



Universiteit
Leiden
The Netherlands

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

Wilby, M.J.

Citation

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

Version: Not Applicable (or Unknown)

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/67531>

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

Cover Page



Universiteit Leiden



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

Author: Wilby, M.J.

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

Issue Date: 2018-11-27

Chapter 2

The coronagraphic Modal Wavefront Sensor

A hybrid focal-plane sensor for the high-contrast imaging of circumstellar environments

Michael J. Wilby, Christoph U. Keller, Frans Snik, Visa Korkiakoski, and Alexander G.M. Pietrow

A&A 597, A112 (2017)

Abstract

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

2.1 Introduction

2.1.1 Scientific motivation

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

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

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

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

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

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

2.1.2 Holographic optics for focal-plane wavefront sensing

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

A modal approach to wavefront sensing has multiple advantages over traditional

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

2.1.3 Content of paper

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

2.2 Theory

2.2.1 The Holographic Modal Wavefront Sensor

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

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

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

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




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

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

2.2.1.1 Generating holograms

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

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

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

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

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

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

$$H_k(x, y) = |O_k(x, y) + R_k(x, y)|^2 \quad (2.4)$$

$$= |O_k|^2 + |R_k|^2 + O_k^* R_k + O_k R_k^* \quad (2.5)$$

$$= 2 + 2\mathbb{R}[O_k^* R_k], \quad (2.6)$$

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

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

$$I_{k\pm}(x', y') = \underbrace{\delta(\mathbf{x}' \pm \mathbf{x}'_k)}_{\text{Carrier Frequency}} * \underbrace{\left| \mathcal{F}[A(x, y)] \right|}_{\text{Telescope PSF}}^2 \quad (2.7)$$

$$* \left| \mathcal{F} \left[\underbrace{e^{i(a_k \pm b_k)M_k(x, y)}}_{\text{Wavefront Bias}} \underbrace{e^{i \sum_{j \neq k} a_j M_j(x, y)}}_{\text{Inter-Modal Crosstalk}} \right] \right|^2,$$

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

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

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

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

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

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

2.2.1.2 Sensor response

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

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

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

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

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

2.2.1.3 Wavefront reconstruction

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

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

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

$$\mathbf{a} = \hat{G}^{-1} \mathbf{I}, \quad (2.12)$$

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




Figure 2.2: Response curve of a 6-mode Zernike HMWS ($b_k = 0.7$ rad) to defocus error a_3Z_3 (all $a_{k \neq 3} = 0$), using photometric apertures of radius $r = 1.22\lambda/D$. The diagonal line (red) denotes a perfect sensor with 1:1 correspondence, which is achieved by the HMWS for $|a_3| \lesssim 0.5$. The response of the remaining sensor modes (Z_4 - Z_8 , grey) are well constrained about zero over the linear sensing regime, with residual non-linear inter-modal crosstalk behaviour manifesting for $|a_3| \gtrsim 1$.

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

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

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

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

2.2.2 Combination with an Apodizing Phase Plate coronagraph

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

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

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

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

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

2.2.3 Impact of multiplexing on mutual performance

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




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




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

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

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

$$\mathbf{a} = \hat{\mathbf{G}}^{-1} \mathbf{I} - \mathbf{a}_c \quad (2.14)$$

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

2.2.4 Impact of structured telescope apertures

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

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

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

2.3 Idealised performance simulations

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

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




Figure 2.5: Example closed-loop performance for a 6-mode Zernike (*left column*) and Fourier (*right column*) mode cMWS for $g = 0.8$. *a)* $N_i = 0$ aberrated PSFs, with $a_k = 1.0$ radians RMS per mode. *b)* Diffraction-limited PSF after N_i closed-loop iterations required to achieve convergence. PSF copies corresponding to each mode bias are labelled $Z_{n\pm}$ for Zernike modes and $F_{nS/C\pm}$ for Fourier modes, where the index S/C denote the sine and cosine mode phases respectively, and mode number corresponds to the circled regions of influence in the APP dark hole. The white circles overlaid on the $Z_{4\pm}$ modes indicated the $r = 1.22\lambda/D$ region of interest used for wavefront measurement. Note also the differing angular extent of each APP, which cover 180° and 90° for the Zernike/Fourier designs respectively. *c)* Azimuthally-averaged residual intensity plots corresponding to the PSFs of panels *a)* (green) and *b)* (blue); shaded regions denote 1σ variance averaged over the APP dark hole. *d)* Science PSF Strehl ratio (black diamonds) and residual RMS wavefront error (red squares) as a function of iteration number N_i . Vertical dashed lines indicate the point of convergence.

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

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

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

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




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

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

2.4 On-sky demonstration

2.4.1 Instrument design

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




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

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

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

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

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




Figure 2.8: 633 nm He-Ne laser calibration data illustrating the response of a 6-mode HMWS to controlled wavefront aberrations. *a)* PSF with no induced wavefront error and *b)* with 1.5 rad of defocus error applied via the SLM, showing the asymmetric response of the $Z_{3\pm}$ defocus WFS spot pair, indicated *top-right/bottom-left*. Filter and SLM reflection ghosts are present below the zeroth-order PSF, though these are sufficiently separated from the PSF copies to ensure no interference with HMWS performance.

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

2.4.2 Characterising HMWS on-sky response

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

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




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

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

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

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

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

2.4.3 Broadband wavefront sensing

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

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

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

2.4.4 Real-time atmospheric wavefront measurements

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




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

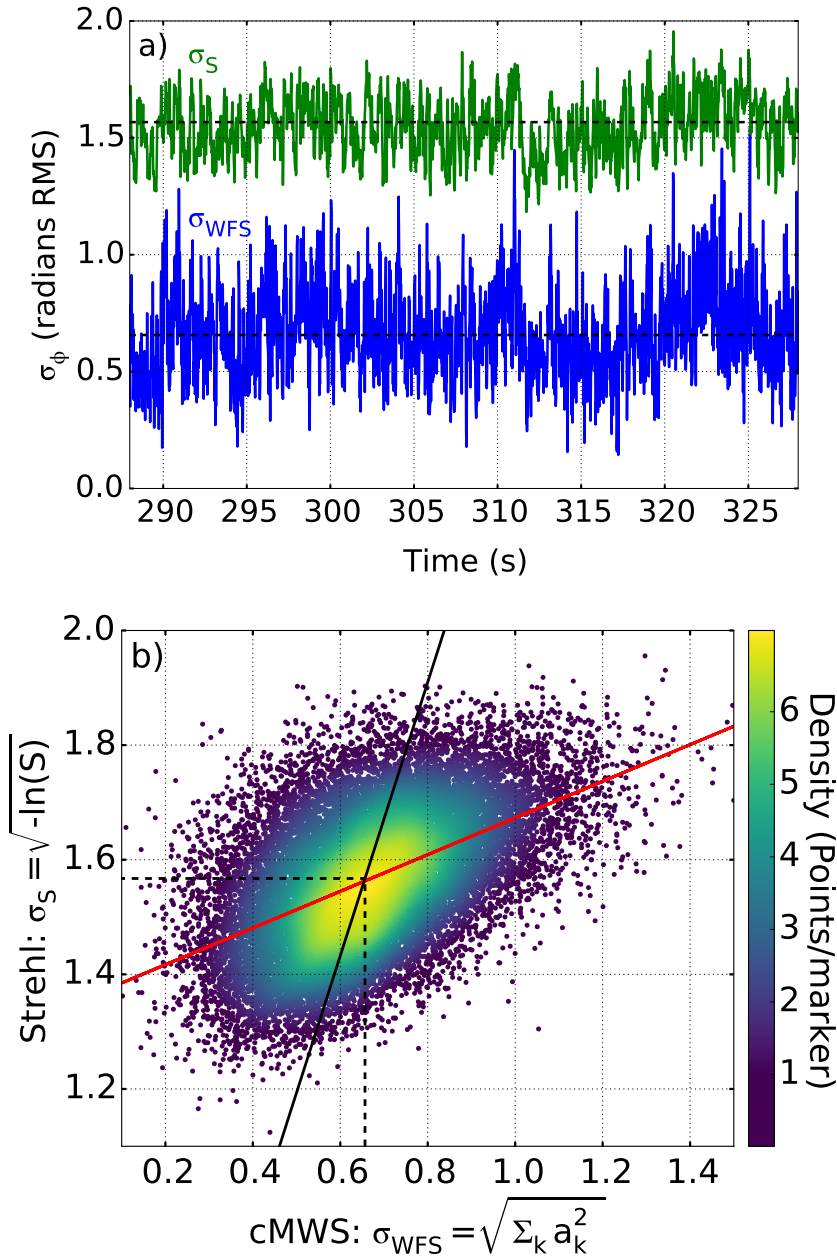


Figure 2.11: *a)* Time series of single-frame RMS wavefront error σ_ϕ as measured by i) the central PSF Strehl ratio under the Maréchal approximation (σ_S , green) and ii) calibrated cMWS measurements (σ_{WFS} , blue). Slow variations in seeing quality seen in σ_S are visibly traced by the sensor. *b)*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

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

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

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




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

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

2.5 Discussion & conclusions

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

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

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

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

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

