly less scatter. The result is shown in Fig. 6.3, where the flux from different CCDs are plotted against each other.

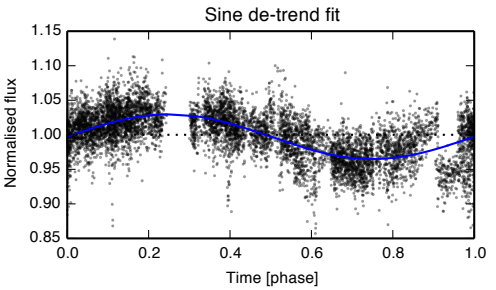


Figure 6.4: Normalised phase-folded *SuperWASP* flux for non-transit 2007 data with period 3.2011 ± 0.0002 d showing the stellar variability. The solid line is a best-fitting sinusoid with variable period.

Table 6.1: Measured period and period variability for J1407 in three years of *SuperWASP* data.

Season	Period [d]	Period change [s d^{-1}]	Amplitude [%]
2006	3.2183(4)	22(2)	4.98(3)
2007	3.2011(2)	31(1)	3.57(2)
2008	3.2066(3)	0.8(6)	5.00(2)

Stellar variability de-trending

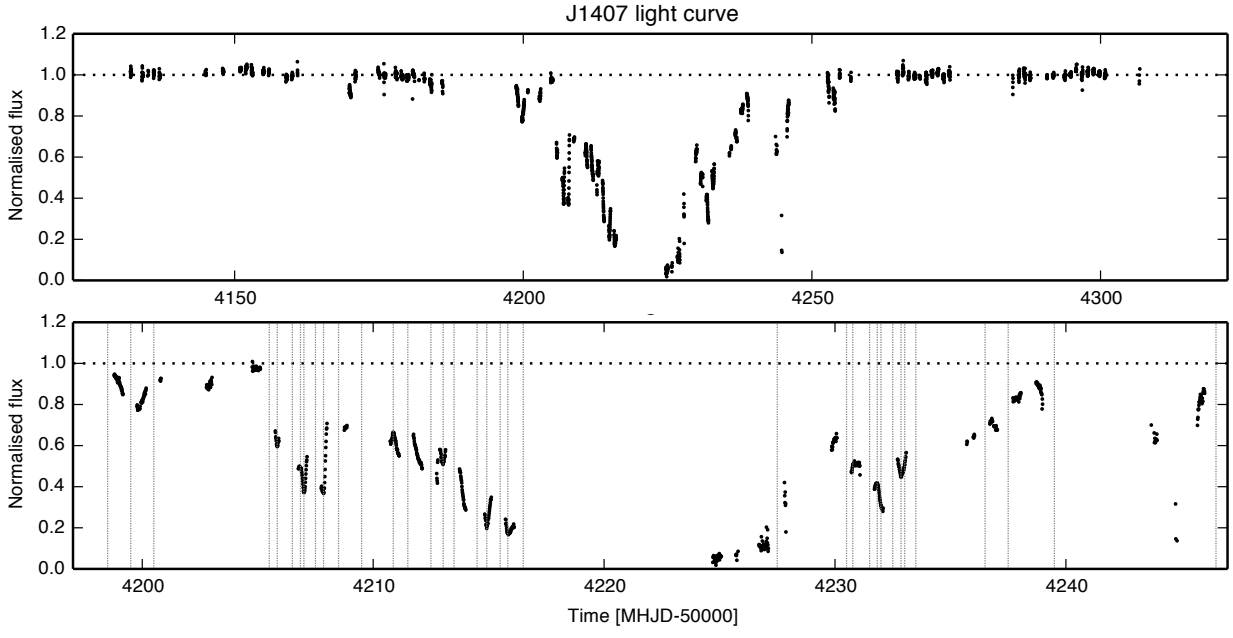


Figure 6.5: J1407 transit light curve as recorded by *SuperWASP* in 2007, corrected for systematics and stellar variability, binned in 15 min bins for clarity. Data are taken during the local night, daytime shows up as gaps. The eclipse midpoint is at approximately MHJD 54 222 and lasts at least 54 d with roughly symmetrical ingress and egress but asymmetry in the nightly variation. *Top:* overview of all of 2007. *Bottom:* zoom in on fine-structure during transit, with inflection points marked with vertical dotted lines.

After successfully decorrelating the signal from the systematics, we then remove the stellar variability. Mamajek et al. (2012) find that the J1407 light curve shows a variability of approximately 3.21 d in the *SuperWASP* data, as expected from rotational modulation of star spots and a young star of ∼ 16 Myr. Since this variability is independent of the transit signal, removing this signal is required for subsequent analysis.

First we investigate the stellar variability for each of the three seasons, 2006, 2007, and 2008. We use phase dispersion minimisation with a variable-period sinusoid as a model and use the Nelder-Mead fitting algorithm (Stellingwerf, 1978; Nelder & Mead, 1965). Our analysis shows that the best fitting periods and their amplitudes differ significantly from year to year (see Table 6.1). Furthermore, the amplitude of the variability is significantly less during the eclipse year 2007 than in the other years.

Initially we fit the variability for the 2006 and 2008 data with a constant period and interpolated this to the 2007 data. However, this interpolated fit mismatches the 2007 data in both phase and amplitude, it appears the period is not constant enough for such interpolation. To mitigate this, we fit the stellar variability using the light curve just before and after the transit in 2007, i.e. before MHJD 54 197 and after 54 255 where we verified by eye that little to no transit signal is present, by plotting the wrapped light curve and increasing the window size until a suitable window was found. Although additionally fitting the 2006 and 2008 data better constrains this variability, this requires modelling of stellar variability outside the scope of this paper.

For the out of transit 2007 data, we find a period of 3.2011 ± 0.0002 d. The phase folded light curve for this stellar variability period is shown in Fig. 6.4. Using the best fitting sinusoid we detrend the 2007 signal of interest as to remove the stellar variability.

The de-correlated and de-trended light curve for 2007 is depicted in Fig. 6.5, the top panel shows the whole year while the bottom panel shows 25 d around the approximate eclipse mid point MHJD 54 222 Ingress and egress show a similar trend of linearly decreasing opacity away from the eclipse mid point, but the nightly structure is very different from night to night.

6.3 Revised Stellar Parameters

Mamajek et al. (2012) presented estimated stellar parameters for J1407. Here we present updated stellar parameters based on slightly revised astrometry and assumed intrinsic stellar colors and temperature.

Using the revised proper motion for J1407 from the UCAC4 catalog (Zacharias et al., 2013, ; UCAC4 252-062736) and the space motion and convergent point solution for the Upper Centaurus-Lupus (UCL) subgroup from Chen et al. (2011), we estimate a revised (predicted) kinematic distance of 133 ±12 pc (ϖ = 7.49 ± 0.65 mas). This is consistent with the mean distance to the UCL subgroup from de Zeeuw et al. (1999): 140 pc.

Pecaut & Mamajek (2013) presented a revised sequence of intrinsic optical-infrared colors, effective temperature ($T_{\rm eff}$), and bolometric corrections as a function of spectral type for pre-main sequence (pre-MS) stars. Adopting the spectral type of K5±1 from Mamajek et al. (2012), and the intrinsic properties of K5 pre-MS stars from Pecaut & Mamajek (2013), we estimate slightly revised stellar parameters: $T_{\rm eff}$ = 4140^{+190}_{-120} K and $BC_J \simeq 1.58 \pm 0.06$ mag. Using the revised intrinsic colors for K5 pre-MS stars, and the $BVJHK_s$ photometry for J1407 compiled in Mamajek et al. (2012), we estimate that the star is statistically consistent with being unreddened, but adopt extinction A_J = 0.02 ± 0.02 (A_V = 0.06 ± 0.06 mag). Combining the revised extinction, parallax, and bolometric magnitude with the 2MASS J magnitude, we estimate the following parameters: $f_{\rm bol}$ = (6.09 ± 0.38) ~ 10^{-13} W m^{-2}, $m_{\rm bol}$ = 11.56 ± 0.07 mag, $M_{\rm bol}$ = 5.93 ± 0.20 mag, log(L/L$_\odot$) = -0.47 ± 0.08 dex. The slight change to the adopted $T_{\rm eff}$, and negligible change to the luminosity, shifts the stellar radius slightly: 1.13 ± 0.14 R$_\odot$. In Table 6.3, we summarize the inferred isochronal ages and masses using grids of pre-MS modern evolutionary tracks and isochrones. The median mass (~0.9 M$_\odot$) has changed negligibly, however the cooler $T_{\rm eff}$ has moved the star from the radiative track to the Hayashi track, and produced correspondingly younger isochronal ages (~10 ± 3 Myr). The isochronal ages of pre-MS K stars appear to be underestimated by nearly a factor of two compared to ages derived from the main-sequence turnoff, main-sequence turn-on, and Li depletion boundary ages (c.f. Pecaut et al., 2012; Bell et al., 2013; Soderblom et al., 2014). We continue to adopt the mean UCL age from Pecaut et al. (2012) (~16 Myr) as representative for J1407, but note that an age somewhere in the range of ~5-20 Myr is possible. We adopt a mass of 0.9 ± 0.1 M$_\odot$ based on the HR diagram position of J1407, noting that four of the five sets of pre-MS evolutionary tracks lead to inferred masses closely bracketing ~0.9 M$_\odot$. However, we note that if the inferred pre-MS isochronal ages may be in error by nearly a factor of ~2×, then the true masses may be systematically off by tens of percent as well (see e.g. Hillenbrand & White, 2004). At this point, we are not in a position to be able to quantify the magnitude of any potential systematic error in mass due to problems with the evolutionary tracks themselves.

In summary, the most significant change to the adopted parameters for

Table 6.2: Revised Stellar Parameters.

(1) Propery	(2) Value	(3) Ref.
α(J2000)	14:07:47.929	1
δ(J2000)	-39:45:42.75	1
μ_α	-24.3 $\pm$ 1.7 mas yr^{-1}	1
μ_δ	-19.4 $\pm$ 2.0 mas yr^{-1}	1
A_J	0.02 $\pm$ 0.02 mag	2
A_V	0.06 $\pm$ 0.06 mag	2
Dist	133 $\pm$ 12 pc	2
$T_{\rm eff}$	4140^{+190}_{-120} K	2
log(L/L$_\odot$)	-0.47 $\pm$ 0.08 dex	2
Age	$\sim$16 Myr	3
Mass	0.9 $\pm$ 0.1 M$_\odot$	2

References: (1) Zacharias et al. (2013), (2) this paper, (3) adopted UCL age from Pecaut et al. (2012), see discussion in text.

J1407 is that the estimate radius is about 18% larger compared to Mamajek et al. (2012). We adopt these new values as compiled in Table 6.2 in the following sections.

6.4 Ring model

After successfully removing instrument systematics and stellar variability, we investigate the fine structure in the light curve. Here we quantitatively investigate this to constrain the geometrical parameters of this system, by modelling the changes in intensity as being due to the passage of a ring edge across the disk of the central star. By measuring the gradient in the light curve and assuming a diameter for the star, we can derive a lower limit to the speed at which the occulting object is moving.

We model the eclipse geometry by assuming that an optically thick ring moves in front of the stellar disk. The stellar disk is approximated with a limb darkening model from Claret & Bloemen (2011b). We assume that the ring is circular and azimuthally symmetric around J1407b, that the diameter of the ring is much larger than the diameter of the star so that the edge of the ring can be approximated as a straight line, and that the edge of the ring is perpendicular to its direction of motion, as seen in Fig. 6.6. The observed light curve gradient dL/dt is given by the rate at which the sur-

Table 6.3: Isochronal Age and Mass Estimates for J1407.

(1) Tracks ...	(2) Age Myr	(3) Mass $M_\odot$
Baraffe et al. (1998)	10	0.90
Siess et al. (2000)	12	0.88
Yi et al. (2003)	5	0.68
Dotter et al. (2008)	12	0.92
Paxton et al. (2011)[1]	9	0.87

[1] Evolutionary tracks and isochrones were generated using the MESA evolutionary code (Paxton et al., 2011), adopting the protosolar abundances from Asplund et al. (2009) (J. Strumpf, priv. comm.).

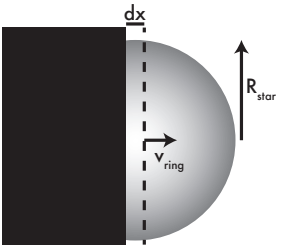


Figure 6.6: Eclipse geometry for an opaque ring crossing the star. In this scenario, no material is in front of the star before the ring arrives.

face of the star is eclipsed, multiplied with the mean intensity of the star in the eclipsed area. Below we calculate light curve gradients for a limb darkened star. These approximations then present a lower limit for the relative velocity of the ring system - if the ring edge is not perpendicular to the direction of motion, then the ring velocity will be higher for a given dL/dt. To explain the data, we require gradients in the light curve of up to 3 $L_\star$ d^{-1}.

For our analysis, we use a linear limb darkening law from Claret & Bloemen (2011b):

$$I(\mu) = 1 - u(1 - \mu), \tag{6.5}$$

with u is the limb darkening coefficient, $\mu = \cos(\gamma)$, where γ is the angle

between the line of sight and the surface normal. For stars such as J1407 ($T_{\rm eff} \sim 4500$ K, $\log g \sim 4.4$, [Fe/H] ~ 0), Claret & Bloemen (2011b) find limb darkening coefficient $u = 0.80 \pm 0.05$ in the *SuperWASP* band. We rewrite μ as linear coordinate $r = \sqrt{1-\mu^2}$ from 0 at the centre of the disk to 1 at the limb to yield:

$$I(r) = 1 - u(1 - \sqrt{1-r^2}). \tag{6.6}$$

The flux in a small vertical strip of width dx at the centre of the stellar disk is then given by

$$\begin{aligned} F_{\rm strip} &= dx \int_{-R}^{R} I(r/R)\, dr \\ &= dx\, 2\, R(1 - u + \frac{\pi u}{4}), \end{aligned} \tag{6.7}$$

while the total flux is given by

$$\begin{aligned} F_{\rm total} &= \int_{0}^{R} 2\pi r\, I(r/R)\, dr \\ &= \frac{1}{6} R^2 \pi\, (6 - 2u), \end{aligned} \tag{6.8}$$

with R the stellar radius. The maximum light curve gradient is then given by the ratio of these, divided by the time dt it takes to cross distance dx, multiplied by the total stellar luminosity

$$\frac{dL}{dt}_{\rm max} = \mathrm{L}_\star \frac{F_{\rm strip}}{F_{\rm total}} \frac{1}{dt}. \tag{6.9}$$

We can rewrite this to obtain the implied speed for a given gradient as

$$\begin{aligned} \frac{dx}{dt}_{\rm min} &= \dot{L}\, R\, \pi \frac{2u - 6}{12 - 12u + 3\pi u} \\ &\approx 13\ \mathrm{km\ s^{-1}} \left(\frac{\dot{L}}{\mathrm{L_\star\ d^{-1}}}\right)\left(\frac{R}{1.13\ \mathrm{R_\odot}}\right) \end{aligned} \tag{6.10}$$

where $v = dx/dt$ is the speed of the ring with respect to the star.

We now assume that J1407b is in a circular Kepler orbit around the central star with velocity v and semi-major axis a:

$$\begin{aligned} a &= G\, M / v^2 \\ &\approx 8.0\ \mathrm{AU} \left(\frac{M}{0.9\ \mathrm{M_\odot}}\right)\left(\frac{v}{10\ \mathrm{km\ s^{-1}}}\right)^{-2}, \end{aligned} \tag{6.11}$$

and a disk diameter of

$$\begin{aligned} r_d &= v\, t_{\text{eclipse}}/2 \\ &\approx 0.16\ \text{AU} \left(\frac{t_{\text{eclipse}}}{54\ \text{d}}\right)\left(\frac{v}{10\ \text{km s}^{-1}}\right). \end{aligned} \tag{6.12}$$

Equation (6.12) is equivalent to Eqn. 10 in Mamajek et al. (2012), although the authors missed a factor π in their equation. We can express the radius of the disk as a fraction of the Hill radius $r_{\text{H}} \equiv a(\mu/3)^{1/3}$, with $\mu = m_2/(m_1 + m_2)$ as

$$\begin{aligned} \frac{r_d}{r_{\text{H}}} \equiv \xi &= t_{\text{eclipse}}\, v\, \frac{3^{1/3}}{a\, \mu^{1/3}} \\ &= 0.28 \left(\frac{t_{\text{eclipse}}}{54\ \text{d}}\right)\left(\frac{v}{10\ \text{km s}^{-1}}\right)^3 \left(\frac{M}{0.9\ \text{M}_\odot}\right)^{-2/3} \left(\frac{m_2}{\text{M}_\text{J}}\right)^{-1/3}. \end{aligned} \tag{6.13}$$

To measure the light curve gradient for J1407, we fit straight lines to manually selected windows with constant gradient. An overview of the light curve is shown in Fig. 6.7, where straight line fits are plotted over the data. The top panel shows the light curve during the transit and each panel below shows one night of data. We plot the measured light curve gradients in Fig. 6.8.

6.5 Results

A wealth of fine structure becomes visible in the light curve (Fig. 6.5 and Fig. 6.7). By decorrelating against reference stars we ensure these are not randomly occurring systematic effects, but are due to J1407. Removing the stellar variability further improves the photometry.

The light curve shows strong gradients of up to 3 $\text{L}_\star\ \text{d}^{-1}$, which means that the eclipsing agent has a speed to occult the star in roughly 1/3 of a day. Furthermore, we find a myriad of fine structures in the light curve, with at least 24 inflection points on ingress, and 16 on egress (see Table 6.4 and Fig. 6.7), and changes of at 15 per cent or more in flux in 6 nights. Some inflection points are directly observed in the nightly data, whilst others are implied by a mismatch in photometric observations between consecutive observation windows, implying an unseen change in light curve gradient. These inflection points indicate the presence of at least 24 different rings in the J1407b system. This ring structure is tentative indication of the presence of shepherding moons, without which the ring structure would be smeared out radially on short dynamical timescales.

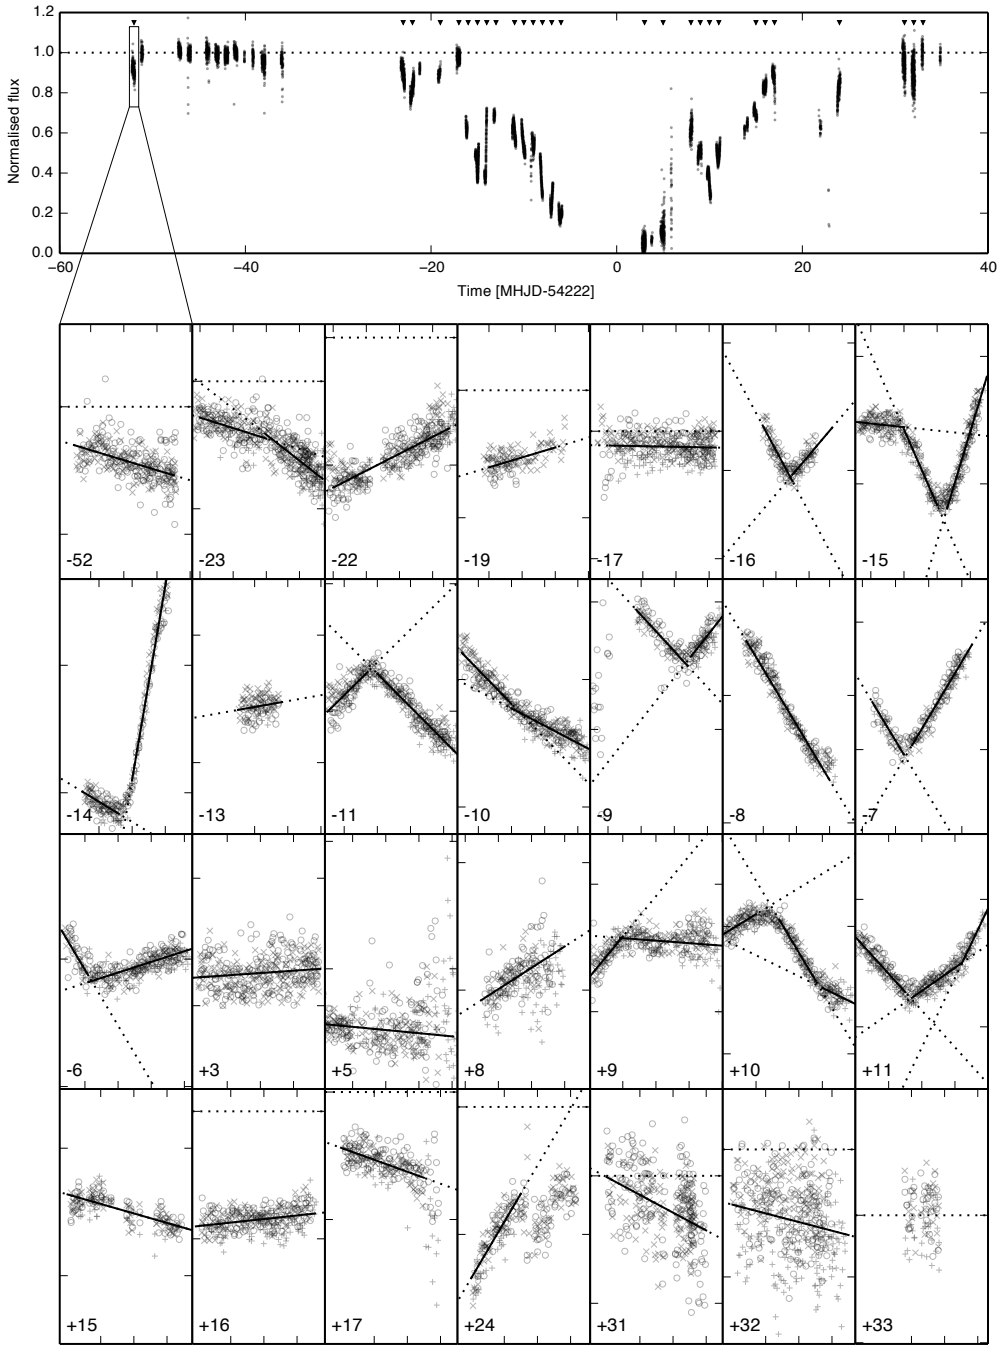


Figure 6.7: Detailed plot of the J1407 light curve during the transit. *Top*: overview of the light curve. *Bottom*: nightly light curve for 28 nights during the transit, indicated in the upper panel by triangles. The straight line fits are plotted over the data, a solid line is the window used for the fit, a dotted line is extrapolated. The y-ticks are spaced at 0.2 AU (as in the top panel) and x-ticks at 0.1 d.

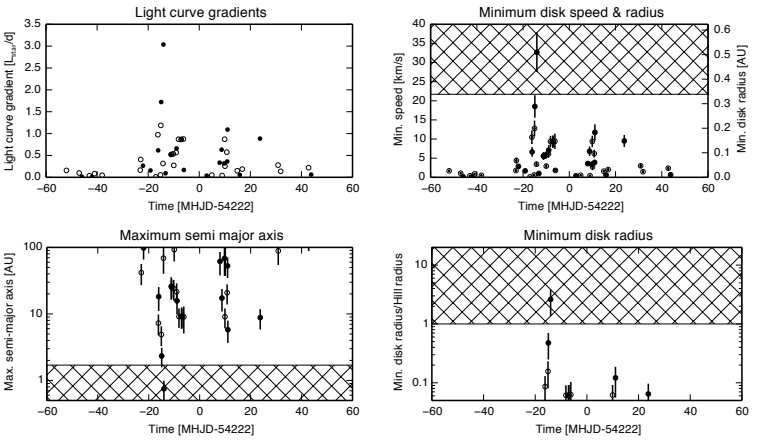


Figure 6.8: Light curve gradients and derived system parameters for circular orbits. *Top-left*: light curve gradients during the transit, *top-right*: disk speeds and radii, *bottom-left*: semi major axes, *bottom-right*: disk radius over Hill radius, assuming $m_2 = 50\,\mathrm{M_J}$. The transit mid point is approximately at MHJD 54 222. Filled circles are positive gradients (i.e. increasing flux), open circles represent negative gradients. The hatched area represents data incompatible with the minimum period of 2.3 yr (lower-left and top-right) or data that predict disks larger than the Hill radius (lower-right).

Given the light curve gradients we measure and an estimate of the diameter of the star, we can compute the physical speed and diameter of the ring system. Using $R = 1.13 \pm 0.14\,\mathrm{R_\odot}$, $M = 0.9 \pm 0.1\,\mathrm{M_\odot}$, $t_{\mathrm{eclipse}} = 54\,\mathrm{d}$ (Mamajek et al., 2012), and $u = 0.80 \pm 0.05$ (Claret & Bloemen, 2011b) gives the results shown in Fig. 6.8. The error bars are dominated by the 16 per cent error in the radius, an improved estimate for this would better constrain the disk parameters.

Most nights have $<15\,\mathrm{km\,s^{-1}}$ occultation speeds, but two nights during ingress have exceptionally high speeds. These gradients are 1.7 and 3.0 $\mathrm{L_\star\,d^{-1}}$, which correspond to 22 ± 3 to $38 \pm 5\,\mathrm{km\,s^{-1}}$, respectively. While the first night is compatible with the $21.7\,\mathrm{km\,s^{-1}}$ upper limit found by Mamajek et al. (2012), the second night is not. Additionally, the highest speeds combined with the 54 d transit time yield a disk radius of 0.3 AU and 0.6 AU, respectively.

Furthermore, we observe significant change in gradients on the timescale of hours in the data, for example during nights –16, –15, –14, –11, 10 and 11 in Fig. 6.7, showing inflection points where the gradient changes. This ratio of the variability timescale of an hour over half the transit duration of 54 d

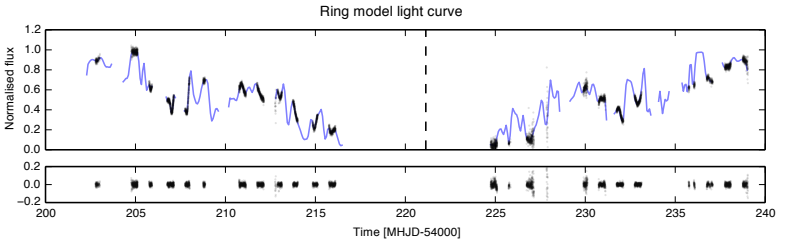


Figure 6.9: Example ring system model light curve corresponding to the geometry shown in Fig. 6.10. The top panel shows the data (black dots) with the model (solid line), the bottom panel shows the residuals of this fit. With 150 rings, the light curve variability can be accounted for, including the steep gradients in some nights. The dashed vertical line indicates the center of the ring system.

implies an aspect ratio of the disk of $h/r \approx 0.0015$, confirming earlier findings by Mamajek et al. (2012).

6.6 Initial ring fitting

After removing instrument systematics and de-trending the signal from stellar variability, we fit a toy ring model to the light curve. The ring system is modeled as a infinitely thin, filled disk with radius $R_{\rm disk}$ consisting of N, equal-width rings each with linear transmission T between 0 and 1. This ring system is rotated under inclination i and azimuth ϕ to match the unknown geometry of the target system. The center of the ring system is located at t_0, the star with diameter d_s impacts at b. The stellar diameter and impact parameter are expressed in days to match the observational data. A best-fitting ring model (detailed below) is shown in Fig. 6.10.

From this ring model we crop out a horizontal section across the full width of the system with the height of the stellar diameter d_s, which in these units is on the order of 0.2 d (see Fig. 6.10). We digitize this section of the model at a resolution of 50 pix d^{-1} and set the pixel values to the transmission T of the ring at that position. We then compute the simulated light curve by convolving the horizontal section with a stellar profile with unit integrated intensity.

A model with fixed radius and number of rings is then fit to the data, using a Nelder-Mead algorithm for all parameters except the ring opacities which are fit using least squares minimisation. The fitting is nested this way because the ring transmissions have a clear optimum value, as opposed to the system orientation. Using least-squares fit where possible

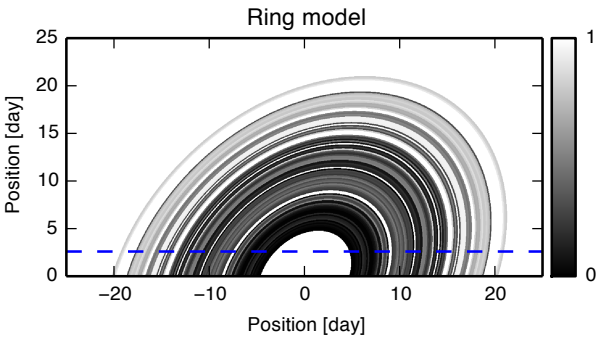


Figure 6.10: Example ring system with 150 rings of variable transmission (indicated by the colourscale) extending up to 25 d out on either side of the eclipse mid point rotated under inclination $i = 0.83$ rad and azimuth $\phi = 2.4$ rad with stellar impact parameter $b = 2.6$ d. The stellar path across the ring system is indicated by the horizontal line. Note that this model is one possible fit, but is by no means unique. The associated light curve is shown in Fig. 6.9.

speeds up the fitting significantly, allowing more rings to be fit in the same time.

We find that using 150 rings for the 25 days around the eclipse mid point gives a satisfactory of the light curve up to the hourly variability, as shown in Fig. 6.10. This best-fitting model is by no means unique, this is only one possible ring model. The sparsity of the data makes ring matching between ingress and egress difficult. Nonetheless, this fit indicates that a ring model is compatible with the data. Future work should include model constraints imposed by resonant gaps introduced by moons as well as a more thorough analysis of the parameter space.

6.7 Discussion

The minimum speeds derived only from the light curve gradients, stellar radius and limb darkening are incompatible with both the minimum period of 2.3 yr found by Mamajek et al. (2012), in spite of the large uncertainties on stellar properties. Also using the stellar mass and a secondary mass upper limit of 50 M_J (Kenworthy, priv. comm.), the disk is also overfilling the Hill radius for the largest minimum speed. Here we briefly consider two scenarios that could mitigate these problems, the first being an elliptical orbit and the second substructure in the rings.

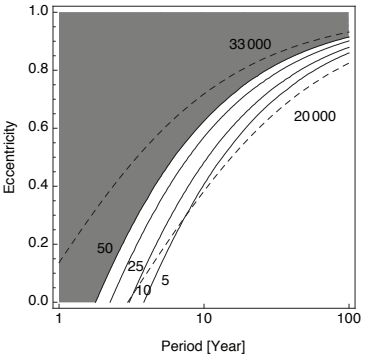


Figure 6.11: Secondary mass for $R_{\rm disk} = R_{\rm Hill}$ in M_J (solid lines) and maximum orbital speed in m s^{-1} (dashed lines), both as function of period and eccentricity. For a given secondary mass (solid line), the area below the line is allowed: in those cases the eccentricity is less, giving a larger minimum separation and thus a larger Hill radius, thus fitting in the disk. For a given speed (dashed line), the area above the curve is allowed: in those cases the eccentricity is higher and thus the maximum speed is larger. For the highest speed calculated from the light curve gradients there is no region compatible with stable rings given that the secondary mass is limited to 50 M_J. The second highest speed could be explained by ellipticity with a <3 yr orbit, or with a secondary mass of 25 M_J above that.

Ellipsoidal orbits for J1407b

An ellipsoidal orbit would allow for larger periastron speeds and explain the anomalous velocities implied in the data, however these are not fast enough to explain the observations. For a Kepler orbit with period P and ellipticity e, we compute the maximum speed at periastron

$$v_{\rm max} = \left(\frac{2\pi\, G\, M}{P}\right)^{1/3} \left(\frac{1+e}{1-e}\right)^{1/2}. \tag{6.14}$$

These speeds are plotted as dashed lines in Fig. 6.11. This maximum speed combined with the transit duration of 54 d gives a disk radius. Additionally, for a planetary mass m_2 we can compute the Hill radius at the periastron

$$R_{\rm Hill} = \left(\frac{m_2\, G\, P^2}{12\pi^2}\right)^{1/3} (1-e). \tag{6.15}$$

We then solve $R_{\rm disk} = R_{\rm Hill}$ using Eqs. (6.12)–(6.15) for m_2, plotted as solid lines in Fig. 6.11. Since the secondary mass is less than 50 M_J (Kenworthy, priv. comm.), ellipticity alone cannot explain the speeds derived from the observed gradients.

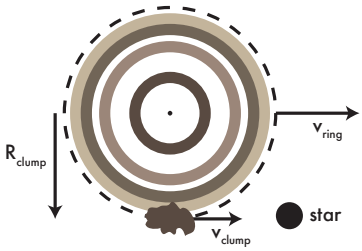


Figure 6.12: Geometry of a clump in a circular orbit around J1407b eclipsing the star. In this case, the speed of the disk with respect to the star ($v_{\rm ring}$) vectorially adds to the speed of the clump in the ring ($v_{\rm clump}$). Outer rings are not shown for clarity.

Azimuthal structure in the rings

If there is non-radial structure in the ring, e.g. a bar or clumps in the ring, this could give rise to an increased gradient dL/dt. Such radial structures ('spokes') have been observed in the B-rings of Saturn (Smith et al., 1982). With azimuthal density structures in the ring, these could vectorially add to the observed light curve gradient, as shown in Fig. 6.12. The night showing the highest light curve gradient is at approximately 14 days before the eclipse mid point. Assuming $m_2 = 50\,{\rm M_J}$ and $P = 3.5\,{\rm yr}$ (Kenworthy, priv. comm.), this would give an orbital speed of the planet of $v_{\rm ring} = 19\,{\rm km\,s^{-1}}$ such that the ring distance would be at $19\,{\rm km\,s^{-1}} \times 14\,{\rm d} = 0.15\,{\rm AU}$. For this radius, the ring would rotate at a speed of $v_{\rm clump} = 17\,{\rm km\,s^{-1}}$ which in total would yield a occulting speed of $36\,{\rm km\,s^{-1}}$, compatible with the observations.

Forward scattering

A third option is that the observed gradient is a combination of the gradient from a ring eclipsing edge added with an additional gradient due to forward scattering from the particles making up the ring wall. Such forward scattering is seen and modelled in Fomalhaut's debris disk (Kalas et al., 2005) and in the light curve produced by a disintegrating exoplanet KIC 12557548 b (Rappaport et al., 2012; Brogi et al., 2012). Since we do not know the thickness of the disk and a unique ring geometry, we do not carry out the modeling in this paper.

6.8 Conclusions

We have reduced the *SuperWASP* extracted photometric flux data for J1407 using an ensemble of photometrically quiet nearby reference stars to remove systematic effect in the data, and detrended for stellar variability, showing an improvement over the raw flux as well as the automatically reduced flux pipeline outputs. We find a periodicity of approximately 3.2 d for the different observation seasons, with significant differences between the seasons. We attribute the difference in observed rotational periods due to the change in the mean latitudinal position of star spots on the surface of J1407 and the differential rotation of the stellar atmosphere.

With the new photometry, we have investigated the nightly structure in the light curve during transit and find strong gradients of up to $\sim$3 $L_\star$ per day. Furthermore, we find a myriad of fine structures in the light curve, which we explain with the transit of a ring system orbiting around J1407b and we identify at least 24 different rings. Given that for a disk alone this structure would be radially smeared out by internal disk interactions on short dynamical timescales, this is indicative of the presence of shepherding exomoons.

We constructed a toy ring model to fit the light curve data and find a good fit can be obtained using 150 rings when using the inner 25 d of light curve around the eclipse mid point. Although we stress that this solution is by no means unique, it does indicate that a ring model can be compatible with the data. An improved model is required as well as a more thorough search through the parameter space for better understanding of the system.

Using a simple model of a limb-darkened star, we calculate the minimum transit speed implied by the observed light curve gradients. We find that these minimum speeds are incompatible with the analysis by Mamajek et al. (2012) who find a minimum period of 2.3 yr. Although an elliptic orbit could alleviate this problem to some extent, it cannot completely explain the largest gradients in the light curve. We propose sub-structure in the rings in forms of clumps (appearing as "spokes") as a plausible explanation for the observed gradients.

Further observations to constrain the orbital period and mass of J1407 are required to better understand this system. ALMA is capable of directly detecting rings down to 0.8 M_{moon} masses (as estimated by Mamajek et al. (2012)) in a few hours, and its spatial resolution would enable a unique confirmation of the bound orbit of J1407b around J1407. Several small telescopes are currently monitoring for the beginning of the next eclipse, and an intensive monitoring campaign to cover the duration of the 60 d eclipse

Table 6.4: Inflection points in the J1407 light curve gradient indicative of opacity change. Each inflection point indicates a change in gradient, which in turn implies change in opacity, and thus a new ring beginning to transit the star. Time is shown in MHJD−54 222. The numbers marked with an asterisk are approximate, indicating the presence of an inflexion point during daylight hours but whose presence is implied by the light curves in adjacent photometric measurements. See also Fig. 6.7.

Inflection point [d]				
Ingress			Egress	
−52.5*	−15.15	−9.5*	5.5*	11.02
−51.5*	−15	−8.97	8.5*	11.5*
−23.5*	−14.5*	−8.5*	8.77	14.5*
−22.5*	−14.15	−7.5*	9.5*	15.5*
−21.5*	−13.5*	−7.08	9.83	17.5*
−16.5*	−12.5*	−6.5*	9.99	24.5*
−16.155	−11.13	−6.17	10.5*	31.5*
−15.5*	−10.5*	−5.5*	10.86	32.5*

would give an unprecedented opportunity for the resolved study of a protoplanetary ring system outside our Solar system. Spectroscopy over a broad range of wavelengths can characterise the density and composition of the rings, providing new insight into their formation and origins. The high rotational of the star, combined with high resolution spectroscopy, leads to the exciting possibility of observing the Rossiter-McLaughlin effect (Rossiter, 1924; McLaughlin, 1924) as the ring edges transit the stellar disk and a subsequent increase in spatial resolution in the ring structure. Initial modelling efforts by our group show that a continuous time coverage using multiple telescopes spread over different latitudes will be essential in constraining the system and determining the ring geometry present there.

Acknowledgments

The WASP project is funded and operated by Queen's University Belfast, the Universities of Keele, St. Andrews and Leicester, the Open University, the Isaac Newton Group, the Instituto de Astrofísica de Canarias, the South African Astronomical Observatory and by STFC.

Table 6.5: Measured light curve gradients greater than 0.1 d^{-1} and derived circular orbital speed. Time is shown in MHJD−54 222.

Time [d]	Gradient [d^{-1}]	Speed [km s^{-1}]	Time [d]	Gradient [d^{-1}]	Speed [km s^{-1}]
−52.00	−0.16(2)	−2.0(3)	−7.12	−0.86(7)	−11(2)
−23.10	−0.16(2)	−2.1(4)	−6.96	0.87(2)	11(1)
−22.92	−0.41(3)	−5.2(7)	−6.23	−0.9(1)	−11(2)
−22.02	0.270(9)	3.4(4)	−6.04	0.17(1)	2.1(3)
−19.10	0.16(3)	2.0(4)	7.97	0.33(4)	4.2(7)
−16.19	−0.97(6)	−12(2)	8.75	0.63(6)	8(1)
−16.09	0.62(7)	8(1)	9.75	0.32(5)	4.0(8)
−15.05	−1.19(4)	−15(2)	9.93	−0.87(3)	−11(1)
−14.91	1.72(4)	22(3)	10.05	−0.26(7)	−3(1)
−14.19	−0.32(4)	−4.0(7)	10.77	−0.58(3)	−7(1)
−14.04	3.04(6)	38(5)	10.94	0.36(2)	4.6(6)
−11.19	0.52(5)	7(1)	11.06	1.09(9)	14(2)
−10.97	−0.52(1)	−6.6(8)	14.90	−0.15(1)	−1.9(3)
−10.18	−0.54(3)	−6.8(9)	16.82	−0.19(2)	−2.4(4)
−9.98	−0.27(1)	−3.5(5)	23.77	0.88(5)	11(2)
−9.04	−0.57(3)	−7(1)	30.85	−0.28(3)	−3.5(6)
−8.90	0.66(6)	8(1)	31.88	−0.14(3)	−1.7(4)
−8.12	−0.87(2)	−11(1)	42.80	−0.22(2)	−2.8(4)

Bibliography

Abbe, E. 1883, *J. Royal Microsc. Soc.*, 3, 790 (p. 7).

Airy, G. B. 1834, *Transactions Camb. Philos. Soc.*, 5, 283 (p. 6).

Albert, O., Sherman, L., Mourou, G., Norris, T. B., & Vdovin, G. 2000, *Opt. Lett.*, 25, 52 (p. 61).

Andersen, M., Dahl, J., Liu, Z., & Vandenberghe, L. 2011, in Optimization for machine learning, ed. S. Sra, S. Nowozin, & S. J. Wright (Cambridge, MA: MIT Press), 55 (p. 74).

Andersen, M., Dahl, J., & Vandenberghe, L. 2013, CVXOPT: A Python package for convex optimization, version 1.1.6, Available at http://cvxopt.org (p. 74).

Antonello, J., Verhaegen, M., Fraanje, R., et al. 2012, *J. Opt. Soc. Am. A*, 29, 2428 (pp. 19, 43, 61, 61, 61, 61, 65, 65, 67, 69, 69, 70, 73, 73, 74, 76, 77, 81, 82, 82).

Artal, P., Iglesias, I., López-Gil, N., & Green, D. G. 1995a, *J. Opt. Soc. Am. A*, 12, 2358 (p. 47).

Artal, P., Marcos, S., Navarro, R., & Williams, D. R. 1995b, *J. Opt. Soc. Am. A*, 12, 195 (pp. 47, 60).

Asplund, M., Grevesse, N., Sauval, A. J., & Scott, P. 2009, *Annu. Rev. Astron. Astrophys.*, 47, 481 (p. 120).

Aviles-Espinosa, R., Andilla, J., Porcar-Guezenec, R., et al. 2011, *Biomed. Opt. Express*, 2, 3135 (p. 60).

Azucena, O., Crest, J., Cao, J., et al. 2010, *Opt. Express*, 18, 17521 (pp. 13, 36, 60).

Azucena, O., Crest, J., Kotadia, S., et al. 2011, *Opt. Lett.*, 36, 825 (pp. 13, 36).

Babcock, H. W. 1953, *Publ. Astron. Soc. Pac.*, 65, 229 (p. 4).

Baglin, A., Auvergne, M., Barge, P., et al. 2006, in ESA Special Publication, Vol. 1306, The CoRoT Mission Pre-Launch Status - Stellar Seismology and Planet Finding, ed. M. Fridlund, A. Baglin, J. Lochard, & L. Conroy, 33 (p. 24).

Baraffe, I., Chabrier, G., Allard, F., & Hauschildt, P. H. 1998, *Astron. Astrophys.*, 337, 403 (p. 120).

Bardell, D. 2004, *BIOS*, 75, 78 (p. 3).

Beck, A., Stoica, P., & Li, J. 2008, *Signal Process. IEEE Transactions on*, 56, 1770 (pp. 61, 67, 67, 68, 68, 68, 69, 69, 69, 69, 82).

Bell, C. P. M., Naylor, T., Mayne, N. J., Jeffries, R. D., & Littlefair, S. P. 2013, *Mon. Notices Royal Astron. Soc.*, 434, 806 (p. 118).

Beuzit, J. L., Feldt, M., Dohlen, K., et al. 2008, in SPIE Proceedings, Vol. 7014, Ground-based and Airborne Instrumentation for Astronomy II, ed. I. S. McLean & M. M. Casali (Marseille, France: SPIE) (pp. 13, 25).

Bifano, T. 2011, *Nat. Photonics*, 5, 21 (p. 14).

Bone, D. J., Bachor, H. A., & Sandeman, R. J. 1986, *Appl. Opt.*, 25, 1653 (p. 37).

Booth, M., Wilson, T., Sun, H.-B., Ota, T., & Kawata, S. 2005, *Appl. Opt.*, 44, 5131 (p. 63).

Booth, M. J. 2006, *Opt. Express*, 14, 1339 (pp. 61, 61, 61, 81).

—. 2007a, *Philos. Transactions Royal Soc. A: Math. Phys. Eng. Sci.*, 365, 2829 (pp. 12, 36).

—. 2007b, *Opt. Lett.*, 32, 5 (pp. 6, 12, 36, 36, 60, 61).

—. 2007c, *Opt. Lett.*, 32, 5 (pp. 61, 61, 81).

Booth, M. J., Débarre, D., & Jesacher, A. 2012, *Opt. Photonics News*, 23, 22 (pp. 12, 13, 14).

Booth, M. J., Neil, M. A. A., & Wilson, T. 2002, *J. Opt. Soc. Am. A*, 19, 2112 (pp. 61, 61, 81).

Booth, M. J., & Wilson, T. 2001, *J. Biomedical Opt.*, 6, 266 (p. 15).

Borucki, W. J., Agol, E., Fressin, F., et al. 2013, *Science*, 340, 587 (p. 24).

Bouchez, A. H., Acton, D. S., Agapito, G., et al. 2012, in SPIE Proceedings, Vol. 8447, Adaptive Optics Systems III, ed. B. L. Ellerbroek, E. Marchetti, & J.-P. Véran, SPIE (Amsterdam, Netherlands: SPIE), 8447II (p. 5).

Brakenhoff, G. J., Visscher, K., & Voort, H. T. M. 1990, in Handbook of Biological Confocal Microscopy, revised edn., ed. J. B. Pawley (New York, USA: Plenum Press), 87 (p. 41).

Brogi, M., Keller, C. U., Ovelar, D. J., et al. 2012, *Astron. Astrophys.*, 545, L5 (pp. 29, 30, 30, 89, 89, 93, 94, 94, 94, 101, 128).

Brown, T. M. 2003, *The Astrophys. J.*, 593, L125 (p. 26).

Budaj, J. 2013, *Astron. Astrophys.*, 557, A72 (pp. 29, 88, 89, 93, 102, 106, 108).

Bueno, J. M., Gualda, E. J., & Artal, P. 2010, *J. Biomedical Opt.*, 15, 066004 (p. 60).

Burrows, C. J., Holtzman, J. A., Faber, S. M., et al. 1991, *The Astrophys. J.*, 369, L21 (p. 7).

Butters, O. W., West, R. G., Anderson, D. R., et al. 2010, *Astron. Astrophys.*, 520, L10 (p. 111).

Cameron, A. C., Bouchy, F., Hébrard, G., et al. 2007, *Mon. Notices Royal Astron. Soc.*, 375, 951 (p. 111).

Campagnola, P. J., Lewis, A., Loew, L. M., Clark, H. A., & Mohler, W. A. 2001, *J. Biomedical Opt.*, 6, 277 (p. 60).

Cha, J. W., Ballesta, J., & So, P. T. C. 2010, *J. Biomedical Opt.*, 15, 046022 (pp. 13, 36, 60).

Chauvin, G., Lagrange, A. M., Dumas, C., et al. 2005, *Astron. Astrophys.*, 438, L25 (p. 25).

Chen, C. H., Mamajek, E. E., Bitner, M. A., et al. 2011, *The Astrophys. J.*, 738, 122 (p. 118).

Claret, A., & Bloemen, S. 2011a, *Astron. Astrophys.*, 529, A75 (p. 103).

—. 2011b, *VizieR Online Data Catalog*, 352, 99075 (pp. 119, 120, 121, 124).

Collier Cameron, A., Wilson, D. M., West, R. G., et al. 2007, *Mon. Notices Royal Astron. Soc.*, 380, 1230 (p. 26).

Crowe, M. J. 1999, The Extraterrestrial Life Debate, 1750–1900 (Courier Dover Publications) (p. 23).

de Grauw, C. J., Vroom, J. M., van der Voort, H. T. M., & Gerritsen, H. C. 1999, *Appl. Opt.*, 38, 5995 (p. 36).

de Zeeuw, P. T., Hoogerwerf, R., de Bruijne, J. H. J., Brown, A. G. A., & Blaauw, A. 1999, *The Astron. J.*, 117, 354 (p. 118).

Débarre, D., Booth, M. J., & Wilson, T. 2007a, *Opt. Express*, 15, 8176 (pp. 12, 20, 36).

—. 2007b, *Opt. Express*, 15, 8176 (pp. 61, 61, 77, 81).

Débarre, D., Botcherby, E. J., Booth, M. J., & Wilson, T. 2008a, *Opt. Express*, 16, 9290 (p. 19).

—. 2008b, *Opt. Express*, 16, 9290 (pp. 61, 61, 61, 62, 66, 81).

Débarre, D., Botcherby, E. J., Watanabe, T., et al. 2009a, *Opt. Lett.*, 34, 2495 (pp. 12, 16, 19, 36, 53).

—. 2009b, *Opt. Lett.*, 34, 2495 (pp. 61, 61, 61, 62, 65, 72, 81).

Débarre, D., Facomprez, A., & Beaurepaire, E. 2012, in SPIE Proceedings, Vol. 8253, MEMS Adaptive Optics VI, ed. S. S. Olivier, T. G. Bifano, & J. Kubby (San Francisco, California, USA: SPIE), 82530F (p. 66).

Deng, S., Liu, L., Cheng, Y., Li, R., & Xu, Z. 2009, *Opt. Express*, 17, 1714 (p. 17).

Denk, W., Strickler, J. H., & Webb, W. W. 1990, *Science*, 248, 73 (pp. 11, 36, 60).

Des Marais, D. J., Harwit, M. O., Jucks, K. W., et al. 2002, *Astrobiology*, 2, 153 (p. 24).

Dotter, A., Chaboyer, B., Jevremović, D., et al. 2008, *The Astrophys. J. Suppl. Ser.*, 178, 89 (p. 120).

Duffner, R. W. 2009, *ITEA J.*, 29, 341 (p. 5).

Eastman, J., Siverd, R., & Gaudi, B. S. 2010, *Publ. Astron. Soc. Pac.*, 122, 935 (p. 93).

Ellerbroek, B. L. 2013, *J. Astrophys. Astron.*, 34, 121 (p. 5).

Facomprez, A., Beaurepaire, E., & Débarre, D. 2012, *Opt. Express*, 20, 2598 (pp. 12, 12, 19, 36, 61, 61, 65, 72, 72, 77, 81, 82, 82).

Feierabend, M. 2004, PhD thesis, Ruperto-Carola University of Heidelberg, Heidelberg, Germany (p. 40, 40).

Feierabend, M., Rückel, M., & Denk, W. 2004, *Opt. Lett.*, 29, 2255 (pp. 13, 18, 36, 40, 44, 44, 45, 55).

Fernandez, E., & Artal, P. 2003, *Opt. Express*, 11, 1056 (p. 63).

Fiolka, R., Si, K., & Cui, M. 2012, *Opt. Express*, 20, 16532 (p. 55).

Foy, R., & Labeyrie, A. 1985, *Astron. & Astrophys.*, 152, L29 (p. 36).

Fraquelli, D., & Thompson, S. E. 2012, Kepler Archive Manual (KDMC-10008-004), Tech. rep., Space Telescope Science Institute (pp. 90, 91, 112).

Fried, D. L. 1965, *J. Opt. Soc. Am.*, 55, 1427 (p. 7).

Friedrich, J., Seidel, C., Ebner, R., & Kunz-Schughart, L. A. 2009, *Nat. Protoc.*, 4, 309 (p. 50).

Fugate, R. Q., Fried, D. L., Ameer, G. A., et al. 1991, *Nature*, 353, 144 (p. 36).

Gelman, A., & Rubin, D. B. 1992, *Stat. Sci.*, 7, 457 (p. 104).

Ghiglia, D. C., & Pritt, M. D. 1998, Two-Dimensional Phase Unwrapping: Theory, Algorithms, and Software, 1st edn. (Wiley-Interscience) (p. 44).

Gilliland, R. L., Jenkins, J. M., Borucki, W. J., et al. 2010, *The Astrophys. J.*, 713, L160 (p. 90).

Goodman, J. W. 1976, *J. Opt. Soc. Am.*, 66, 1145 (p. 44).

Gould, T. J., Burke, D., Bewersdorf, J., & Booth, M. J. 2012, *Opt. Express*, 20, 20998 (p. 17).

Greenwood, D. P. 1977, *J. Opt. Soc. Am.*, 67, 390 (p. 7).

Guinan, E. F., & Dewarf, L. E. 2002, in ASP Conference Proceedings, Vol. 279, Exotic Stars as Challenges to Evolution, ed. C. A. Tout & W. van Hamme (San Francisco: Astronomical Society of the Pacific), 121 (p. 110, 110).

Hansen, P. C. 2010, Discrete inverse problems: insight and algorithms, Vol. 7 (SIAM) (p. 65).

Hartmann, J. 1900, *Zeitschrift Für Instrumentenkunde*, 20, 17 (p. 8).

Hecht, E. 2002, Optics, 4th edn. (Addison-Wesley) (p. 7).

Hell, S., Reiner, G., Cremer, C., & Stelzer, E. H. K. 1993, *J. Microsc.*, 169, 391 (p. 15).

Hell, S. W. 2003, *Nat. Biotechnol.*, 21, 1347 (p. 17, 17).

Hell, S. W., & Wichmann, J. 1994, *Opt. Lett.*, 19, 780 (p. 17).

Helmholtz, & Fripp, H. 1876, *The Mon. Microsc. J.*, 16, 15 (p. 7).

Hermann, B., Fernández, E. J., Unterhuber, A., et al. 2004, *Opt. Lett.*, 29, 2142 (p. 6).

Hillenbrand, L. A., & White, R. J. 2004, *The Astrophys. J.*, 604, 741 (p. 118).

Hinkley, S., Oppenheimer, B. R., Brenner, D., et al. 2008, in SPIE Proceedings, Vol. 7015, Adaptive Optics Systems, ed. N. Hubin, C. E. Max, & P. L. Wizinowich, The International Society for Optical Engineering (Marseille, France: SPIE) (p. 13).

Hodapp, K. W., Suzuki, R., Tamura, M., et al. 2008, in SPIE Proceedings, Vol. 7014, Ground-based and Airborne Instrumentation for Astronomy II, ed. I. S. McLean & M. M. Casali (Marseille, France: SPIE) (pp. 5, 13).

Hough, G. W. 1885, *Observatory*, 8, 275 (p. 4).

Huang, D., Swanson, E. A., Lin, C. P., et al. 1991, *Science*, 254, 1178 (p. 36).

Itoh, K. 1982, *Appl. Opt.*, 21, 2470 (p. 44).

Jenkins, J. M., Caldwell, D. A., Chandrasekaran, H., et al. 2010a, *The Astrophys. J.*, 713, L120 (p. 90).

—. 2010b, *The Astrophys. J.*, 713, L87 (p. 90).

Jesacher, A., Marshall, G. D., Wilson, T., & Booth, M. J. 2010, *Opt. Express*, 18, 656 (p. 6).

Jesacher, A., Thayil, A., Grieve, K., et al. 2009, *Opt. Lett.*, 34, 3154 (pp. 16, 61, 61, 69, 72, 81).

Ji, N., Milkie, D. E., & Betzig, E. 2009, *Nat. Methods*, 7, 141 (pp. 13, 60).

Ji, N., Sato, T. R., & Betzig, E. 2011, *Proc. Natl. Acad. Sci.*, 109, 22 (p. 13).

Kalas, P., Graham, J. R., & Clampin, M. 2005, *Nature*, 435, 1067 (p. 128).

Kalas, P., Graham, J. R., Chiang, E., et al. 2008, *Science*, 322, 1345 (pp. 3, 25).

Kane, S. R., Cameron, A. C., Horne, K., et al. 2004, *Mon. Notices Royal Astron. Soc.*, 353, 689 (p. 111).

Kasper, M., Beuzit, J.-L., Verinaud, C., et al. 2010, in SPIE Proceedings, Vol. 7735, Ground-based and Airborne Instrumentation for Astronomy III, ed. I. S. McLean, S. K. Ramsay, & H. Takami, SPIE (San Diego, California: SPIE), 77352E (p. 9).

Kawahara, H., Hirano, T., Kurosaki, K., Ito, Y., & Ikoma, M. 2013, *The Astrophys. J.*, 776, L6 (pp. 93, 98).

Koch, D. G., Borucki, W. J., Webster, L., et al. 1998, in SPIE Proceedings, Vol. 3356, Space Telescopes and Instruments V, ed. P. Y. Bely & J. B. Breckinridge (Kona, HI: SPIE), 599 (p. 24).

Koch, D. G., Borucki, W. J., Basri, G., et al. 2010, *The Astrophys. J.*, 713, L79 (p. 89).

Kolmogorov, A. N. 1941a, *Akademiia Nauk SSSR Doklady*, 32, 16, translated into English by V. Levin in Kolmogorov (1991a). (pp. 7, 138).

—. 1941b, *Akademiia Nauk SSSR Doklady*, 30, 301, translated into English by V. Levin in Kolmogorov (1991b). (pp. 7, 138).

—. 1991a, *Royal Soc. Lond. Proc. Ser. A*, 434, 15, english translation of Kolmogorov (1941a) by V. Levin. (p. 138).

—. 1991b, *Royal Soc. Lond. Proc. Ser. A*, 434, 9, english translation of Kolmogorov (1941b) by V. Levin. (p. 138).

Korff, D., Dryden, G., & Leavitt, R. P. 1975, *J. Opt. Soc. Am.*, 65, 1321 (p. 7).

Kriss, T. C., & Kriss, V. M. 1998, *Neurosurgery*, 42, 899 (p. 3).

Kromann, E. B., Gould, T. J., Juette, M. F., Wilhjelm, J. E., & Bewersdorf, J. 2012, *Opt. Lett.*, 37, 1805 (p. 17).

Lagarias, J., Reeds, J., Wright, M., & Wright, P. 1998, *SIAM J. on Optim.*, 9, 112 (p. 72).

Langley, S. P. 1903, *Observatory*, 26, 249 (p. 4).

Layard, A. H. 1853, Discoveries in the Ruins of Nineveh and Babylon; With Travels in Armenia, Kurdistan and the Desert (London: John Murray), 197 (p. 3).

Lenzen, R., Hartung, M., Brandner, W., et al. 2003, in SPIE Proceedings, Vol. 4841, Instrument Design and Performance for Optical/Infrared Ground-based Telescopes, ed. M. Iye & A. F. M. Moorwood (Waikoloa, HI: SPIE), 944 (p. 5).

Liang, J., Williams, D. R., & Miller, D. T. 1997, *J. Opt. Soc. Am. A*, 14, 2884 (p. 6).

Linhai, H., & Rao, C. 2011, *Opt. Express*, 19, 371 (pp. 61, 81).

Linnik, V. P. 1957, *Opt. Spectrosc.*, 3, 401, translated into English by M. Slade in Linnik (1994). (pp. 4, 139).

Linnik, V. P. 1994, in ESO Conference and Workshop Proceedings, Vol. 48, Active and adaptive optics, ed. F. Merkle, ESO (Garching: ESO), 535, english translation of Linnik (1957) by M. Slade and edited by Jacques M. Beckers. (p. 139).

Macintosh, B. A., Graham, J. R., Palmer, D. W., et al. 2008, in SPIE Proceedings, Vol. 7015, Adaptive Optics Systems, ed. N. Hubin, C. E. Max, & P. L. Wizinowich (Marseille, France: SPIE) (pp. 5, 13, 25).

Macy, W. W. 1983, *Appl. Opt.*, 22, 3898 (p. 37).

Mamajek, E. E., Quillen, A. C., Pecaut, M. J., et al. 2012, *The Astron. J.*, 143, 72 (pp. 31, 31, 31, 31, 109, 110, 110, 112, 117, 117, 118, 118, 119, 122, 124, 124, 125, 126, 129, 129).

Mandel, K., & Agol, E. 2002, *The Astrophys. J.*, 580, L171 (p. 104).

Marois, C., Macintosh, B., Barman, T., et al. 2008, *Science*, 322, 1348 (pp. 3, 25).

Marsh, P., Burns, D., & Girkin, J. 2003, *Opt. Express*, 11, 1123 (pp. 12, 36, 61).

Martin, J. A., & Roorda, A. 2005, *Ophthalmology*, 112, 2219 (p. 13).

Mauclair, C., Mermillod-Blondin, A., Huot, N., Audouard, E., & Stoian, R. 2008, *Opt. Express*, 16, 5481 (p. 6).

Mayor, M., & Queloz, D. 1995, *Nature*, 378, 355 (p. 23, 23).

Mayor, M., Pepe, F., Queloz, D., et al. 2003, *Messenger*, 114, 20 (p. 23).

McCray, P. 2000, Interview with Dr. Robert Q. Fugate, Interview (p. 4).

McLaughlin, D. B. 1924, *The Astrophys. J.*, 60, 22 (p. 130).

McPherson, A., Gilmozzi, R., Spyromilio, J., Kissler-Patig, M., & Ramsay, S. 2012, *The Messenger*, 148, 2 (p. 5).

Mikolajewski, M., & Graczyk, D. 1999, *Mon. Notices Royal Astron. Soc.*, 303, 521 (p. 110).

Milkie, D. E., Betzig, E., & Ji, N. 2011, *Opt. Lett.*, 36, 4206 (p. 60).

Minsky, M. 1961, Microscopy apparatus, US patent 3,013,467 (pp. 11, 36).

Moré, J. J. 1993, *Optim. Methods Softw.*, 2, 189 (pp. 68, 68, 68, 69, 69).

Murray, L. P., Dainty, J. C., & Daly, E. 2005, in SPIE Proceedings, Vol. 5823, Opto-Ireland 2005: Imaging and Vision, ed. F. D. Murtagh (Dublin, Ireland: SPIE), 40 (p. 61).

Neil, M. A. A., Booth, M. J., & Wilson, T. 2000, *Opt. Lett.*, 25, 1083 (pp. 61, 61, 81).

Nelder, J. A., & Mead, R. 1965, *The Comput. J.*, 7, 308 (pp. 43, 117).

Noll, R. J. 1976, *J. Opt. Soc. Am.*, 66, 207 (pp. 7, 44, 47, 61).

Norton, A. J., Payne, S. G., Evans, T., et al. 2011, *Astron. Astrophys.*, 528, A90 (p. 112).

Olivier, N., Débarre, D., & Beaurepaire, E. 2009, *Opt. Lett.*, 34, 3145 (p. 16).

Ooto, S., Hangai, M., Takayama, K., et al. 2011, *Ophthalmology*, 118, 873 (p. 13).

Ooto, S., Hangai, M., Sakamoto, A., et al. 2010, *Ophthalmology*, 117, 1800 (p. 13).

Oppenheimer, B. R., & Hinkley, S. 2009, *Annu. Rev. Astron. Astrophys.*, 47, 253 (pp. 24, 25, 28).

Paterson, C., Munro, I., & Dainty, J. C. 2000, *Opt. Express*, 6, 175 (pp. 37, 61, 65).

Pawley, J. B., ed. 2006, Handbook Of Biological Confocal Microscopy, 3rd edn. (Boston, MA: Springer US) (pp. 11, 36).

Paxton, B., Bildsten, L., Dotter, A., et al. 2011, *The Astrophys. J. Suppl. Ser.*, 192, 3 (p. 120, 120).

Pecaut, M. J., & Mamajek, E. E. 2013, *The Astrophys. J. Suppl. Ser.*, 208, 9 (p. 118, 118).

Pecaut, M. J., Mamajek, E. E., & Bubar, E. J. 2012, *The Astrophys. J.*, 746, 154 (pp. 118, 118, 119).

Peck, M. 2010, Interferometry mathematics, algorithms and data (pp. 44, 47).

Perez-Becker, D., & Chiang, E. 2013, *Mon. Notices Royal Astron. Soc.*, 433, 2294 (pp. 29, 29, 89, 98, 108).

Pollacco, D. L., Skillen, I., Cameron, A. C., et al. 2006, *Publ. Astron. Soc. Pac.*, 118, 1407 (p. 111).

Potter, A. E., & Morgan, T. H. 1990, *Science*, 248, 835 (p. 98).

Rahman, S. A., & Booth, M. J. 2013, *Appl. Opt.*, 52, 5523 (pp. 13, 21, 39, 45, 46, 47, 55, 57, 60, 60, 60).

Rappaport, S., Levine, A., Chiang, E., et al. 2012, *The Astrophys. J.*, 752, 1 (pp. 29, 88, 88, 93, 95, 98, 128).

Lord Rayleigh, J. W. S. 1902, in Scientific papers, Vol. 3, Scientific papers (Cambridge: Cambridge University Press), 47 (p. 7, 7).

Reed, S. G., Van Atta, R. H., & Deitchman, S. J. 1990, DARPA Technical Accomplishments: An Historical Review of Selected DARPA Projects, Tech. rep., Defense Advanced Research Projects Agency (p. 4).

Rocca, A., Roddier, F., & Vernin, J. 1974, *J. Opt. Soc. Am.*, 64, 1000 (p. 6).

Roorda, A., Romero-Borja, F., Donnelly, W., et al. 2002, *Opt. Express*, 10, 405 (p. 6).

Rossiter, R. A. 1924, *Astrophys. J.*, 60, 15 (p. 130).

Rückel, M., & Denk, W. 2007, *J. Opt. Soc. Am. A*, 24, 3517 (pp. 13, 18, 47).

Rückel, M., Mack-Bucher, J. A., & Denk, W. 2006, *Proc. Natl. Acad. Sci.*, 103, 17137 (pp. 13, 18, 36, 40, 41, 44, 44, 44, 45, 55, 55, 56, 56, 60).

Sackett, P. D. 1999, in NATO Science Series, Vol. 532, Planets Outside the Solar System: Theory and Observations, ed. J. M. Mariotti & D. Alloin (Dordrecht: Springer Netherlands), 189 (p. 26, 26).

Schwertner, M., Booth, M. J., Neil, M. A. A., & Wilson, T. 2004a, *J. Microsc.*, 213, 11 (pp. 14, 36).

Schwertner, M., Booth, M. J., & Wilson, T. 2004b, *Opt. Express*, 12, 6540 (p. 14).

—. 2007, *J. Microsc.*, 228, 97 (pp. 14, 36).

Shack, R. V., & Platt, B. C. 1971, in Spring Meeting of the Optical Society of America, ed. D. S. Chairman, Optical Society of America (Optical Society of America) (p. 8).

Shaw, M., O'Holleran, K., & Paterson, C. 2013, *Opt. Express*, 21, 19353 (p. 60).

Sheppard, C. J. R., & Török, P. 1997, *J. Microsc.*, 185, 366 (p. 15).

Sherman, L., Ye, J. Y., Albert, O., & Norris, T. B. 2002, *J. Microsc.*, 206, 65 (p. 12).

Shotton, D. M. 1989, *J. Cell Sci.*, 94, 175 (pp. 11, 36).

Siess, L., Dufour, E., & Forestini, M. 2000, *Astron. Astrophys.*, 358, 593 (p. 120).

Simmonds, R. D., & Booth, M. J. 2013, *J. Opt.*, 15, 094010 (pp. 14, 15).

Smith, B. A., Soderblom, L., Batson, R., et al. 1982, *Science*, 215, 504 (pp. 31, 128).

So, P. T. C., Dong, C. Y., Masters, B. R., & Berland, K. M. 2000, *Annu. Rev. Biomedical Eng.*, 2, 399 (p. 43).

Soderblom, D. R., Hillenbrand, L. A., Jeffries, R. D., Mamajek, E. E., & Naylor, T. 2014, in Protostars and Planets VI, ed. H. Beuther, R. Klessen, C. Dullemond, & Th (University of Arizona Press) (p. 118).

Song, H., Fraanje, R., Schitter, G., et al. 2010, *Opt. Express*, 18, 24070 (pp. 61, 81).

Soule, H. D., Vazguez, J., Long, A., Albert, S., & Brennan, M. 1973, *J. Natl. Cancer Inst.*, 51, 1409 (p. 50).

Southwell, W. H. 1980, *J. Opt. Soc. Am.*, 70, 998 (p. 44).

Stellingwerf, R. F. 1978, *The Astrophys. J.*, 224, 953 (pp. 93, 117).

Struve, O. 1952, *Observatory*, 72, 199 (p. 25).

Swain, M. R., Vasisht, G., & Tinetti, G. 2008, *Nature*, 452, 329 (p. 28).

Takeda, M. 1990, *Ind. Metrol.*, 1, 79 (pp. 37, 40, 41).

Takeda, M., Ina, H., & Kobayashi, S. 1982, *J. Opt. Soc. Am.*, 72, 156 (pp. 18, 37, 37, 39, 41, 41, 42, 44, 47).

Tamuz, O., Mazeh, T., & Zucker, S. 2005, *Mon. Notices Royal Astron. Soc.*, 356, 1466 (pp. 111, 112).

Tang, J., Germain, R. N., & Cui, M. 2012, *Proc. Natl. Acad. Sci.*, 109, 8434 (p. 12).

Tao, X., Azucena, O., Fu, M., et al. 2011a, *Opt. Lett.*, 36, 3389 (pp. 13, 36).

Tao, X., Crest, J., Kotadia, S., et al. 2012, *Opt. Express*, 20, 15969 (pp. 13, 36).

Tao, X., Dean, Z., Chien, C., et al. 2013a, *Opt. Express*, 21, 31282 (p. 60).

Tao, X., Fernandez, B., Azucena, O., et al. 2011b, *Opt. Lett.*, 36, 1062 (pp. 13, 36).

Tao, X., Norton, A., Kissel, M., Azucena, O., & Kubby, J. 2013b, *Opt. Lett.*, 38, 5075 (pp. 13, 36, 60).

Thayil, A., & Booth, M. 2011, *J. Eur. Opt. Soc. - Rapid publications*, 6 (pp. 61, 62).

Tinetti, G., Beaulieu, J. P., Henning, T., et al. 2012, *Exp. Astron.*, 34, 311 (p. 28).

Török, P., Hewlett, S. J., & Varga, P. 1997, *J. Microsc.*, 188, 158 (p. 15).

Torrieri, D. 1984, *Aerosp. Electron. Syst. IEEE Transactions on*, AES-20, 183 (p. 66).

Trauger, J. T., Ballester, G. E., Burrows, C. J., et al. 1994, *The Astrophys. J.*, 435, L3 (p. 7).

Tuohy, S., & Podoleanu, A. G. 2010, *Opt. Express*, 18, 3458 (pp. 13, 36, 55).

Tyson, R. 2010, Principles of adaptive optics (CRC Press) (p. 60).

Urban, N. T., Willig, K. I., Hell, S. W., & Nägerl, U. V. 2011, *Biophys. J.*, 101, 1277 (p. 17).

van Werkhoven, T., Antonello, J., Truong, H. H., et al. forthcoming, *Opt. Express* (pp. 60, 60, 71, 71).

van Werkhoven, T., Truong, H., Antonello, J., et al. 2012, in SPIE Proceedings, Vol. 8253, MEMS Adaptive Optics VI, ed. S. S. Olivier, T. G. Bifano, & J. Kubby (San Francisco, California, USA: SPIE), 82530E (pp. 60, 71).

Čižmár, T., Mazilu, M., & Dholakia, K. 2010, *Nat. Photonics*, 4, 388 (p. 6).

Vdovin, G., Soloviev, O., Loktev, M., & Patlan, V. 2013, Oko guide to adaptive optics, 4th edn. (Flexible Optical B.V.) (pp. 61, 63, 70, 70).

Vdovin, G. V. 1998, in SPIE Proceedings, Vol. 3353, Adaptive Optical System Technologies, ed. D. Bonaccini & R. K. Tyson (Kona, HI: SPIE), 902 (p. 61).

von Zernike, F. 1934, *Physica*, 1, 689 (pp. 7, 44).

Vorontsov, M. A. 2002, *J. Opt. Soc. Am. A*, 19, 356 (p. 61).

Wang, B., & Booth, M. J. 2009, *Opt. Commun.*, 282, 4467 (p. 63).

Wang, J., & Gh. 2012, *Opt. Lett.*, 37, 4862 (p. 13).

Wang, J., Léger, J.-F., Binding, J., et al. 2012, *Biomedical Opt. Express*, 3, 2510 (pp. 40, 44, 45, 45, 45, 55, 57, 57).

Willig, K. I., Rizzoli, S. O., Westphal, V., Jahn, R., & Hell, S. W. 2006, *Nature*, 440, 935 (p. 3).

Wolfing, J. I., Chung, M., Carroll, J., Roorda, A., & Williams, D. R. 2006, *Ophthalmology*, 113, 1014 (p. 13).

Wolszczan, A., & Frail, D. A. 1992, *Nature*, 355, 145 (p. 23, 23).

Wright, A. J., Burns, D., Patterson, B. A., et al. 2005, *Microsc. Res. Tech.*, 67, 36 (p. 61).

Yi, S. K., Kim, Y.-C., & Demarque, P. 2003, *The Astrophys. J. Suppl. Ser.*, 144, 259 (p. 120).

Yoo, H. W., Verhaegen, M., van Royen, M., & Schitter, G. 2012, in Instrumentation and Measurement Technology Conference (I2MTC), 2012 IEEE International, 1083 (p. 61).

Zacharias, N., Finch, C. T., Girard, T. M., et al. 2013, *The Astron. J.*, 145, 44 (pp. 118, 119).

Zeng, J., Mahou, P., Schanne-Klein, M.-C., Beaurepaire, E., & Débarre, D. 2012, *Biomedical Opt. Express*, 3, 1898 (pp. 14, 56, 61, 61, 65, 72, 72, 77, 81).

Zipfel, W. R., Williams, R. M., & Webb, W. W. 2003, *Nat. Biotechnol.*, 21, 1369 (pp. 11, 36).

Zuidervaart, H. J. 2010, in History of Science and Scholarship in the Netherlands, Vol. 12, The Origins of the Telescope, ed. A. Van Helden, S. Dupré, R. van Gent, & H. Zuidervaart (Amsterdam: KNAW Press), 9 (p. 3, 3).