

Durham E-Theses

Point-diffraction interferometry for wavefront sensing in adaptive optics

NICOLAS SEBASTIEN DUBOST

How to cite:

DUBOST, NICOLAS SEBASTIEN (2020) Point-diffraction interferometry for wavefront sensing in adaptive optics. Doctoral thesis, Durham University.

Use policy




This work is licensed under a [Creative Commons Attribution Non-commercial 3.0 \(CC BY-NC\)](https://creativecommons.org/licenses/by-nc/3.0/)

**Point-diffraction interferometry
for wavefront sensing in
adaptive optics**

Nicolas Dubost

A thesis presented for the degree of
Doctor of Philosophy



Centre for Advanced Instrumentation
The University of Durham
United Kingdom
April 2020

Para mis padres

Point-diffraction interferometry for wavefront sensing in adaptive optics

Nicolas Dubost

Abstract

The work presented in this thesis aims at the development and validation of a wavefront sensor concept for adaptive optics (AO) called the pupil-modulated Point-Diffraction Interferometer (m-PDI). The m-PDI belongs to a broader family of wavefront sensors called Point-Diffraction Interferometers (PDIs), which make use of a small pinhole to filter a portion of the incoming light, hence generating a reference beam. This allows them to perform wavefront sensing on temporally incoherent light, such as natural guide stars in the context of astronomical AO.

Due to their high sensitivity, PDIs are being developed to address several difficult problems in AO, namely measuring quasi-static aberrations to a high degree of accuracy, the cophasing of segmented apertures, and reaching a high correction regime known as extreme AO. But despite their advantages, they remain limited by their narrow chromatic range, around $\Delta\lambda = 2\%$ relative to central bandwidth, and short dynamic range, generally of $\pm\pi/2$. The purpose of developing the m-PDI is to explore whether this new concept has any advantages regarding these limitations. Indeed, we find that the m-PDI has a maximum chromatic bandwidth of 66% relative to the central wavelength and a dynamic range at least 4 times larger than that of other PDIs.

Although the m-PDI concept had been proposed previously, it had not been explored to the extent reached in this manuscript. This thesis presents an initial investigation into the m-PDI, beginning with the development of the theory. Here the theoretical framework is laid out to explain how interference fringes are modulated by the wavefront, how to then demodulate the propagating electric field's phase and then finally how to measure the signal-to-noise ratio (SNR).

After building analytical and numerical models, a prototype is designed, built and characterised using CHOUGH, a high-order AO testbed in the lab. This incarnation of the m-PDI is called the Calibration & Alignment Wavefront Sensor (CAWS). The characterisation of the CAWS shows two things. The first one is that the CAWS' response is approximately flat across its spatial frequency domain. The second one is that its dynamic range decreases at higher frequencies, suggesting that it depends, amongst other things, on the wavefront's slope.

In order to prove that m-PDIs can be used for AO, a control loop is closed using the CAWS and CHOUGH's deformable mirror, with both monochromatic and broadband light. The results show that the final Strehl ratio increases from 0.2 to 0.66, at a wavelength of 633 nm. The difference in residual aberrations seen separately by the imaging camera and by the CAWS is about 20 nm RMS. This is explained by non-common path aberrations and low order aberrations which are invisible to the CAWS.

Finally, the instrument was tested on the CANARY AO bench at the William Herschel Telescope. The CAWS was successful at characterising the quasi-static aberrations of the system and at demodulating the phase of wavefronts produced with the deformable mirror. When demodulating on-sky residual aberrations at the back of CANARY's single-conjugate AO loop, the SNR remained too low for effective wavefront demodulation, only sporadically increasing above 1. These results are not discouraging as the CAWS was only a first prototype and CANARY is not a high-order system, reaching a Strehl ratio of around 0.5% at 675 nm. The lessons and improvements for future designs are to increase the diameter of the instrument's pinhole by at least twice, and deliver it a higher Strehl ratio by moving towards longer wavelengths and employing a higher order AO system.

Acknowledgements

PhD students are not, for the most part, the intellectual lone wolves we sometimes dream we are, cold-heartedly disrupting the establishment with rogue sparks of brilliance. We are instead industrious little workers, the oompa loompas of science, to quote a famous TV show, inching our way through research with the help of a veritable army of people. These people can not go without recognition, as so much of the work and learning I did over these many years would not have been possible without them.

At the top of the list must always be *mis padres*. *Ustedes nos dedicaron a sus dos hijos, incansables años de amor y esfuerzo. Junto con Julien hemos sido capaces de desenvolvernos con éxito por el mundo gracias a su apoyo y a su herencia. Gracias totales a ambos.*

I also want to thank my wonderful girlfriend Charlotte for all her patience and encouragement. This adventure, as all others, was so much better with you.

I can never thoroughly or fairly thank all of the people in CfAI who so generously helped me with their knowledge, skill and time. I can only list them in alphabetical order and cross my fingers that I didn't miss anyone. My gratitude goes to Alastair, Andrew, Andy, Ariadna, Graham, James, Lazar, Marc (but not Mark), Matt, Ollie, Penny, Tim B, and Tim M.

Special thanks to my office mates Daniel, Huizhe, David and Saavi for the technical and emotional support.

Quick thanks to Steve and to all the technicians at the Mechanical Engineering Services for their work and advice.

Also thanks to the people behind Matplotlib and to NumPy for providing open-source libraries that have become essential to science.

Being able to test the CAWS on-sky was an unexpected opportunity and an excellent experience. Thanks to the CANARY team and to our collaborators from CRAL, Maud, Magali and Clément for making it all possible.

Finally, the biggest of thanks go to my supervisors Ali and Richard. Thank you both for the countless hours of explanations, discussions and guidance. The best parts of this thesis are as much yours as they are mine.

Contents

Declaration	ix
List of Figures	x
List of Tables	xxii
Nomenclature	xxiii
1 Astronomical adaptive optics	1
1.1 Basic principles	2
1.1.1 Wavefront & phase aberration	3
1.1.2 Atmospheric turbulence	6
1.1.3 Optical performance metrics	11
1.1.4 Main AO components	13
1.1.5 Open-loop & Closed-loop configurations	16
1.2 Current challenges	21
1.2.1 Quasi-static aberrations	21
1.2.2 Phasing segmented apertures	23
1.2.3 Extreme Adaptive Optics	25
2 Point-diffraction interferometers & other solutions	29
2.1 A potential approach: point-diffraction interferometers	30

2.1.1	Definitions	30
2.1.2	The self-referencing Mach-Zehnder	33
2.1.3	The Zernike wavefront sensor	34
2.2	Comparative analysis: the point-diffraction interferometer (PDI) and other alternative techniques	36
2.2.1	Quasi-static aberrations	37
2.2.2	Phasing segmented apertures	44
2.2.3	Extreme Adaptive Optics	47
2.3	Summary	53
3	The pupil-modulated Point-Diffraction Interferometer	54
3.1	Concept	55
3.1.1	Principle	55
3.1.2	Formalism	56
3.2	Simulated model & phase demodulation	61
3.3	Noise propagation	67
3.4	Visibility loss from chromatic bandwidth and a decrease in Strehl . .	70
3.5	Accuracy and dynamic range	71
3.6	Application to a real system	76
3.7	Conclusions	78
4	Design & implementation	80
4.1	From requirements to optical design	81
4.1.1	Optical relay	83
4.1.2	Bias	88
4.1.3	Focal plane filter masks	91
4.1.4	Mechanical implementation	94
4.2	Testing & characterisation	96
4.2.1	Control software & data pipeline	97
4.2.2	Spatial transfer function	98

4.2.3 Spectrograph mode	107
5 Closing the control loop	114
5.1 Strehl ratio measurements	115
5.2 Control matrix	119
5.3 Closing the control loop	122
6 On-sky tests	131
6.1 Objectives	132
6.2 Experimental setup & requirements	135
6.3 Phase 0: Characterisation of quasi-static aberrations	146
6.4 Phase 1: On-sky AO residuals	153
7 Discussion	165
Appendix A Analytical derivation for the propagation of photon- shot and read-out noise	169
Appendix B Control matrix for the CAWS in CHOUGH	172
Appendix C Cophasing of segmented apertures	175
Bibliography	176

Declaration

The work in this thesis is based on research carried out at the Centre for Advanced Instrumentation, Department of Physics, University of Durham, England. No part of this thesis has been submitted elsewhere for any other degree or qualification, and it is the sole work of the author unless referenced to the contrary in the text.

Some of the work presented in this thesis has been published in a journal - the relevant publication is presented below.

Relevant Publication

Chapter 3

N. Dubost, N. A. Bharmal, and R. M. Myers. Calibration of quasi-static aberrations in high-contrast astronomical adaptive optics with a pupil-modulated point-diffraction interferometer. *Optics express*, 26(9):11068–11083, 2018.

Copyright © 2020 by Nicolas Dubost.

“The copyright of this thesis rests with the author. No quotation from it should be published without the author’s prior written consent and information derived from it should be acknowledged”.

List of Figures

1.1	Propagation of spherical wavefront into a flat wavefront.	4
1.2	First 10 Zernike polynomials numbered using Noll's sequential indices (Noll, 1976).	5
1.3	Point-spread function of an optical system with a circular entrance pupil. (a) Simulation of the intensity pattern across a surface. (b) Diagram of longitudinal section showing the interval defining the angular resolution.	6
1.4	Effects of the turbulence on the angular resolution of a telescope, for a wavelength λ . (a) Propagation without atmospheric turbulence and formation of a diffraction limited PSF for a telescope of diameter D . (b) Propagation through atmospheric turbulence with a Fried's parameter $r_0 < D$	8
1.5	Power spectral density for von Karman turbulence with inner scale of 5 mm and and outer scale of 20 m. The left and right dotted vertical lines indicate the wave-numbers for outer and inner scales, respectively.	11
1.6	Effect of atmospheric turbulence on SR. (a) System PSF without turbulence and (b) <i>seeing limited</i> PSF with turbulence.	12
1.7	2D section of a Shack-Hartmann WFS.	14

1.8	Effect of a deformable mirror on the phase of a wavefront. (a) Left - Incoming wavefront. (b) Right - Reflected wavefront. In this case the DM is not made of a continuous surface but of a segmented one.	14
1.9	Adaptive optics system in open-loop configuration.	16
1.10	Adaptive optics system in closed-loop configuration.	19
1.11	Non-common path aberrations. The common path, the imaging and WFS paths are respectively noted as (0), (1) and (2).	23
1.12	Approximation of the aperture of ESO's ELT, with telescope spiders included. In reality the gaps between segments are smaller.	24
2.1	Mach-Zehnder's working principle, based on image by aurantiaca.	31
2.2	PDI version of the Mach-Zehnder interferometer, based on image by aurantiaca.	33
2.3	Zernike WFS, where d is the depth of the transmissive focal plane phase mask.	35
2.4	Schematic presentation of SAXO. Adapted from Fusco et al. (2014).	48
2.5	Kolmogorov phase screens for (<i>left</i>) good <i>seeing</i> and (<i>right</i>) for bad <i>seeing</i> . Phase screens are 4.2 m wide.	53
3.1	Layout of the pupil-modulated point-diffraction interferometer. The grating splits the beam into modes. Mode 0, shown in a solid red line, goes through a point-diffraction pinhole in the focal plane. Mode +1 goes through a larger aperture.	56
3.2	Focal plane filter mask, in Plane B and with a square large aperture M_{+1}	58

- 3.3 Graphical representation of the Fourier transform of the intensity pattern I_C on the detector plane (*bottom*), using an electric field Ψ_0 for which the Fourier transform $\widehat{\Psi}_0$ has an arbitrary triangular shape (*top*). Filtering $\widehat{\Psi}_0$ by apertures M_0 and M_{+1} produces $\Psi_{B,0}$ and Ψ_{LP} respectively (*middle*). The width of filter M_{+1} was chosen to be as wide as possible, without incurring in aliasing between the central lobe $\Psi_{B,0} \otimes \Psi_{B,0}$ and the side lobes. This figure shows the absolute value of all represented terms. 59
- 3.4 Flow diagram of the monochromatic optical simulation of the m-PDI. This diagram is simplified and does not show all minor operations. The input of the simulation is a wavefront (panel *a*), which in this case follows a Kolmogorov law. The wavefront is turned into an electric field before being multiplied by a pupil plane mask (*b*), which includes the pupil aperture and the Ronchi ruling. The result is then zero-padded in order to control the scale of the sampling at the focal plane. Panel *c* shows the masked angle of the padded electric field that will be propagated to the focal plane. In this case $N_T/N_A = 5$. In all panels black pixels have a value of zero. The padded electric field is propagated to the focal plane (*d*) with a DFT. The electric field is then multiplied by a binary focal plane filter mask (*e*), before being propagated to the exit pupil by another DFT and then multiplied by its complex conjugate to obtain the intensity pattern on the detector plane (*f*). 63

- 3.5 Flow diagram of the demodulation process, showing actual results from the optical simulation and the demodulation. After a wavefront (*a*) is fed into the m-PDI (or in this case a simulation thereof) and the corresponding intensity pattern is obtained on the focal plane (*b*), the demodulation starts by computing the DFT of this pattern. On the Fourier plane (*c*), a sideband is cropped, as shown by an orange square and an inverse DFT is used to retrieve an estimation of the electric field at the entrance pupil. Panel *d* shows the phase of the estimated electric field. Although the resulting phase estimation has less spatial resolution than the input wavefront, in this figure it has been resampled in the Fourier plane to have the same number of pixels N_A as that input. 65
- 3.6 Demodulation outputs to three different known phase aberrations: (*a*) defocus, (*b*) astigmatism and (*c*) coma. Simulation performed with 30 line-pairs across the pupil, a padding factor of 5, and a pinhole of diameter $D_{B,0} = \lambda/D_A$. The RMS error between the input and the demodulated aberrations is between 5% and 7% for all three aberrations. 66
- 3.7 Propagation of read-out and photon-shot noise as a function of sensor flux available before the entrance pupil. Simulated data is presented for two sinusoidal wavefronts fed into the WFS, with frequencies $f = 1/D_A$ and $f = N_G/3D_A$ (*black dotted solid*). Simulations are compared to estimations using Eq. (3.21) for 3 values of b normalized by P_0 . The unravelling limit (*black dash-dotted*) is a limit on the standard deviation of noise, as its contribution to phase can not be unravelled beyond $\pm 0.5\lambda_0$. Consequently, as the noise increases the resulting phase contribution converges to a uniform distribution with a standard deviation of $1/\sqrt{12} \simeq 0.29$ in units of λ_0 69
- 3.8 Propagation of read-out and photon-shot noise as a function of normalized reference beam amplitude, for a given sensor flux. 69

3.9	(a) Input wavefront at entrance pupil A. The sinusoidal input is perpendicular to the grating's line-pairs and to the resulting fringes on the exit pupil. (b) Intensity at exit pupil C	71
3.10	Loss of visibility as a function of wavefront phase RMS in units of wavelength, for a sinusoidal aberration, perpendicular to the grating's line-pairs. Solid lines represent simulated polychromatic cases and the dashed line represents the theoretical monochromatic prediction. Visibility is measured in the place of maximum aberration $\phi_{0,LP}$. The vertical line marks the largest disagreement between the theoretical model and the simulation.	72
3.11	Response of the m-PDI to wavefront errors for different low-order aberrations and in different configurations. Each configuration is produced by changing one parameter from (i) the initial configuration. The 4 parameters that are changed are: (ii) the number of line-pairs in the pupil N_G , (iii) the number of pixels sampling an interference fringe N_P , (iv) the chromatic bandwidth $\Delta\lambda$ in units of λ_0 , and (v-a) and (v-b) the focal plane pinhole's diameter $D_{B,0}$	73
3.12	Transfer function with respect to spatial frequency for different aberration amplitude and in different configurations. Columns have different pinhole sizes and rows have different preexisting tilts.	75
3.13	Exposure time required to achieve $SNR = 1$ as a function of aberration spatial scale. (left) $N_G = 60$, (right) $N_G = 180$. The exposure time is presented for different amplitudes of aberration and for different chromatic bandwidths. The parameters of the simulation can be found in Table 3.1.	77

4.1	Layout of the high-order AO testbed CHOUGH. The HOWFS is a 31-by-31 SH, the NFSI is a narrow FOV imaging camera, the DM is a 32-by-32 Boston Micromachines Kilo DM, the dichroic is a 647 nm long-pass one, and light is provided by a fibre-fed illumination interface. For all tests presented in this thesis the Atmospheric Dispersion Corrector (ADC) has been removed.	82
4.2	M_{+1} sideband region on DFT of the interferogram. N_T is the total number of pixels across the pupil.	85
4.3	Diagrams of the first half of the CAWS going from the diffraction grating to the first lens' focal plane, with a defocused focal plane filter mask, with a (a) a 4f telecentric and (b) a non-telecentric optical relay design.	86
4.4	Quantum Efficiency of the Imperx Bobcat B0620M, based on data found in the camera's manual.	88
4.5	Optical model of the CAWS using OpticStudio. D_A and D_C are the diameters of the entrance and exit pupils respectively.	89
4.6	Estimation of the differential chromatic defocus term in the CAWS' focal plane, between wavelengths spanning its chromatic range and the central wavelength $\lambda_0 = 675$ nm. Curve was produced with OpticStudio	89
4.7	Biases introduced by (<i>top left</i>) ℓ_1 , (<i>top right</i>) ℓ_2 and (<i>bottom</i>) $\ell_1 + \ell_2$, for $\lambda_0 = 675$ nm.	91
4.8	b_0 as a function of pinhole diameter $D_{B,0}$, for a circular pupil.	92
4.9	Layout of focal plane filter masks.	93
4.10	Mechanical assembly of the CAWS made with Inventor.	95
4.11	Image of the CAWS.	96

- 4.12 Layout of the data acquisition, processing and control software for experiments of CAWS on the CHOUGH high-order AO bench. In the interface layer, the real-time controller DARC interfaces with the DM, the HOWFS, the NFSI and the CAWS. DARC loads pixel data from the cameras of these instruments into its pixel buffer and pushes commands found in its actuators buffer onto the DM. DARC also computes HOWFS centroids and loads them into its centroid buffer. All these buffers are accessible from the user layer by a Python DARC client. This client then delivers this data to a phase demodulation algorithm or to any other processing and control script defined by the user. . . . 97
- 4.13 (a) Interference fringes produced with a DM in its zero position, (b) the DFT of the fringes and (c) the corresponding demodulated reference phase. Panel *b* shows the square root of the absolute value of the DFT for display purposes. The light source is a HeNe laser at $\lambda = 633$ nm. The physical scale of the pupils in the fringe image and the demodulated phase is the same as they span the length of the pupil, but the sampling is lower on the demodulated phase. The reason for this is a line-pair is being sampled by $N_P = 6.8$ pixels and the width of 3 line-pairs equals the width of 2 phase data points. The figure also shows the demodulation process, including the cropping of the sideband in the Fourier plane, the inverse DFT, the computing of the angle of the resulting electric field and the unwrapping of the phase. The cropped region in the Fourier plane is rounded down to 35 pixels in order to avoid aliasing. . . 99
- 4.14 (*top*) Amplitude of sinusoidal wavefronts measured by the CAWS. The DM command axis shows the uncalibrated amplitude commanded onto the DM, using a simple volts-to-radians proportionality constant. (*bottom-left*) Cut along the DM command axis. (*bottom-right*) Cut along the DM frequency axis. The light source is a HeNe laser at $\lambda = 633$ nm. . . 100

4.15	(<i>top</i>) Amplitude of reconstructed sinusoidal wavefront estimated by the HOWFS. The DM command axis shows the uncalibrated amplitude commanded onto the DM, using a simple volts-to-radians proportionality constant. (<i>bottom-left</i>) Cut along the DM command axis. (<i>bottom-right</i>) Cut along the DM frequency axis. The light source is a HeNe laser at $\lambda = 633 \text{ nm}$	102
4.16	CAWS' monochromatic transfer function, calibrated by the HOWFS for $\lambda = 633 \text{ nm}$	103
4.17	Layout of the CAWS in spectrograph mode. The detector is pushed back into Plane D and an auxiliary lens is placed after ℓ_2 , in order to re-image the focal plane filter mask, located on Plane B.	108
4.18	Focal plane images of diffraction modes 0 and +1 with (<i>top</i>) OSL2B halogen lamp and (<i>bottom</i>) HeNe laser as reference for the spectral calibration. Red lines show, from left to right, the positions of $\lambda = 0 \text{ nm}$, $\lambda = 633 \text{ nm}$, and $\lambda = 647 \text{ nm}$, which correspond to the pinhole, the HeNe laser, and the dichroic's cutoff respectively. Original values have been square-rooted to increase the contrast of the images presented in this figure.	109
4.19	Normalised chromatic spectrum of the OSL2B halogen light source, (<i>blue</i>) as seen through CHOUGH and measured by the CAWS on mode +1, and (<i>black</i>) as measured by an independent spectrograph. The CoM and FWHM are presented for the spectrum as sensed on mode +1.	111
5.1	(<i>left</i>) Zero-padded discrete model of the pupil and (<i>right</i>) a cropped image of the resulting PSF, with values in arbitrary units. The binary model of the pupil has ones inside of it and zeros outside. N_d and N_p are the number of array cells across the diameter of the pupil and across the entire array once zero-padded, respectively. On the PSF image, array elements have the same plate scale as pixels in the NFSI camera.	116

5.2	(<i>top-right</i>) Synthetic PSF with no aberrations, (<i>top-left</i>) same PSF intentionally centred around the junction between four pixels, (<i>bottom-left</i>) real image from the NFSI, and (<i>bottom-right</i>) real image, numerically processed to have its peak value centred around a pixel. All images have been normalised so that the synthetic centred PSF has a peak value of 1, and also so they all have the same total energy, i.e. the sum of the intensity.	118
5.3	Map of temporal frequencies assigned to active actuators for the construction of the control matrix. There are 812 active actuators, with cycles per period of array A from 1 to 812. Inactive actuators are masked.	121
5.4	Normalised singular values, sorted from highest to lowest. The <i>left</i> panel shows the 15 largest values in detail and the <i>right</i> panel shows the rest.	125
5.5	CAWS phase measurements of the initial static aberrations (<i>left</i>), the final aberrations the system converges to after closing the control loop (<i>centre</i>), and a section at the centre of the pupil for those central aberrations (<i>right</i>).	126
5.6	NFSI focal plane images with (<i>left</i>) the initial static aberrations, (<i>centre</i>) the final aberrations the system converges to after closing the control loop, and (<i>right</i>) the synthetic PSF used as reference to calculate SR, for comparison. All three images have been normalised to have the same cumulative intensity, and have been centred around a single pixel. . . .	126
5.7	Temporal evolution of (<i>black</i>) residual RMS aberrations and (<i>blue</i>) SR from the moment the control loop is closed until it converges. Data with the same type of marker are equivalent and were converted from (<i>black</i>) RMS aberrations to (<i>blue</i>) SR, or <i>vice versa</i> , using Maréchal's approximation. The origins of the data are (<i>solid line</i>) CAWS phase measurements including tip-tilt, (<i>dotted line</i>) the same measurements excluding it, and (<i>dashed line</i>) NFSI focal plane images.	127

5.8	Temporal evolution of (<i>black</i>) residual RMS aberrations and (<i>blue</i>) SR from the moment the control loop is closed until it converges. These results correspond to the same experiment presented in Figure 5.7, but with a broadband halogen lamp rather than a HeNe laser. The equivalent central wavelength is 667 nm and the bandwidth is 76.5 nm.	128
6.1	CANARY’s SCAO loop, with CAWS, an imaging camera and the iMZ.	136
6.2	Image of the experimental setup, showing the CAWS, the iMZ, the PSF imaging arm, the beginning of the TS and a part of CANARY.	137
6.3	CANARY’s optical model (CANARY team, 2019), from the telescope’s focal plane to the first 10R/90T beamsplitter. The reflected arm containing the TS is included, whereas the transmitted arm alongside the iMZ and CAWS are not.	139
6.4	Reflectance of Protected Aluminium, for an angle of incidence of up to 12°, as provided by Thorlabs (2019).	139
6.5	Estimated throughput of the WHT + CANARY + 10R/90T beamsplitter	140
6.6	Estimated throughput of the WHT + CANARY + CAWS, including camera’s QE. Rapidly decreasing values after 750 nm are explained in part by the approaching sensitivity limit of silicon detectors in the NIR.	140
6.7	Optical model of the common-path interface, the optical 4F relay and CAWS.	142
6.8	Minimum exposure time to achieve $\text{SNR}_d > 1$, for different filters. Solid and dash lines show the exposure time with a respective pixel binning of a $\times 1$ and $\times 2$	144
6.9	List and elevation of target guide stars, as seen at the Roque de los Muchachos Observatory on the night of July the 17 th , 2019. Star elevations processed by the Isaac Newton Group of Telescopes (2019). . . .	145

6.10	(<i>left</i>) Residual static aberrations on CAWS in nanometers and (<i>right</i>) corresponding PSF image measured on its focal plane mask, mode +1. Reference phases and focal plane measurements were produced using the 635 nm diode laser internal source.	146
6.11	Interference fringes obtained using (<i>left</i>) the monochromatic diode laser and (<i>right</i>) the broadband halogen lamp.	148
6.12	Focal plane images of diffraction modes 0 and +1 with (<i>top</i>) halogen lamp and (<i>bottom</i>) diode laser for reference. Original values have been square-rooted to increase the contrast of the images presented in this figure.	149
6.13	Zemax estimation of the differential chromatic defocus term in the CAWS' focal plane, between wavelengths spanning its chromatic range and the diode laser's wavelength (635 nm). The model includes CANNARY and the optical interfaces leading to the CAWS.	149
6.14	Chromatic spectra (<i>blue</i>) as measured on mode +1 and (<i>green</i>) as reconstructed for mode 0. The values of these curves have been normalised to the same reference so that the maximum common value between them is equal to 1. This normalisation preserves their relative magnitudes. Both spectrums are used to estimate (<i>red</i>) the normalised amplitude of interference fringes as a function of wavelength. The CoM and FWHM are presented for the fringes' amplitudes. R is the theoretical spectral resolution.	151
6.15	(top left) Monochromatic references, (top right) broadband references, corrected for differential defocus, and (bottom) error between them, all in nanometers. Blue central circles mark the location of the telescope's central obscuration.	152

6.16	RMS error between monochromatic and broadband measurements, corrected for chromatic defocus, as a function of central wavelength. Vertical dashed lines show (<i>blue</i>) the wavelength for minimum error and (<i>red</i>) the central wavelength found through the spectrum reconstruction method.	153
6.17	Off-sky (<i>left</i>) interference fringes with (<i>right</i>) their respective Fourier transform (the square-root of the absolute value is displayed for contrast enhancing purposes). Fourier transform panels also show (<i>blue</i>) the sideband region M_{+1} , (<i>green</i>) the noise estimating region $M_{\text{estim.}}$, and the resulting estimation of $\text{SNR}_f^{(\text{empiric})}$	157
6.18	Off-sky sequence of images from <i>replay</i> data, with decreasing SNRs, from left to right and then from top to bottom.	159
6.19	On-sky images with the highest $\text{SNR}_f^{(\text{empiric})}$. Images were produced while observing Beta Pegasi, with filter 6 and an integration time of 400 ms.	160
6.20	SNR of fringes for different exposure times and filters, and the expected value of an unfitted model for $\text{SR}=0.012$. The bars extend $\pm 1.5\sigma$ from the average value, spanning a total of 3σ . Results were produced on-sky while observing Beta Pegasi.	160
B.1	Transpose of A . This is the shortest possible A for which $M = 2L_{\text{active}}$. Green columns show inactive actuators.	173
B.2	Control matrix mapping the relationship between CAWS' phase points and DM actuators.	174
C.1	Preliminary simulation of (<i>top</i>) random phase aberrations through an arbitrary ELT-like pupil and (<i>bottom</i>) the error of the wavefront reconstructed by the CAWS, in units of λ_0	175

List of Tables

1.1	Summary: extreme adaptive optics (XAO) benchmarks and their performance.	27
2.1	Summary: Approaches to reduce quasi-static aberrations.	43
2.2	Summary: Approaches for segment phasing.	47
2.3	Summary: XAO systems, their wavefront sensor (WFS)s and their performances	50
3.1	Parameters used for simulating the exposure time.	77
4.1	List of focal plane filter masks.	94
6.1	List of cameras by instrument.	137
6.2	List of internal sources.	138
6.3	List of components by position on the filter wheel.	141
6.4	Configuration parameters for CAWS	142
6.5	List of target guide stars, their coordinates and magnitudes.	144

Nomenclature

ADI angular differential imaging

ASDI angular and spectral differential imaging

AO adaptive optics

CAWS Calibration and Alignment Wavefront Sensor

CHOUGH Canary Hosted-Upgrade for High-Order Adaptive Optics

CoM centre of mass

DARC Durham Adaptive Optics Real-time Controller

DFT Discrete Fourier Transform

DM deformable mirror

ELTs extremely large telescopes

FOV field of view

FWHM full-width at half-maximum

iMZ integrated Mach-Zehnder

MEMS micro electro-mechanical systems

m-PDI pupil-modulated PDI

NCPAs non-common path aberrations

NGS natural guide star

NIR near-infrared

OPD optical path difference

PDI point-diffraction interferometer

PWFS pyramid wavefront sensor

PSF point-spread function

PtV peak-to-valley

QE quantum efficiency

RTC real-time controller

SCAO single-conjugate adaptive optics

SH Shack-Hartmann

SNR signal-to-noise ratio

SDI spectral differential imaging

SR Strehl ratio

SRI self-referencing interferometer

SVD singular value decomposition

TS truth sensor

VLT Very Large Telescope

WFE wavefront error

WHT William Herschel Telescope

WFS wavefront sensor

XAO extreme adaptive optics

ZWFS Zernike WFS

*“Every valley shall be exalted,
and every mountain and hill shall
be made low: and the crooked
shall be made straight, and the
rough places plain.”*

Isaiah 40:4, King James Bible

CHAPTER **1**

Astronomical adaptive optics

The aim of this thesis is to present the development of a new wavefront sensing technique for astronomical adaptive optics (AO). The new WFS, dubbed the pupil-modulated PDI (m-PDI) belongs to the broader family of point-diffraction interferometers. Despite being originally offered as a solution to a particular problem, this hammer has since found new nails in the field of AO. Indeed, it is this area of research that motivates the development of this new technique.

In order to provide context and imbue further explanations and experiments with a sense of purpose, the potential applications of this novel approach will be presented first in the form of “challenges”. Therefore, before fully delving into this thesis’ main proposition, the m-PDI, this chapter lays out the context and main incentives for this work. First, Section 1.1 presents basic principles in astronomical adaptive optics. Some of these definitions and concepts might seem to be unprompted, but they will allow us to establish a common language that will be useful to present the challenges facing the field of AO. Following this, Section 1.2 explains these challenges, which will serve as the principal motivations for the development of the m-PDI.

1.1 Basic principles

Adaptive optics is a technology that enhances the angular resolution of ground-based astronomical observations by correcting the random optical distortions induced by the atmosphere in real time. This technology was first developed in the 1970s for defense applications, specifically laser beam compensation and satellite imaging (Hardy, 1998). But it wasn't until 1991, around the time it was declassified by the United States Air Force, that civilian applications started to be developed. Nowadays AO can be found in diverse fields such as astronomy, microscopy, ophthalmology, telecommunications and even fusion research. Amongst these, the astronomical community was one of the first to recognise its potential for dealing with one of the major hurdles to producing high resolution images in the visible and near-infrared (NIR) parts of the electromagnetic spectrum: atmospheric turbulence. Indeed, today this is AO's main application and provides the motivations for this thesis.

After traveling mostly undisturbed from what can sometimes be the furthest reaches of the observable universe, light rays get distorted in their final moments before reaching us on the surface of the Earth. The atmosphere, despite its transparency in the visible part of the spectrum, does not present itself as a homogeneous medium through which light can travel unscathed. Instead, different temperature air masses are constantly mixed by buoyancy and wind, the later of which produces vortices of heterogeneous temperature that break up and decay into smaller ones. As the index of refraction of a gas depends on the temperature, two parallel light rays going through zones of different temperatures will experience different indices of refraction, and will therefore travel different optical path lengths. As a result the light rays get shifted out of phase. This is how phase aberrations become imprinted onto propagating electric fields, reducing the angular resolution of astronomical telescopes. In order to explore the meaning of the previous explanation, some basic definitions need to be made with respect to optics, the characteristics of optical

systems and atmospheric turbulence.

1.1.1 Wavefront & phase aberration

Light being an electromagnetic wave, its electric field can be described by its wavelength λ , its *celerity* or phase velocity c , and its complex amplitude of absolute amplitude P and phase φ . The phase $\varphi(p, t)$ is a function of both time t and position p . However, in the visible part of the spectrum, electromagnetic waves have frequencies between 405-789 THz. At these frequencies and due to its fast variation, it is very difficult to measure the phase. Instead, most instruments measure the phase difference between two points, which, for a distant point-like source, only changes under the action of a perturbation. Since the phase difference between two points is as persistent as the perturbation producing it, in the case of it being atmospheric turbulence requiring sampling rates of between 200Hz-2kHz, it is instrumentally feasible to determine the phase difference between all points contained in an arbitrary surface and a reference point in said surface. From this data, the set of all points, or locus, of equal phase (not necessarily contained in the arbitrary surface) can be determined. This set of points, known as the wavefront, is the knowledge required to correct for the optical aberrations introduced by the turbulence.

In astronomy, instruments and telescopes are all designed to collect and focus wavefronts describing flat planes. This is because point-like sources, such as stars, produce spherical wavefronts which then become flat after propagating great distances. This is shown in Figure 1.1. But as mentioned earlier, atmospheric disturbances can change the phase of light around a certain area, hence changing the shape of the wavefront. Any deviation between the actual and the expected shape of the wavefront is called a wavefront error (WFE), and is measured in units of length. For monochromatic light, a WFE is equivalent to a phase aberration. Given a surface S representing the expected shape of the wavefront, the phase aberration at a point $p \in S$ is formally measured as the difference between $\varphi(p)$ and the average




Figure 1.1: Propagation of spherical wavefront into a flat wavefront.

phase across S , $\bar{\varphi} = \langle \varphi(p) \rangle_{p \in S}$. The phase aberration can be expressed as

$$\phi(p) = \varphi(p) - \bar{\varphi}. \quad (1.1)$$

In order to restore the quality of an astronomical image, the task of an AO system is to reduce the phase aberrations to a minimum, for all points p across the pupil. Since here the WFE is described for every point or “zone” p , this form of representation is called a zonal representation. But depending on the application and technology, sometimes it is useful to describe the aberrations in terms of modes. A modal representation describes a function as the weighted sum of other functions called modes. For every mode, its assigned weight has a unique value which will make the sum of modes equal to the target function. In optics, a set of modes commonly used are Zernike polynomials. These polynomials constitute an orthogonal basis inside the unitary disc and can therefore describe any function ϕ defined inside a pupil of radius R as

$$\phi[R\rho, \theta] = \sum_j a_j Z_j[\rho, \theta], \quad (1.2)$$

where j is the polynomials index, Z_j is the polynomial, ρ and θ are respectively the radius and angle inside the unitary disk, and a_j is the weight attributed to the mode.

For the purposes of this thesis, it is important to know that the first few Zernike polynomials describe some of the most common aberrations found in optical systems and also the strongest aberrations caused by atmospheric turbulence, as quantified by Noll (1976). For example, Figure 1.2 shows the polynomials corresponding

to tip, tilt, defocus, vertical and oblique astigmatism, vertical and horizontal coma, and vertical and oblique trefoil. When assessing the accuracy of a wavefront sensor

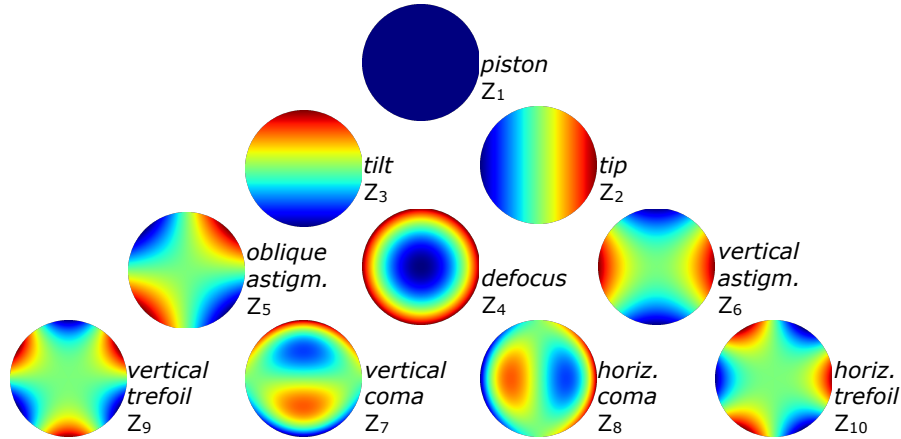


Figure 1.2: First 10 Zernike polynomials numbered using Noll's sequential indices (Noll, 1976).

or the correction of an AO system, it is important to do so against these aberrations, as they tend to be the biggest and most detrimental to the angular resolution of an imaging system.

Point-spread function & angular resolution A defining feature of an imaging system is its angular resolution. The angular resolution is the minimum angle between two point-like objects for which they can be resolved. At a smaller angular distance, the two objects will be indistinguishable from a single point. For example, if two binary stars are closer to each other than the angular resolution, they will look like a single star. In the case of other extended objects, small features will be blurred into single larger features. In order to increase the image quality and to be able to resolve as many details as possible, an imaging system's angular resolution needs to be as small as possible.

Given an entrance pupil with a fixed size, the best possible angular resolution is achieved when the incoming wavefront is flat. In this case, the system is said to be in the diffraction limit, because the only thing limiting the angular resolution is the phenomenon of diffraction. Under these conditions, the image of the point

source obtained through an imaging system is said to be that system's point-spread function (PSF). For all existing telescopes, most stars are far enough away that, if there was no atmospheric turbulence, their image would be the PSF of that telescope. The PSF produced by a telescope with a circular aperture of any diameter D and at any wavelength λ is shown in Figure 1.3. For such a telescope, the angular




Figure 1.3: Point-spread function of an optical system with a circular entrance pupil. (a) Simulation of the intensity pattern across a surface. (b) Diagram of longitudinal section showing the interval defining the angular resolution.

resolution is defined as $1.22\lambda/D$, or as the angular distance or spread between the central maximum and the first minimum of the PSF. Normally this means that a larger telescope should have smaller (better) resolution. Unfortunately, this is not always the case in presence of atmospheric turbulence, which will deflect light away into an area larger than the system PSF.

1.1.2 Atmospheric turbulence

Achromaticity An important aspect of atmospheric turbulence is that it affects visible and NIR wavelengths of light in a similar manner. The optical path length travelled by a ray of light is dependent on the index of refraction of the medium. The index of refraction is closely approximated by

$$n = (P/T)77.6 \cdot 10^{-6} - 1, \quad (1.3)$$

where P is the pressure in millibars and T is the temperature in Kelvin. Note that this expression does not consider the light's wavelength λ . As a result, rays on the

same path but with different wavelengths will travel the same optical path length. In other words, if phase aberrations ϕ_r in units of phase (e.g. radians) are instead converted into aberrations ϕ_z in units of length (e.g. meters, nanometers) such that

$$\phi_z = \frac{\lambda}{2\pi} \phi_r, \quad (1.4)$$

then all wavelengths will have the same aberrations. Therefore, any device measuring or correcting the WFE directly in units of length will have the benefit of being able to do so for all wavelengths at once. Despite this, some of the different metrics that characterise the strength of turbulence were first defined with respect to aberrations measured in radians and must therefore specify the wavelength they use.

Fried's parameter One of the metrics that is used to characterise the optical strength of atmospheric turbulence and that needs to specify a working wavelength is Fried's parameter. Due to the distortions imprinted by atmospheric turbulence, the resolution of a telescope does not improve after enlarging its aperture size beyond a certain limit. This limit is given as a measure of the optical strength of the turbulence known as Fried's parameter or r_0 . By definition, r_0 is the diameter over which the phase aberration has a mean-square value of near 1 rad^2 at a wavelength of 500 nm. Fried's parameter is sometimes also called the length of coherence. The reason for this is that two beams separated by a distance shorter than this length will have a similar phase and can constructively interfere when focused, i.e. they are spatially coherent. When the WFE has a mean-square value higher than 1 rad^2 the image quality, which can be described by the angular resolution or other optical performance metrics (see Section 1.1.3), deteriorates rapidly. A useful interpretation of this parameter is that it represents an upper threshold after which increasing the diameter of a telescope does not significantly increase its angular resolution in the presence of turbulence.

The long exposure, atmospherically distorted image of a star is called the seeing




Figure 1.4: Effects of the turbulence on the angular resolution of a telescope, for a wavelength λ . (a) Propagation without atmospheric turbulence and formation of a diffraction limited PSF for a telescope of diameter D . (b) Propagation through atmospheric turbulence with a Fried's parameter $r_0 < D$.

disk. As with the PSF, the angular resolution of this image depends upon the disk's angular spread. The relationship between the angular spread of the seeing disk and Fried's parameter is illustrated on Figure 1.4. When $r_0 < D$, the angular resolution becomes $0.98\lambda/r_0$. In a good site, Fried's parameter can take values ranging from <5 cm in poor seeing conditions to >20 cm in excellent seeing, and can change by a factor of 2 within a few seconds (Hardy, 1998). As an example, for a telescope of diameter $D = 4$ m in good seeing conditions, r_0 is almost always at least 20 times smaller, resulting in a proportionally diminished angular resolution. For a wavelength $\lambda = 500$ nm and $r_0 = 20$ cm, the resolution with turbulence would be 0.6 arcsec, in contrast with 0.03 arcsec without turbulence. The potential gains of

an adaptive optics system that can, even partially, correct for the distortions of the turbulence is therefore considerable. This is especially true when contemplating even larger telescopes like the already well-deployed 8 m class and future extremely large telescopes (ELTs) of over 20 m.

A second benefit of employing an AO system is to produce brighter images. Distant astronomical objects are not just small in angular size but can also be very faint. When trying to observe faint objects, there are two main approaches to increasing the observed signal. The first one is to set a longer exposure time on the detector, in order to integrate a proportionally larger number of photons. The second one is to have a telescope with a larger diameter. The effect of enlarging the telescope is twofold. On one hand, the number of collected photons is proportional to the telescope's surface area, which in turn is proportional to the square of the diameter. For both a large enough telescope and a long enough exposure time this is true regardless of the turbulence. On the other hand, if there were no turbulence, increasing the telescope's diameter increases the angular resolution and squeezes the same number of photons into an ever smaller PSF. By this effect alone the intensity at the centre of the image increases to the square of the diameter. For detector pixels of constant field of view (FOV) and a well sampled PSF, both effects combine so that the detected central intensity increases to the power of 4 of the diameter. Assuming that sources of noise such as detector read-out noise and background light remain constant when the PSF size is reduced, and considering that the astronomical signal is proportional to the detected central intensity, then the signal-to-noise ratio (SNR) would also increase to the power of 4, should the aberrations be corrected for.

Kolmogorov's turbulence model Despite its usefulness, Fried's parameter does not fully characterise the optical effects of atmospheric turbulence. For a more complete description one requires statistical models. These models generally describe the turbulence's power spectral density. The most commonly used is Kol-

mogorov's model, which states that the three-dimensional phase spectral density is

$$\Phi(\kappa) = 0.023r_0^{-5/3}\kappa^{-11/3}, \quad (1.5)$$

where, κ is the spatial wave-number defined as $\kappa = 2\pi/\ell$, with ℓ being the spatial scale, and r_0 is Fried's parameter. This expression holds valid within the inertial range for $\mathcal{L}_0 > \ell > \ell_0$, where \mathcal{L}_0 is known as the outer scale and ℓ_0 is the inner scale. Whereas the range of the inner scale has been found to be about 1 mm close to the ground and about 1 cm near the tropopause (Roddier, 1981), the size of the outer scale is less constrained. Beckers (1993) suggests that turbulent energy is fed into the atmosphere by various source such as wind shear, solar heating and the surrounding topography. As a result, the outer scale is not a single number and, although it is usually considered to be around a few tens of meters (about 10~20 m), it has been found to vary from between 1 m to 100 m.

Extensions to Kolmogorov's model exist that consider these variables. The most important one is von Kármán's model, which prevents the amount of turbulent energy from diverging to infinity by incorporating the outer scale (Rao et al., 2000). The model in terms of r_0 is

$$\Phi(\kappa) = 0.023 \frac{\|\kappa^2 + \kappa_0^2\|^{-11/6}}{r_0^{5/3}}, \quad (1.6)$$

where $\kappa_0 = 2\pi/\mathcal{L}_0$. The energy below the inner scale is small enough that not considering it does not cause any fundamental problems when modelling phase aberrations produced by turbulence. But in case a more accurate power spectrum is required, the power spectrum with the inner scale becomes

$$\Phi(\kappa) = 0.023 \frac{|\kappa^2 + \kappa_0^2|^{-11/6}}{r_0^{5/3}} e^{-\kappa^2/\kappa_m^2}, \quad (1.7)$$

where $\kappa_m = 5.92/\ell_0$ (Azoulay et al., 1988). An example of the resulting power spectrum is presented in Figure 1.5.

The shape of this power law informs the design of AO systems. Since most of the turbulence's power is concentrated at large spatial scales, there is a higher

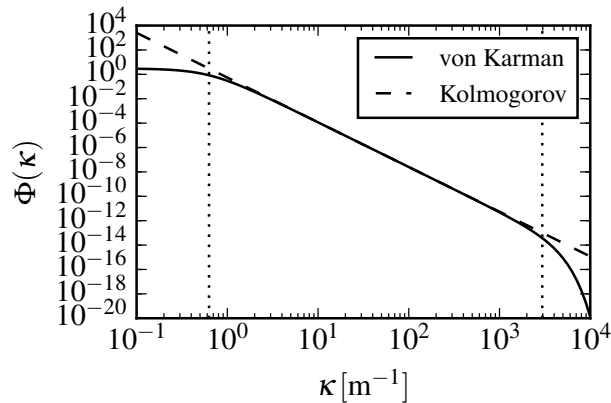


Figure 1.5: Power spectral density for von Karman turbulence with inner scale of 5 mm and and outer scale of 20 m. The left and right dotted vertical lines indicate the wave-numbers for outer and inner scales, respectively.

potential gain from being able to correct the aberrations at these scales first. This, combined with larger scale aberrations being easier to measure, led to efforts being first directed towards designing systems that could accurately sense at these scales and that had the necessary dynamic range to deal with these levels of power. As large scale aberrations have been successfully addressed, the technology has gradually been pushed into removing the remaining higher order (lower spatial scale) aberrations, where corrections are more difficult and where fewer gains are to be made in terms of corrected wavefront phase variance. The result is better angular resolution and also improvements quantified by other optical performance metrics, such as image brightness and SNR.

1.1.3 Optical performance metrics

Strehl ratio Phase aberrations spread the photons away from the diffraction limited PSF and into a large seeing disk, cancelling out one of the aforementioned benefits of having a larger telescope diameter. The result is an image with a smaller peak intensity relative to that of the PSF, obtained with the same telescope and exposure time. The ratio between the peak intensity of the seeing disk and that of the theoretically perfect PSF is called the Strehl ratio (SR). A graphical description

is provided in Figure 1.6. In the figure, both the PSF and the seeing disk have




Figure 1.6: Effect of atmospheric turbulence on SR. (a) System PSF without turbulence and (b) *seeing limited* PSF with turbulence.

been normalised by the same value, so that the peak intensity of the system PSF equals 1. Having done this, the values measured at their peak intensities encodes the SR, which in the case of the PSF equals 1 by definition. Due to diffraction, it is impossible to achieve a SR higher than 1. In other words, there is no phase pattern $\phi(x, y)$, where the coordinate pair (x, y) describes a position in the pupil, for which the maximum intensity of the image produced by a point source can be brighter than that achieved for $\phi(x, y) = 0$. Instead, any ϕ other than 0 will result in a SR smaller than 1, and for increasingly large aberrations, its value will converge to 0.

Despite not being necessarily true for all scientific cases, the performance or success of an AO system attempting to correct the phase aberrations is primarily assessed by the achieved SR. This metric is preferred above others for being the most direct and unbiased, as it measures the focal plane image quality on the image itself. Furthermore, in astronomy there are abundant point sources (stars) to produce seeing disks on which to measure the SR. This is not to say that this is a completely unbiased metric. Estimating the PSF of a telescope can be particularly difficult, leading to an uncertainty in the absolute value of the SR. Nonetheless, for any estimated PSF, the SR of any pair of images can still be compared, and by extension, the performance of AO at the instants the images were acquired. In this case, instead of absolute, their SRs are said to be relative.

Maréchal’s approximation The SR can also be used to estimate the variance σ^2 of the phase across the pupil, defined as $\sigma^2 = \langle (\varphi - \bar{\varphi})^2 \rangle$. Maréchal’s approximation, the relationship linking them together, is

$$S \simeq e^{-\sigma^2}, \quad (1.8)$$

where S is the SR and σ^2 is in units of rad^2 . Despite underestimating its value, this approximation is widely used to estimate the residual aberrations of an AO system due to its small error of less than 10% for $S > 0.6$, and converging to 0% for S approaching 1 (Mahajan, 1982, 1983).

1.1.4 Main AO components

Wavefront sensor In order to correct the phase aberrations, first they need to be estimated. Estimating $\phi(x, y)$ from the seeing disk is problematic. Even assuming a uniform amplitude of light across the pupil, the inverse problem of computing the wavefront from the intensity pattern at the focal plane has no single solution. Instead, some light needs to be diverted by a beam-splitter into a WFS. These sensors are normally (but not uniquely) located on a pupil conjugate plane and can measure the wavefront either directly or indirectly. When the object they observe in order to measure the wavefront is a star, it is usually referred to as a natural guide star.

At the time of writing, the most commonly used WFS in astronomical AO is the Shack-Hartmann (SH) WFS. SHs are composed of a square array or grid of identical positive lenslets placed at the pupil plane. Each lenslet samples a small fraction of the total pupil or aperture, as shown in Figure 1.7. A region of the pupil sampled by a lenslet is called a subaperture. For any given subaperture, the light is focused onto a detector by a lenslet of focal length f , producing a spot. By measuring the centroid of the spot, the average local slope $\Delta\theta$ of the wavefront sample can be estimated. All of the local average slopes can be finally used to reconstruct




Figure 1.7: 2D section of a Shack-Hartmann WFS.




Figure 1.8: Effect of a deformable mirror on the phase of a wavefront. (a) Left - Incoming wavefront. (b) Right - Reflected wavefront. In this case the DM is not made of a continuous surface but of a segmented one.

the original wavefront, by several alternative methods, such as different forms of integration or by fitting modal functions.

Deformable mirror The most widely used wavefront corrector in AO is the deformable mirror (DM). A DM usually consists of a reflective metallic sheet which is pushed up or down in different locations by actuators. The whole reflective surface can either be continuous or segmented. In case of being continuous it is thin and flexible enough that it can deform into different shapes on command. If it is not, then every segment is disconnected from those adjacent to it and is controlled by one or more actuators. In either case, the fundamental principle with which they interact with a wavefront is the same and is shown in Figure 1.8. For both types of DM, the correction applied on the wavefront in units of length is

$$\phi_{cor} = 2\phi_{dm}, \quad (1.9)$$

where ϕ_{dm} is the position adopted by the mirror.

As seen previously, atmospheric turbulence imprints a similar ϕ_z to all wavelengths. Since both ϕ_z and ϕ_{dm} are in units of length, then a single DM can correct the WFE of all wavelengths at once. Due to several sources of noise present in AO systems, the correction applied by a DM is never perfect and leaves a residual error computed as

$$\phi_{res} = \phi_z - \phi_{cor}. \quad (1.10)$$

Finally, the principle describing the interaction between the mirror and the wavefront is also independent on the type of actuators. There are multiple types of actuators, such as micro electro-mechanical systems (MEMS) or piezoelectric ones. For all types of actuators, a first order linear approximation can be used to quickly and effectively map the relationship between the voltage u applied to them and the position of the mirror just above. This linear approximation, which does not consider other effects such as hysteresis, is

$$\phi_{dm} = G_{dm}u, \quad (1.11)$$

where G_{dm} is ideally a constant gain applicable to all actuators. In reality, the gain is slightly different for each actuator and requires calibration. For continuous mirrors, where a point can partially interact with all adjacent actuators, ϕ_{dm} and u can be regarded as vectors, while G_{dm} is a non-diagonal (although sparsely populated) matrix. More complex non-linear models exist, but they are generally more computationally expensive. As a consequence, they are more difficult to implement and will generally require more processing time, therefore increasing the time lag from the moment the desired shape of the DM has been computed and the moment that shape is implemented by sending voltage signals. For most AO systems it is preferable to use linear models, since the accuracy gained by using non-linear ones is not deemed sufficient to offset the resulting lag error, i.e. the difference between the passed shape of the wavefront currently being corrected for by the DM and the actual current shape of the wavefront.

1.1.5 Open-loop & Closed-loop configurations

Open-loop Having defined all the basic components of an AO system, it is now possible to describe the configurations in which they can be organised. This section only mentions two broad classes of configurations: open-loop and closed-loop.

After entering the telescope, the light of the star heads into the AO system. In an open-loop system, the first optical element in the light's path is a beamsplitter, as shown in Figure 1.9. This element, as its name suggests, splits the optical path into




Figure 1.9: Adaptive optics system in open-loop configuration.

two. In the first path is a WFS, which will measure in full the aberrations produced by the atmosphere. The measurement will be sent to a control computer which will estimate the appropriate actuator commands and transmit them to the DM. On the second path after the beamsplitter, the DM will apply the commands and relay the corrected beam to an imaging camera or to any other scientific instrument that requires it.

Open-loop configurations have a few disadvantages, which stem from the WFS's

inability to *see* the DM. The first one of these disadvantages is not being able to measure and map any biases in the correction. Not all WFSs and DMs are created equal. Some DM actuators respond differently to the same command, others might have a different initial offsets. On the other hand, WFS may also have biases across the pupil. Finally, there is the problem of optically aligning both to each other. If misaligned, the actuators will not correct the phase aberrations that should correspond to them.

The second disadvantage is the existence of time-variant phenomena, which can not be mapped beforehand. Instead, they need to be measured and compensated for in real-time. For example, depending on the technology, some actuators *creep*, i.e. their offsets changes in time. In turn, this can be the result of hysteresis, which happens when an actuator does not come back to its initial position after being commanded to do so. When this happens, the correction leaves a larger residual aberration, which could be corrected in the next iteration of the AO system, if only they could be *seen* by the WFS.

Despite their drawbacks, open-loop systems can be useful when trying to determine optical turbulence parameters from open-loop slopes, which would not be affected by the DM, or when there is no guide star in the area of sky to be corrected. The latter is the case when observing faint extended objects like galaxies. In these situations, the wavefront needs to be estimated from observing off-axis guide stars which will encounter a partially different turbulence. Closed-loop solutions have been devised that can make use of tomography to correct for a wide FOV that includes both the guide star and the scientific object, such as multi-conjugate AO (Beckers, 1988) and ground-layer AO (Rigaut et al., 2002). But these approaches have disadvantages of their own.

In the case of ground-layer AO, as the name suggests, only the turbulence produced at the ground-layer is corrected for. On one hand, this turbulence is common to all angles in the sky, which enables the wide field correction with a single DM. On the other hand, the ground-layer is not the only source of turbulence and several more

layers can appear at multiple altitudes. As a result, a fraction of the turbulence, which depends on the vertical profile of the turbulence and the angle between the guide star and the direction of correction (Tokovinin, 2004), can go uncorrected, resulting in a low overall performance of the system.

Multi-conjugate AO solves this by using several DMs, each conjugated to a different altitude. By doing so, it can achieve a relatively higher SR than ground-layer AO, all while correcting for a wide FOV. But in order to do so it requires larger DMs for higher layers. This becomes increasingly problematic as telescopes become larger themselves due to the Lagrange invariant. Multi-conjugate systems, such as MAORY (Diolaiti et al., 2017), can still be implemented in the next generation of ELTs, with diameters around and above 30 m. Unfortunately, MAORY's FOV remains limited to $75''$ due to, in part, technological constraints regarding the size of the DMs.

Finally, this is solved by moving into an open-loop configuration. Multi-object AO systems (Dekany et al., 2004; Hammer et al., 2004) use a multitude of small DMs, each to correct for a specific target in a small FOV. For comparison, the MOSAIC system which is being designed to operate on the same telescope as MAORY, can place small DMs over a FOV of $7.4'$ (Morris et al., 2018), a correctable solid angle almost 6 times larger. But since the DMs operate in open-loop, the system's biases and time-variant phenomena need to be well characterised if a relatively high degree of correction is to be attained.

The pros and cons of open-loop systems will therefore change depending on multiple variables such as the availability of guide stars, the number and altitude of turbulence layers, the size of the telescope, and the desired level of correction and corrected FOV. Because of this, depending on the circumstances, closed-loop systems will be favored over open-loop ones.

Closed-loop The main difference with an open-loop system is that a closed-loop one has the DM before the the beamsplitter, as shown in Figure 1.10. The imaging




Figure 1.10: Adaptive optics system in closed-loop configuration.

camera and the WFS are still in different paths, separated by the beamsplitter. But because the DM precedes it, its action is seen in both paths. As a consequence, the WFS does not receive the original atmospheric aberrations, but the residual aberrations left by the DM.

The main reason for sensing the residual aberrations rather than the open-loop wavefront is to make adjustments in order to reduce these residuals even further, hence achieving a greater level of correction. The existence of residual aberrations on the other hand can be attributed to three factors. The first one is sampling, both temporal and spatial. For limited samplings there will be high temporal and spatial frequencies that are not sampled and that will therefore not be corrected. In turn, sampling is limited in part by the second factor: noise.

Noise issues from a variety of sources such as the WFS's detector read-out noise or the photon-shot noise. Signal and noise are indistinguishable from each other on an individual frame. Furthermore, noise is uncorrelated from frame to frame and is unbiased, i.e. it averages to zero. Therefore, it will not compound into a larger residual error over multiple frames. As a result, the contributions to residual errors

from noise sources are similar in both open-loop and closed-loop systems.

The third source of residual aberrations are systematic biases. This is where having a closed-loop becomes an advantage. One example is, as will be seen later, AO systems can have differential aberrations between the imaging arm and the WFS known as non-common path aberrations (NCPAs). In the presence of this form of bias, an AO system can correct for all the aberrations being measured by the WFS and still leave aberrations on the imaging arm. Closed-loop systems can be calibrated to remove them by applying offsets. Without turbulence, if a set of actuator commands is somehow found to produce a good quality PSF, cancelling out any systematic aberrations in the imaging arm, then the response of the WFS can be recorded to determine the system's offset. These can be later set as the reference to which the control loop should converge once it is closed in the presence of turbulence, in order to recover the best possible PSF.

Besides offsets, closed-loop systems can also map the gain between the WFS and the DM. When an actuator is poked, the response of the WFS can be recorded into an interaction matrix. This matrix not only records the amplitude of the WFS's response relative to that of the actuator command, but also its spatial distribution, allowing it to account for actuator and subaperture heterogeneities and misalignments. Finally, any errors in the estimated gain value will lead to a bias in the AO correction which will compound in the form of a residual error if not sensed. A closed loop will be able to measure this growing residual error and will correct for it.

This section introduced some of the basic physical phenomena, optical performance metrics and technologies relevant to AO. The following section will explore a few of the main challenges driving the forefront of this field. These comprise problems that have already been briefly mentioned, such as the compensation of non-common path aberrations, as well as other that have not, because they are harder to directly draw from simple AO principles.

1.2 Current challenges

Astronomical AO technology will have to solve a wide spectrum of new problems, some of which stem from the large size of telescopes in the era of ELTs. These telescopes call for finer calibrations over larger degrees of freedom. Furthermore, they also enable new scientific cases with more stringent correction budgets. Compiling an exhaustive account of all of the challenges and problems facing astronomical AO would be too long and besides the point. Instead, the object of this section is to provide some of the context that will motivate the development of a new type of WFS. This will be done by showcasing three selected cases of interest that could be tackled by using the m-PDI. The three problems or challenges that can be addressed by a new WFS are the presence of quasi-static aberrations, the phasing of segmented apertures and achieving extreme adaptive optics, each of which is presented in its own subsection.

1.2.1 Quasi-static aberrations

Quasi-static aberrations are slowly changing aberrations introduced by the optics of the telescope, the AO system or the science instrument. These aberrations have a myriad of sources that include errors in the control law, small movements of the pupil due to elements like atmospheric dispersion correctors or the pupil derotator, and dilations of optical and mechanical components caused by temperature variations, amongst others. Some even have unknown origins which are difficult to untangle from the large number of simultaneously acting effects that drive them.

The multiplicity of effects driving them act on different timescales, making them difficult to characterise. For example, it was reported that at Gemini with Altair NIRI, such effects make the PSF evolve on timescales of 10-60 minutes (Marois et al., 2006). While on one hand short-term residual aberrations are quickly averaged out, on the other, slowly evolving quasi-static ones are more persistent. In

such cases, if they are not compensated by the AO system, they can cause speckles on the focal plane (Marois et al., 2000). When searching for exoplanets, special and difficult solutions need to be implemented to distinguish them from such speckles, or to remove the speckles altogether.

In AO systems, the WFS path and the science path will contain different optics and will therefore see different aberrations. These are NCPAs and they can be static or quasi-static, varying slowly over time. In systems such as the Gemini Planet Imager, Fusco et al. (2006) estimate that they would have to be corrected to less than 10 nm RMS within the first 100 modes in order to achieve direct exoplanet detection. This task is not straightforward and usually requires iterative methods.

The reason for this is the aberrations in the scientific path are invisible to the WFS. The aberrations in the common path leading up to the beam splitter, in the science path and in the WFS path can respectively be dubbed ϕ_A , ϕ_B and ϕ_C , as shown in Figure 1.11. The aberrations seen by the science path are $\phi_{(1)} = \phi_A + \phi_B$, while those seen by the WFS are $\phi_{(2)} = \phi_A + \phi_C$. When the control loop is closed, the aberrations seen by the WFS are sent to the DM to be subtracted from the system. While the resulting aberrations in (2) are $\phi_{(2')} = \phi_{(2)} - \phi_{(2)} = 0$, the remainder in the imaging path is $\phi_{(1')} = \phi_{(1)} - \phi_{(2)} = \phi_B - \phi_C$. This result is explained by two reasons. First, aberrations ϕ_B are not seen by the WFS and therefore can not be subtracted. Second, the WFS can also not distinguish between common aberrations ϕ_A and its own aberrations ϕ_C . As a result, by subtracting the whole of $\phi_{(2)}$ into the common path, it introduces the additive inverse of its own local aberrations.

When quasi-static aberrations can not be corrected by a WFS, iterative focal plane methods are often used. These methods such as Phase Diversity, the Simplex algorithm, Speckle Nulling and Angular Differential Imaging will be the subject of further discussion.

1.2.2 Phasing segmented apertures

In the past, most telescope primary mirrors used to be monolithic, i.e. they were build of a single piece of glass. But mirrors become harder to manufacture, handle and hold in the right shape when their size increases, up to the point when it is actually impossible to do so at the scale of ELTs (above 30 m in diameter). In order to solve this, primary mirrors can be segmented into a large number of smaller, disconnected mirrors. An example is shown in Figure 1.12. This is the first reason why a telescope aperture could be segmented. The production of small mirrors can be carried out in several preexisting facilities at once (no single giant factory has to be built just to then manufacture one single mirror), in a serialised way, and they can be transported without much interruption to roads and traffic. The above is already a problem for 8 m mirrors. The last advantage of segmented mirrors is spares can be kept on site for when they need to be serviced (re-polished and recoated), which with the 798 segments of ESO's ELT in Cerro Armazones will be constantly happening.

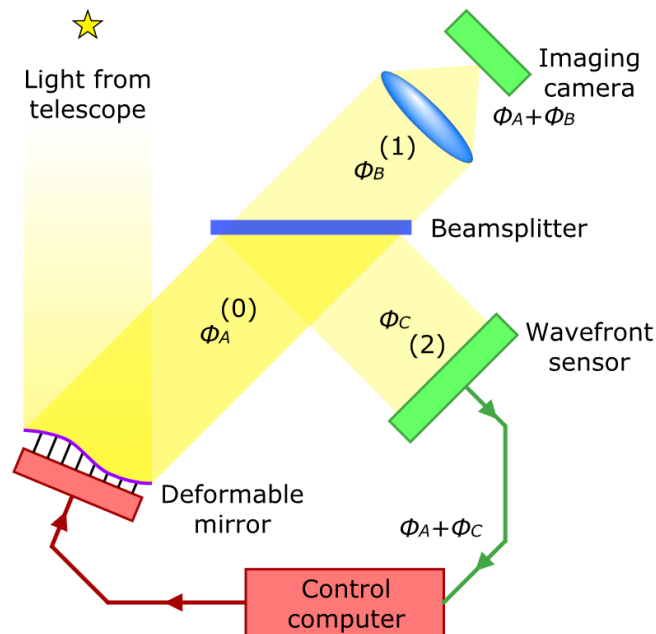


Figure 1.11: Non-common path aberrations. The common path, the imaging and WFS paths are respectively noted as (0), (1) and (2).




Figure 1.12: Approximation of the aperture of ESO’s ELT, with telescope spiders included. In reality the gaps between segments are smaller.

The optical problem with segmented mirrors is that, for them to optically behave as one monolithic piece, they all need to be in phase. This is to say they need to form a desired shape, often similar to a parabola. The act of putting them in phase is referred to as cophasing. The mirrors are directly cophased by the actuators of the stages holding them. The difficulty in cophasing the segments comes from WFSs either not being able to measure the differential piston term between them or from not having enough dynamic range. This will be discussed further in Section 2.2

As seen in Figure 1.12, the aperture is not only divided into multiple segments by the segmentation of the primary mirror, but also by the obstruction of the spiders. In smaller telescopes the projection of the spiders onto the pupil is below the spatial resolution of WFSs. The result is the wavefront can be partially measured, and estimated in the obscured area. But in the case of ELTs, spiders will have a much wider section, of 50 cm (Schwartz et al., 2018), as they need to hold a heavier secondary mirror over a longer distance. The slope of the wavefront can not be accurately sensed over wide gaps, as WFSs with sampling areas about the size of the spider beams or smaller do not coherently mix much light from one side of the beam’s obscuration with light from the other in the same sample (or same subaperture in the case of a Shack-Hartmann WFS) (N’Diaye et al., 2018). This makes it difficult to estimate the relative piston between two disconnected areas.

This is sometimes referred to as the island effect and it can be harder to solve than the cophasing problem between adjacent mirror segments. The reason for this is it is usually worsened by what is known as the low wind effect, which introduces local strong turbulence around the spiders (Sauvage et al., 2016). The impact of these phenomena is disastrous for exoplanet imaging as they occur during the best-suited observing conditions. High-contrast instruments have had a significant loss of observation time, of up to 20% in the case of the AO system SPHERE (Sauvage et al., 2016).

1.2.3 Extreme Adaptive Optics

In astronomy, when the level of correction of an AO system surpasses a given threshold, it is said to be an XAO system. Unfortunately there is no hard consensus on what the level of correction should be by definition. Instead, a system is deemed to be “extreme” if it enables high-contrast applications, such as high-contrast imaging for direct exoplanet detection.

Most specifications for XAO are indeed set by the requirements of exoplanet detection. The performance of instruments seeking to image exoplanets is usually defined by the level of contrast they can tolerate between the host star and its planet. For example, the detection limit of the SPHERE instrument at the Very Large Telescope (VLT) is 10^6 (i.e., 15 magnitudes between the planet and the star), with a goal around 10^8 (Fusco et al., 2006). This contrast needs to be achieved around a limited angular distance from the star. Again for SPHERE, the highest contrast is required at $0.5''$ from the star (for stars with magnitude <6 in the J band) and there needs to be access to small angular separations of between $0.1''$ to $3''$ (Beuzit et al., 2008).

In order to achieve this level of contrast, high-contrast imagers use coronagraphs in conjunction with XAO systems. The better the AO correction, the better the coronagraph extinction, which in turn leads to a reduction of flat-field and photon

noises and a decrease in the number and intensity of residual speckles. Unfortunately, each XAO planet imager has a complex error budget, with precise error allowances being allocated in different ways for each system design. As a result, there is no direct fixed relationship between contrast goals and subsequent AO requirements for the various existing systems. However, by examining the actual AO specifications and achievements, we can arrive at a practical working definition of XAO-level performance.

SAXO, SPHERE's AO system has a 41x41 actuator DM and a 40x40 SH WFS (Sauvage et al., 2010). Without turbulence and with an internal source, the maximum possible SR achieved after correcting for NCPAs was $99.0 \pm 0.1\%$ in H band (Fusco et al., 2014). Through Maréchal's approximation this is estimated to be a residual error of 25 ± 10 nm RMS. At nominal conditions with turbulence in the lab, a SR of $90.3 \pm 2\%$ or 81 ± 10 nm RMS of residual error were achieved at a wavelength of 1589 nm. Here, nominal conditions refer to a star of magnitude 9 producing 100 photo-e- per WFS subaperture per frame, at 1.2 kHz, a *seeing* of $0.85''$ and a wind speed of 12.5 m/s. Finally, preliminary on-sky results achieve a SR of 80%, 70% and 60% in the H , J and Y bands respectively and with an estimated *seeing* of $0.7''$. The star that was used, HD165054, has a magnitude of 6.9 and 7.2 in bands H and J . This is brighter than nominal conditions and should allow the system to perform as expected.

Equivalent high-contrast systems such as the Subaru Coronagraphic Extreme Adaptive Optics (SCEXAO) system or the Gemini Planet Imager (GPI) have similar constraints. After lab testing, SCEXAO expects to achieve a 1-hr long contrast between $2 \cdot 10^{-7}$ and $2 \cdot 10^{-8}$ over a few λ/D of separation (Guyon et al., 2010). On the other hand, during first light science observations GPI produced a $5 - \sigma$ contrast of 10^{-6} at $0.75''$ and 10^{-5} at $0.35''$ (Macintosh et al., 2014).

In terms of AO performance, on-sky observations of HD 36546 with SCEXAO achieved Strehl ratios of 70% – 80% in H band, for a *seeing* of $1.0''$ and a 12 m/s wind speed (Currie et al., 2017). This was attained by correcting for 1080 modes

at 2 kHz. For GPI, observations on β Pictoris ($I = 3.8$ mag.) yielded a SR of 89% in the H band (Macintosh et al., 2014). This was enough for the direct detection of the planet β Pictoris b. Despite *seeing* conditions not being provided, these results serve as a good benchmark to define XAO levels of correction for high-contrast imaging applications.

The aforementioned performances in terms of SR are summarised in Table 1.1. By no means is this to be considered an exhaustive summary of all existing XAO systems. Nonetheless, it manages to provide a soft definition of XAO at a glance, by presenting what different research groups have aimed for regarding the necessary level of correction.

Table 1.1: Summary: XAO benchmarks and their performance.

System	Strehl ratio in the H band
SAXO	80% <i>on-sky</i>
SCExAO	70%-80% <i>on-sky</i>
GPI	89% <i>on-sky</i>

Although it might seem from these working systems as if XAO was a solved problem, there are still many improvements to be made. The two main fronts where work is being done are extending the sky coverage and increasing the ultimate Strehl ratio. With respect to sky coverage, for some of the examples given above, nominal conditions are reached with stars of approximately magnitude 9 or brighter. Just by pushing this limit by one magnitude, the number of observable stars increases by about 288%, from 217,689 at magnitude 9.49 to 626,883 at magnitude 10.49 in the visible. With respect to increasing the Strehl ratio, this would lead to fewer residual speckles and a higher SNR when detecting exoplanets. The number of detectable exoplanets this would unlock is difficult to estimate and it might remain so until the necessary technological improvements are made and we launch ourselves into their discovery. But it is expected that both these fronts could benefit from more sensitive, accurate and faster WFSs. This would produce a decrease in the AO error budget, which is what increasing both the limiting magnitude and the

ultimate SR need.

This chapter has introduced the field of AO, laying out the context and problems that will serve as the incentives for the development of a new type of WFS. It has been done in two distinct steps. First, fundamental concepts in AO were introduced, such as the definition of wavefront error, models of atmospheric turbulence, metrics of optical performance, and the main components and potential configurations these systems can take. Second, once these concepts had been introduced, three important challenges facing the next generation of instruments were presented. These challenges are the characterisation and correction of quasi-static aberrations, the phasing of segmented apertures and the achievement and improvement of XAO.

The next chapter will introduce a potential solution to these problems: point-diffraction interferometers. The achievements of this family of interferometers will be compared to other solutions, in terms of both qualitative advantages and, mostly, in terms of quantitative results.

Point-diffraction interferometers & other solutions

Having overviewed some of the challenges facing the field of AO, the question arises of what the different approaches to solve them are and how they compare in terms of their potential and the actual progress made. The purpose of this chapter is then to introduce a kind of interferometer known as PDI, that has only recently started to be considered as a potential solution. The reasons for this are PDIs' main limitations: small chromatic bandwidth and dynamic range. By comparing the performance of PDIs to that of other existing solutions, their disadvantages and especially their benefits can be put into context in order to justify their development.

PDIs and the necessary principles to understand them are presented in Section 2.1. Section 2.2 then goes on to lay out the technological landscape into which they are to be considered. Some of the alternative methods focus on well established techniques such as traditional WFSs or algorithmic solutions such as phase diversity or speckle nulling. All of these are summarily explained and their advantages and disadvantages contrasted against those PDIs.

2.1 A potential approach: point-diffraction interferometers

Amongst the many approaches proposed to address the challenges described in Section 1.2, a class of WFS, the PDI, shows great promise. The main objective of this section is to present PDIs by introducing some basic definitions and exemplify them through two particular designs. First, Subsection 2.1.1, lists basic definitions surrounding general interferometers and those of the point-diffraction kind. The ultimate focus of this is to highlight the defining features of the point-diffraction configuration, and explain why these features allow it to be used in astronomical instrumentation, whereas other interferometers are either not well suited or outright impossible to use in this context.

Then Subsections 2.1.2 and 2.1.3 go on to present two of the most developed PDIs for AO. These are respectively the self-referencing Mach-Zehnder and the Zernike WFS. This serves the dual purpose of illustrating and putting into practice the previously introduced interferometry principles, as well as facilitating the comparison of benefits and achievements of these and other competing methods in the subsequent section.

2.1.1 Definitions

Intensity signal and interferogram As their name indicates, interferometers work by interfering two beams of light, therefore encoding information into an intensity signal. An ordered collection of intensity signals from multiple pixels forms an intensity pattern, also known as an interferogram. By measuring and analysing these signals and patterns, interferometers can precisely measure optical path differences between two beams.

Reference and test beams If one of the two beams has a “known” wavefront, whereas the other one has an unknown wavefront, then the interference pattern can be used to reconstruct the phase of the unknown beam, therefore acting as a WFS. In reality nothing can be known in absolute terms, but rather relative to some reference. This is why the “known” beam is instead dubbed the reference beam whereas the other is called the test beam.




Figure 2.1: Mach-Zehnder’s working principle, based on image by aurantiaca.

The Mach-Zehnder interferometer In many lab applications, such as metrology, the light source is internal. Both beams are produced together and then split so that one will go on to reflect on or get transmitted through a specimen, such as a surface, that needs to be measured, hence becoming the test beam. The beams can then easily be re-interfered due to their high temporal coherence. An example is the Mach-Zehnder interferometer, first proposed by Zehnder (1891), later refined by Mach (1892), and shown in Figure 2.1. When the test beam travels through the specimen, in this case a candle’s flame, variations in the index of refraction make light rays travel different optical path lengths. When they re-interfere, the intensity on the screens will be a function of the phase difference. In turn, the

intensity signal can be used to compute a physical property of the specimen, such as the index of refraction or temperature.

Self-referencing However, in other interferometers both the test beam and the specimen are external, such as with a star and atmospheric turbulence. In this case there is no possible temporal coherence between the external test beam and a reference beam generated with a local oscillator. Instead, in order to ensure coherence, some of the incoming light needs to be taxed and used in order to produce the reference beam. Such interferometers are called self-referencing.

Point-diffraction interferometry A possible way of achieving a self-referencing instrument is by using a small point-like feature, such as a pinhole. The diameter of the feature should roughly be the size of the core of a diffraction-limited PSF, around λ/D and $2\lambda/D$. When using such a pinhole, an interferometer is said to be a PDI. In these devices some of the light is focused onto the pinhole, producing a spherical wavefront at the exit. The spherical wavefront is then collimated into a flat one which will serve as the reference beam. The action of the pinhole can alternatively be described using Fourier optics, as that of a very low-pass filter that subtracts most high-frequency information from the wavefront. In doing so, most aberrations are removed, hence allowing for a highly flattened wavefront. More details can be found in Chapter 3. Finally, because the reference beam was created with a portion of light coming from the test beam, both beams are temporally coherent, allowing for interference to become visible.

In a historical note, the first version of a PDI, the Zernike phase-contrast test, was introduced by Zernike (1934). The theory was further generalized and the term coined by Smartt and Steel (1975). The difference between the Zernike phase-contrast test and Smartt's PDI is the first introduces a phase change in the pinhole, while the second uses an opaque film to reduce in different amounts the amplitude of the light transmitted through and around the pinhole.

Common-path interferometer In Mach-Zehnders and other interferometers such as Michelson’s, the reference and test beams propagate through different optical paths or “arms”, hence they are said to be non-common path. Since interferometers measure the optical path difference between both beams, vibrations, temperature dilations and any internal turbulence that affects both arms differently would register as wavefront aberrations. Since these variations would only have to be in the order of tens of nanometers in the visible to be well detected by the sensor, and many interferometers have nanometric accuracy targets, if not damped they represent an important source of bias and noise. A solution is to have both beams traverse the same optical path. In this case the interferometer is said to be common-path.

2.1.2 The self-referencing Mach-Zehnder

The regular Mach-Zehnder can be modified into becoming a self-referencing PDI, as shown in Figure 2.2. In this configuration, the test beam is free to propagate




Figure 2.2: PDI version of the Mach-Zehnder interferometer, based on image by aurantiaca.

through the instrument, while the reference beam is produced by filtering the aberrations with the pinhole. The main application of this modification is wavefront

sensing for astronomical AO, in which case the candle's flame is replaced by atmospheric turbulence. An other application is to characterise laser beams (Du et al., 2013; Feldman et al., 1991).

Just as in the original Mach-Zehnder, there are two interferograms, usually referred to as the symmetric and the anti-symmetric output. With equally balanced arms, the monochromatic outputs are respectively

$$I_1 = \frac{P^2}{2}(1 + \cos \phi) \quad (2.1)$$

and

$$I_2 = \frac{P^2}{2}(1 - \cos \phi), \quad (2.2)$$

where P is the amplitude of the electric field at the instrument's entrance pupil. The advantage of these two outputs is both the cosine of the phase ϕ and the amplitude P can be retrieved. On the one hand

$$\cos \phi = \frac{I_1 - I_2}{I_1 + I_2}, \quad (2.3)$$

and on the other

$$P^2 = I_1 + I_2. \quad (2.4)$$

The disadvantage of having two outputs is photons are spread over twice as many pixels, increasing the detection noise by $\sqrt{2}$.

2.1.3 The Zernike wavefront sensor

An example of a common-path interferometer is the Zernike WFS (ZWFS), based on the Zernike phase-contrast test. As seen in Figure 2.3, all of the light is transmitted through the same optical path, resulting in fewer optical components and making the instrument less prone to vibrations or internal turbulence induced aberrations. The central component of the ZWFS is a focal plane mask in the form of a glass plate with a small circular depression of depth d . The depression, also referred to as the phase disk, introduces a phase shift on the central part or core




Figure 2.3: Zernike WFS, where d is the depth of the transmissive focal plane phase mask.

of the star image. The phase shift is

$$\theta = \frac{2\pi}{\lambda}(n-1)d, \quad (2.5)$$

where λ is the wavelength and n is the glass' refractive index. This leads to two separate propagating electric fields, going through and outside of the phase disk. Instead of blocking high-frequency information as in the Mach-Zehnder, this technique just dissociates it from the core of the star image, hence producing both the reference and test beams without wasting any light. Both beams are later collimated by the same lens and made to interfere on a screen located on a conjugate pupil plane.

Given an entrance pupil-conjugate plane A , a focal plane B and an exit pupil-conjugate plane C , the resulting intensity signal at plane C is

$$I_C = P^2 + 2b^2(1 - \cos \theta) + 2Pb[\sin \theta \sin \phi - \cos \phi(1 - \cos \theta)], \quad (2.6)$$

where P is the amplitude of the electric field and b is another electric field amplitude which is a function of the phase disk's geometry. Glückstad and Mogensen (2001) suggest that for the generalised phase-contrast test using a circular pinhole, its radius should be smaller than 0.627 times the Airy radius. In the simulations presented by N'Diaye et al. (2013) the preferred phase shift and pinhole diameter

respectively are $\theta = \pi/2$ and $1.06\lambda/D$, resulting in $b \simeq 0.5P$. It is argued that these values lead to a quasi-linear intensity response. For small aberrations ϕ , the first-order Taylor expansion of the intensity is

$$I_C = P(\phi + 0.5). \quad (2.7)$$

An advantage of this formulation is that it allows for quick wavefront reconstruction. Unfortunately, this expression also reveals one of the fundamental restrictions of using this instrument as a WFS: the amplitude P must stay constant across the pupil. Different to the self-referencing Mach-Zehnder, scintillation, i.e. variations in P , is indistinguishable from wavefront aberrations ϕ and will contribute to the sensor's error budget.

Not being able to distinguish variations in amplitude comes with an advantage. The ZWFS is a highly specialised and efficient WFS. In fact, in terms of detection noise, which is comprised of photon-shot and read-out noises, this is the most optimal WFS of any existing kind (Guyon, 2005; N'Diaye et al., 2013). The detection noise of this sensor for the design specifications shown above is

$$\sigma_R^2 + \sigma_P^2 = \frac{\delta S_R^2}{I_0^2} + \frac{0.5}{I_0}, \quad (2.8)$$

where σ_R^2 and σ_P^2 are the power of the noise contributions of the read-out noise and the photon-shot noise, respectively, to the phase estimation error; S_R^2 is the power of the read-out noise and I_0 is the average light intensity at the entrance pupil. This high-degree of efficiency is achieved by not vignetting any light and by concentrating all of the light into a single output pupil.

2.2 Comparative analysis: the PDI and other alternative techniques

This section compares the performance of PDIs against that of other existing techniques, when addressing the challenges listed in Section 1.2. For any technique that

has not been presented already, only a summary explanation and bibliographic references will be provided. Instead, the main focus will be to provide simulation or experimental results in order to assess the strengths, deficiencies and prospects of different methods relative to each other. Such a compilation will allow us to identify those aspects or capabilities needing improvement in PDIs.

2.2.1 Quasi-static aberrations

All the different methods to deal with quasi-static aberrations can be grouped into two categories. The first category comprises iterative focal plane algorithms.

Phase diversity One of the earliest examples of iterative focal plane algorithms in AO is phase diversity, first proposed by Gonsalves (1982), which involves measuring the focal plane image before and after introducing some known aberration, or phase diversity, in order to estimate preexisting aberrations through an optimisation process. When used in conjunction with a coronagraph, it was shown by Sauvage et al. (2012) that the right mix of defocus and astigmatism enlarges the global minimum and pushes away local minima, increasing the chances of success and allowing for more accurate estimations.

An advanced application of this principle is the COFFEE (Paul et al., 2013) focal plane wavefront sensing method. When tested on SPHERE, it was shown to reach contrasts between 10^{-5} and 10^{-6} , with its best performance at angular separations between $9\lambda/D$ and $17\lambda/D$ (Paul et al., 2014). Given $\lambda = 1589$ nm and $D = 8.2$ m, these correspond to separations between $0.36''$ and $0.68''$. It is important to note that these results are also dependent on the SAXO closed loop which made the corrections with the estimations provided by COFFEE.

An advantage of this method over other iterative approaches is it only requires two images to successfully estimate quasi-statics both upstream and downstream of the coronagraph. Other methods, as will be shown later, can take up to a few

minutes to converge to a solution. Current disadvantages of this technique are it can not estimate amplitude aberrations, required to reach even higher levels of contrast, and the estimation of aberrations can only be done *offline*, prior to scientific observations. This last problems means that if quasi-static aberrations were to change during observations, this would go unnoticed.

Differential imaging Some methods can suppress speckles by computing the difference between two or multiple images and, in so doing, increase contrast for the detection of exoplanet companions. This is done so without having to estimate quasi-static aberrations. For example, angular differential imaging (ADI) (Marois et al., 2006) uses a series of images that are taken with an altitude/azimuth telescope, all while keeping the instrument rotator off at the Cassegrain focus, or adjusted at the Nasmyth, in order to keep the optics of the telescope and the instrument aligned. As a result, the FOV will rotate, but instrument-born speckles will not. If the images are simply added, any rotating feature in the FOV, such as a companion, will be smeared out, reinforcing static features and allowing one to reconstruct the instrument's PSF. This PSF can then be subtracted from all images, leaving behind short-lived turbulent speckles (born from atmospheric turbulence) and rotating features. Finally, when the resulting images are rotated to compensate for the FOV's rotation and then added, rotating features are instead reinforced and short-lived speckles are averaged out.

On-sky results show a speckle attenuation factor of between 20 and 100 (Marois et al., 2006). In terms of contrast, the best results are about 10^{-5} at an angular separation of $0.8''$ and as much as 10^{-8} at $8''$ (achieved with 225 images of Vega over a 2 hr period and good *seeing* conditions). The contrast improves with angular separation because the smearing effect also increases proportionally to it. This can be a disadvantage as the technique is optimised for relatively large separations, but it can also be an advantage as it can complement other methods that perform better at shorter separations.

ADI is sometimes complemented with other forms of differential imaging in order to achieve a better contrast at small angular separations. A common complement of ADI is spectral differential imaging (SDI), first conceived by Racine et al. (1999) and then studied by Marois et al. (2000). The combination of both techniques is called angular and spectral differential imaging (ASDI). SDI relies on the fact that exo-planets have large molecular features in their chromatic spectrum, while host stars have a relatively flat spectrum containing mostly atomic features. These features are absorption bands, the most prominent of which are caused in cool companions by atmospheric H₂O, CH₄ and NH₃ (Fortney et al., 2008). By simultaneously taking two or more images of a system at close wavelengths located around a sharp molecular feature and subtracting them, the contribution of the star can be partially suppressed, and the planet signal revealed. The ultimate contrast level achieved will depend on the abundance of these molecules in the exo-planet’s atmosphere and on the selection of pairs of chromatic filters.

When used by itself, SDI achieves similar contrasts to ADI at an angular separation of 0.8'', of 10^{-5} (Zurlo et al., 2014). Results at a separation of 8'' are not reported, perhaps because SDI does not have a strong advantage at large angular separations as ADI does. When combined, ASDI reaches contrasts of 10^{-7} at 0.8'' and 10^{-8} as close as 2.5''. This improvement prompts the question of whether, since differential imaging methods can be combined with each other, perhaps their final performance can be improved even further by the estimation and correction of quasi-static aberrations.

Speckle nulling One method that can either be used as stand-alone or in conjunction with differential imaging, but always after a coronagraph, is speckle nulling, first proposed by Malbet et al. (1995). By using a DM, this technique iteratively searches to create diffraction features that interfere destructively with speckles originating from starlight. Speckles are nulled in a limited region around the star creating a D-shaped high-contrast region, also known as a “dark hole”. This relies

on the self-coherence between the light of a speckle and the rest of the light from the star. The other side of the coin is that, because the light from two separate astronomical objects, e.g. a star and a planet, is incoherent, it is impossible to interfere them. Hence, if a small feature can not be nullified, then it is likely to be an separate object.

An advantage of this technique is it can act on very small angular separations and is only limited to a maximum angular separation. For example, in results at Palomar and Keck presented by Bottom et al. (2016) and obtained using a white light internal source, the best contrast, of about 10^{-5} , is achieved at $0.36''$. Since angular separations on the focal plane relate to spatial frequencies, the maximum angular separation where an anti-speckle can be generated is determined by the highest spatial frequency achievable by the DM. This is the sole variable defining the DM's control region on the focal plane and therefore the size of the "dark hole". Martinache et al. (2014) shows that with speckle nulling the PSF mean level over this region is lower by a factor of 2, and also more stable, with a standard deviation reduced by a factor of 3.

Another advantage is it can deal with short-lived speckles, characterised by a timescale of tens of minutes. Martinache et al. (2014) show that over the course of the twelve minutes covered by their speckle nulling control loop, the fluctuations on prominent image features were on the order of 25–50% of the local median value. At these timescales, and with such variations, they are invisible to ADI and manifest themselves in PSF subtracted images as speckle noise. Indeed, reported results suggest that the high-contrast detection limits achieved with ADI are set by long-lived aberrations, at the timescale of about 1 hour or less (Marois et al., 2006). This reenforces the idea that ADI and speckle nulling can not only complement each other in terms of angular separation, but also in terms of timescale and speckle lifetime.

Phase-shifting Zernike interferometer The second category of methods to tackle quasi-static aberrations comprises pupil plane WFSs. Different PDIs have been suggested to measure quasi-static aberrations. One version, the phase-shifting Zernike PDI, was proposed to correct for high-order quasi-statics in GPI (Wallace et al., 2010). This non-common path interferometer, for which a modified version is described in more detail in Wallace et al. (2011), works by adding a series of piston shifts to the reference beam, recording an interferogram in each case. All interferograms are required to reconstruct the electric field’s phase and amplitude. This instrument is not to be confused with the common-path phase-shifting interferometer presented by Medecky et al. (1996). An advantage of this WFS over all other previously mentioned techniques is that, despite having to acquire multiple interferograms, the proposed configuration allows for *online* measurements. These are measurements taken during on-sky operation. The tradeoff for this to be possible is that light has to be diverted into the PDI with a beamsplitter, incurring some amount of NCPAs that, even if completely removed with an iterative method, could later drift into an unknown value.

In terms of accuracy, lab tests against a Zygo interferometer show an agreement of 4.5 nm RMS between them, when measuring the phase errors in a microscope slide. About the precision, it achieves 1 nm RMS and 12 nm RMS of error when integrating for 1 minute with stars of magnitude 5 and 8 respectively. It is estimated this performance will allow GPI’s coronagraph to reach its 10^{-7} target contrast, but at the current time of writing no test results on GPI were found. The only results found were correcting low-order aberrations with a SH (Hartung et al., 2014). It was also not indicated what the chromatic bandwidth or the dynamic range of the PDI are or will be.

Zernike WFS The best results of a PDI correcting a real astronomical instrument’s quasi-statics were present by N’Diaye et al. (2016), for the SPHERE hosted tests of their instrument ZELDA, a Zernike WFS (N’Diaye et al., 2013). For the

experiment, the Zernike phase mask was put in the place of the coronagraphic mask, the instrument was set to pupil imaging mode, and the system was illuminated using an internal source. The advantage of this setup is that, by being right in the science path, ZELDA minimises the number of NCPAs with the rest of the coronagraph and science instrument. The disadvantage is that, unless ZELDA is put in a separate path like the phase-shifting PDI, only *offline* measurements can be performed, i.e. before observations, a calibration step would have to be undergone in order to set the offsets of the system and correct quasi-static aberrations. If any of these aberrations drifted during observations, except for tip and tilt, they would go unnoticed.

Experiments were divided into two sets, all of them using an internal source and no turbulence. The first set attempts to determine ZELDA's accuracy when measuring known aberrations. The input aberrations were in the form of both low-order Zernike modes, up to spherical, and Fourier spatial frequencies in the X and Y directions, up to 10 cycles per pupil. The measurements were compared against simulated predictions of the instrument's response. Results show that for most modes and spatial frequencies (results are not presented for all tests) at amplitudes under 100 nm RMS, the accuracy remains beneath 5 nm RMS, i.e. accuracy is nanometric.

Most tests were carried out at a central wavelength $\lambda_c = 1642$ nm, with a chromatic bandwidth $\Delta\lambda = 24$ nm (relative spectral bandwidth of $\Delta\lambda = 1.46\%$). This small spectral bandwidth is to be expected when working with most interferometers. But one small test, done with a chromatic bandwidth of $\Delta\lambda = 290$ nm (17.7%) and only for one spatial frequency, shows that the accuracy remains under 5 nm RMS. This suggests the chromatic bandwidth could be somewhat extended without detriment to the instrument's accuracy, allowing for higher throughputs and consequentially for a potentially higher SNR.

The second set of tests quantifies the gains in contrast on the coronagraph's focal plane when using the instrument to correct for NCPAs. At their best, there was a

contrast gain of 10 in the coronagraphic image at $0.2''$, almost reaching the coronagraph's theoretical (without aberrations) contrast limit of 10^{-6} between $0.2''$ and $0.4''$ of separation. These numbers are not translated into estimated residual aberrations and so cannot easily be projected onto other applications or science cases. Nonetheless, an initial comparison with phase-diversity is possible through the results presented by COFFEE. Although results are of the same order of magnitude, additional studies are required to get a decisive verdict.

Summary Major advantages and disadvantages for all techniques are summarised in Table 2.1. The table does not comprehensively include all possible attributes such as dynamic range, chromatic bandwidth and compatibility with other techniques, as there is not good information available for all methods. Instead, the compilation considers crucial attributes for which there is sufficient data for most approaches. As can be seen in the table, there is a wide variety of angular separa-

Table 2.1: Summary: Approaches to reduce quasi-static aberrations.

Method	Achieved Contrast	Optimal Angular Separation	On-line	Verification
Phase diversity	$10^{-5} - 10^{-6}$	$0.36'' - 0.68''$	No	Telescope, internal source
ADI	10^{-8}	$8''$	No	On-sky
SDI	10^{-5}	$0.8''$	No	On-sky
ASDI	10^{-8}	$2.5''$	No	On-sky
Speckle nulling	10^{-5}	$0.36''$	No	Telescope, internal source
Phase-shifting Zernike interferometer	10^{-7}	No information	Yes	Testbed
Zernike WFS	10^{-6}	$0.2'' - 0.4''$	No	Telescope, internal source

tions at which these techniques perform best. It is important to notice that most of these results, with the exception of differential imaging methods, were obtained

without atmospheric turbulence. Otherwise, they can be difficult to compare, as their ultimate performance will depend on the *seeing* conditions, as well as on the AO system they are coupled to; both can vary greatly and will pull the performance down. In the case of differential imaging methods, the ultimate performance on-sky is already given, as this technique is not dependent on an AO system and is easy to test on-sky without any major hardware modifications.

Finally, since they span different optimal angular separations and timescales, the conclusion is most of these methods can be used in conjunction and their development and exploration must therefore continue in parallel.

2.2.2 Phasing segmented apertures

Approaches dealing with cophasing fall largely into the category of pupil plane WFSs. So far, the only focal plane approach uses phase diversity.

Single-image phase diversity As described previously in Section 2.2.1, phase aberrations can be estimated from a pair of focal plane images by introducing a known aberration, or diversity, in one of them. Furthermore, the same publication (Gonsalves, 1982) which introduces phase diversity also includes another concept: phase retrieval. The difference is that, whereas phase diversity can be used to estimate both the aberrations and the object if given two images, phase retrieval estimates aberrations with only one image, provided the object is known. Because of this, phase retrieval is also referred to as single-image phase diversity. By working under the assumption that images can be acquired near-instantaneously, i.e. with “frozen” turbulence (not to be confused with the frozen flow approximation Guesalaga et al. (2014)), Lamb et al. (2017) manage to use phase retrieval to estimate phase segment piston errors in simulated images. Simulation conditions emulated VLT-SPHERE and Keck-NIRC2, with a defocus term of 20 nm peak-to-valley (PtV), as found in SPHERE’s Differential Tip-Tilt Sensor (DTTS). Results show that a 153 nm RMS cophasing WFE can be estimated within an error of

29 nm RMS. Assuming a perfect correction this would represent a SR increase of 12% for near infrared NIRC2 images.

One advantage of phase diversity and phase retrieval is that very modest hardware modifications are required. In the case of phase diversity, either the imager is mounted on a translating mount to defocus the image, or a second imager (such as DTTS) can be installed. For phase retrieval no modifications are required. Furthermore, by using a single image, the problem of turbulence evolution between images is circumvented. Finally, the technique can be enhanced by using non-simple pupils (Lamb et al., 2016), such as those found in coronagraphs or segmented telescopes.

WFSs The first example of a pupil plane WFS being used for cophasing is found at the 36-segment Keck telescope, with the use of a modified SH. The concept, proposed by Chanan et al. (1998), places circular subapertures on the intersegment edges. But instead of measuring the spots' centroid, the proposed algorithms analyse the resulting diffraction pattern. Using filters with a variety of coherence lengths, this approach has a maximum dynamic range of $\pm 30 \mu\text{m}$ and can reduce the piston error down to 30 nm. The phasing was done at intervals of around four weeks and took between half an hour and up to two hours using relatively bright guide stars.

The success of this approach prompted the development of further wavefront sensors, notably PDIs. Notaras and Paterson (2007) was able to demonstrate that under weak aberrations and unlike gradient based sensors using conventional least-squares reconstructors, PDIs are robust against strong scintillation and optical vortices. As a consequence, they are well suited to measuring random piston and tip-tilt differences between telescope segments (Janin-Potiron et al., 2017; Yaitskova et al., 2005).

Further research focused on solving PDIs short dynamic range. The two main interferometers presented so far, the Mach-Zehnder and the ZWFS, have an *a priori*

or *base* dynamic range of $\pm\lambda/4$ or $\pm\pi/2$ in terms of phase. This is insufficient when trying to correct for piston errors in the range of microns to tens of microns. An approach called the multi-wavelength scheme was hence devised by Vigan et al. (2011), that uses a pair of narrow bandpass chromatic filters with different central wavelengths to extend a PDI's dynamic range. By using the relationship between both phases, the piston error was extended to a range of $\pm 4.0\ \mu\text{m}$.

This was developed and tested in the context of ESO's Active Phasing Experiment (Gonté et al., 2009), set out to compare the performance of four WFSs when phasing segmented mirrors at the VLT. The instruments tested were SHAPS (Mazzoleni et al., 2008), an improved version of the SH presented by Chanan et al. (1998), DIPSI (Chueca et al., 2008), a curvature WFS (Roddier, 1988), PYPS (Pinna et al., 2008), a pyramid WFS (Ragazzoni, 1996) and ZEUS (Dohlen et al., 2006; Surdej et al., 2008), a ZWFS that employs the multi-wavelength scheme.

Final closed-loop residual errors for SHAPS, DIPSI, PYPS and ZEUS were 9.2 nm RMS, 31 nm RMS, 23 nm RMS and 6.7 nm RMS respectively (Gonté et al., 2009; Surdej et al., 2010). In addition, most wavefront sensors managed to converge with star magnitudes between 14 and 18 and for initial WFEs of 55 nm RMS for SHAPS, 550 nm RMS for DIPSI and 45 nm RMS for PYPS and about 50 nm RMS for ZEUS. These results should not be considered as the best possible values, since as it was shown earlier, methods like the multi-wavelength scheme can greatly increase an instrument's dynamic range. Instead, they should be taken as worst case scenarios and the starting points for further developments.

Summary Major advantages and disadvantages for all techniques are summarised in Table 2.2. The table does not comprehensively include all possible attributes such as required seeing conditions and star magnitudes, as there is not good information available for all methods. Instead, the compilation considers crucial attributes for which there is sufficient data for most approaches. As can be seen in the table SHs and PDIs outperform all other methods by between one and two or-

Table 2.2: Summary: Approaches for segment phasing.

Method	Residual Piston Error	Dynamic Range	Verification
Single-image phase diversity	29 nm RMS	153 nm RMS	On-sky
SH	9.2 nm RMS	$\pm 30 \mu\text{m}$	On-sky
Curvature WFS	31 nm RMS	550 nm RMS	On-sky
Pyramid WFS	23 nm RMS	50 nm RMS	On-sky
PDI	6.7 nm RMS	$\pm 4.0 \mu\text{m}$	On-sky

ders of magnitude in both residual error and dynamic range. The advantage of the latter over the former is that it remains non-pupil specific, i.e. a PDI does not have to be designed to comply with a specific pupil's segment geometry, whereas a modified SH is built to match the intersegments of a given telescope. As a consequence the same PDI could be used to perform tasks other than phasing.

2.2.3 Extreme Adaptive Optics

Currently there are five XAO systems implemented on telescopes that can achieve high Strehl ratios: SPHERE's SAXO, SCEXAO, GPI, PALM-3000 (Dekany et al., 2013) and FLAO (Esposito et al., 2010). The on-sky performance with turbulence was already presented in Section 1.2.3 for the first 3 systems, but nothing was said about the technologies being used. This section introduces the core standard AO configuration found in most systems, using SAXO as the main example. It then presents each system's main WFS, primarily used to achieve XAO, as well as performance results. Finally, potential reasons are provided as to why PDIs remain unrepresented in XAO and what improvements are required to change this predicament.

XAO systems are based fundamentally on the standard single-conjugate adaptive optics (SCAO) configuration, but with some additional parts. This is the most basic closed-loop configuration. It was presented in Section 1.1.5 and shown in Figure 1.10, but it was not specifically categorised as SCAO. The name comes from

the fact that the main high-order DM and the high-order WFS are both conjugated to the same altitude, usually the ground layer, without there being other WFSs or DMs conjugated to other altitudes. In these system the AO correction is centered around a single guide star. The SAXO system (Fusco et al., 2014), described in Figure 2.4, is an example of an XAO system with a SCAO configuration at its core. The system has 4 acting devices and 2 sensors. The acting devices are a pupil tip-

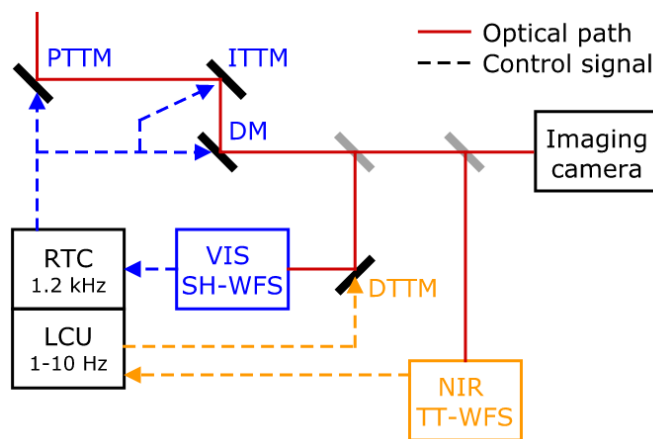


Figure 2.4: Schematic presentation of SAXO. Adapted from Fusco et al. (2014).

tilt mirror (PTTM), a fast image tip-tilt mirror (ITTM), a 41-by-41 DM and a differential tip-tilt mirror (DTTM) for fine centering on the SH-WFS. The last of these is controlled using measurements from an NIR tip-tilt sensor, while the first three ones are controlled using information provided by a visible high-order WFS, in this case a 40-by-40 SH. The signals produced by the sensors are sent to one of two computers, the real-time controller (RTC) which handles fast control loops in the order of kilohertz, and the telescope’s local control unit (LCU) in charge of slow loops in the order hertz.

This is a SCAO configuration in that the bulk of the AO corrections are handled by the high-order WFS, the ITTM and the DM, all of which observe and correct for the same guide star. The PTTM, the DTTM and the NIR tip-tilt sensor are added features to enable coronagraphic capabilities after the AO on the science path, where the imaging camera is located.

What turns this and, indeed, any other system from being a conventional SCAO system into an XAO one is the heightened level of correction. This quantitative rather than qualitative difference is enabled by operating at higher frequencies, both temporally and spatially. For example, SAXO's high-order WFS can run at >1.3 kHz with relatively small 20.5 cm subapertures at the telescope pupil. Having small subapertures allows one to sample spatially small aberrations, as well as the larger ones. And smaller aberrations vary faster, which is what prompts higher frame-rates. Both, less collecting surface and higher frame-rate, diminish the number of photons collected, hence decreasing the signal intensity. In order to compensate for this and maintain good levels of SNR, the WFS's detector must have very low levels of read-out noise, in this case <0.1 e-.

Most XAO systems have similar subaperture sizes, around 20 cm or smaller, and can reach similar frame-rates, around 1 kilohertz. The differentiating aspect between them tends to be the type of high-order WFS they use. In a similar way to SAXO, other instruments use SH-WFSs. GPI uses a 43-by-43 SH in the visible with which they can achieve 89% SR in the H band, with atmospheric turbulence on-sky (Macintosh et al., 2014). This translate to a $5-\sigma$ coronagraphic contrast of 10^{-6} at $0.75''$ and 10^{-5} at $0.35''$. The PALM-3000 system at the 5.1 m Hale telescope uses a reconfigurable SH, which can either have 64, 32, 16 or 8 subapertures across the pupil, depending on the guide star's brightness. In bright conditions and using the 64-by-64 configuration it is able to reach around 85% SR in the K band for stars brighter than magnitude 8, on-sky (Dekany et al., 2013). For comparison, this would translate to a 75% SR in the H band.

SCEXAO contemplated the possibility of using a non-linear curvature WFS (Guyon et al., 2010), but instead a pyramid WFS was chosen. No curvature WFS can be found at the heart of any XAO system, whereas pyramids have become increasingly more ubiquitous, due in part to their flat frequency response (Ragazzoni and Farinato, 1999). SCEXAO's pyramid sensor can work at 1.5-3.5 kHz (Jovanovic et al., 2016), with a read-out noise of 1 e- at 1.7 kHz (Guyon et al., 2011), and reach 70%-

Table 2.3: Summary: XAO systems, their WFSs and their performances

System	WFS	Strehl ratio in the H band
SAXO	40-by-40 SH	80% <i>on-sky</i>
SCExAO	PWFS	70%-80% <i>on-sky</i>
GPI	43-by-43 SH	89% <i>on-sky</i>
PALM-3000	64, 32, 16, 8 SH	75% <i>on-sky</i>
FLAO	30-by-30 PWFS	>80% <i>on-sky</i>

80% SR in the H band on-sky (Currie et al., 2017). The spatial resolution of this pyramid could not be found, presumably because the resolution of a pyramid can be changed through several means, such as modulation, resizing of the plate scale or binning of the pixels. Nonetheless, it is sometimes still provided, as in the case of the First Light LBT AO (FLAO) system’s 30-by-30 pyramid. With this setup it was able to produce an 83% SR in the H band in lab experiments (Esposito et al., 2010), above 80% *on-sky* (Esposito et al., 2011), and a $2.3 \cdot 10^{-5}$ contrast at $0.235''$ of separation (Skemer et al., 2012).

The performance of these systems is summarised in Table 2.3. As the error budget of large AO systems depends upon many variables, it is difficult to draw a clear conclusion from these results as to which type of WFS, the pyramid or the SH, is better. On the positive side, the decision of which to choose does not seem to be crucial, at least from this narrow point of view.

But why not PDIs? An important question remains. Why, if PDIs have shown promise when dealing with quasi-static aberrations and phasing telescope segments, do they remain unused in XAO applications? The reasons have to do with two of the biggest limitations facing interferometers, and specially point-diffraction ones. These are chromatic bandwidth on the one hand, and dynamic range on the other.

Chromatic bandwidth Historically, interferometers of most kinds have been designed to work with lasers, which could provide stable and temporally coherent

reference beams for applications ranging from metrology to physics experiments. In these applications, having the narrowest and most stable chromatic bandwidth possible was an advantage when trying to measure the smallest defects on an optical surface, or in other high precision measurements. But for AO, limiting the bandwidth reduces the sensor's throughput and therefore reduces the SNR. A WFS that can accept a bandwidth of 400 nm, or most of the visible spectrum, receives 13 times more light than a PDI with a bandwidth of 30 nm, or a 2.7 magnitude increase.

There are three effects that force PDIs to be narrowband. These effects have to do with the optical path difference (OPD) between the reference and the test beams, the dependency between phase and interferogram, and the size of the pinhole relative to the point-diffraction PSF.

The OPD problem PDIs that were originally conceived for optical metrology, like the ZWFS, are meant to be illuminated with monochromatic light. In this case, for the intensity at the exit pupil to be proportional to the sine of the phase aberrations, the OPD introduced at the focal plane mask must be exactly $\pm\lambda/4$ for all wavelengths. In turn, this required a material with an index of refraction proportional to the wavelength. If the OPD equals $\pm\lambda/2$, then the signal is proportional to the cosine of the phase and therefore does not yield any sign information. In the worst possible case, the OPD equals $\pm\lambda$ and there is no interference. In order to avoid this problem, self-referencing interferometers must produce a reference beam without relying on a phase shift. This is the main improvement to be made on PDIs, as will be shown on Chapter 3.

Phase dependency The next chromatic effect is common to all interferometers when measuring atmospheric distortions and can not be fixed. When two beams of light interfere, the resulting interferogram is a function of the phase differences between the reference and the test beams. But as seen earlier, atmospheric wavefront

distortions are similar for wavelengths in the visible and NIR when measured in units of length, i.e. they are inversely proportional to wavelength when measured in phase. This turns out to be an advantage for WFSs that measure slope, such as SHs and pyramids. Since the slopes are proportional to wavefront distortions in meters, then, in principle, these WFSs have an unlimited chromatic bandwidth, which is only limited by the real optics of the instrument. But in the case of interferometers, since they measure phase, every wavelength will produce signals of different magnitudes. These signals are difficult to disentangle by traditional silicon chips as they can only count photons, but can not discern their wavelengths.

Relative pinhole size The last chromatic effect is common to all PDIs. As seen earlier in Section 1.1, the size of a diffraction-limited PSF is proportional to the wavelength. This means that the pinhole in a PDI has a different relative size for different wavelengths. As a result, the reference beam will have a different amplitude and flatness for each wavelength. In order to avoid this, self-referencing interferometers would have to produce the reference beam by using something other than a pinhole. But as will be seen in the next chapter, the chromatic bandwidth of a PDI can be increased without having to fix this problem or that of the phase dependency.

Dynamic range The other big hurdle keeping PDIs from being used in XAO is that of usually having a limited dynamic range. Some interferometers have a dynamic range of $\pm\pi$, while others such as the ZWFS and Mach-Zehnders have a dynamic range of $\pm\pi/2$. In the case of the first, when aberrations exceed $\pm\lambda/2$, traditional unwrapping is possible. This allows for an extended dynamic range.

But most self-referencing PDIs fall in the second group. In this case, conventional unwrapping techniques can not be used, which is a problem when trying to close an AO control loop. As shown in Figure 2.5, uncorrected atmospheric aberrations can be on the order of $\pm 1 - 2 \mu\text{m}$. For a PDI with a range of $\pm\lambda/2$ working in




Figure 2.5: Kolmogorov phase screens for (*left*) good *seeing* and (*right*) for bad *seeing*. Phase screens are 4.2 m wide.

the R band, this is about 9 times bigger than its range. This could be partially addressed by increasing the WFS's working wavelength, for instance into the NIR range. But even then, for example when working in the Y band, aberrations would still be about 6 times greater than the dynamic range. In order to solve this problem, an interferometer must allow for easy unwrapping, i.e. have a dynamic range of $\pm\pi$.

2.3 Summary

This chapter introduced PDIs and how they fit into the AO landscape. Their advantages and disadvantages were compared to those of other technologies when dealing with some of the major challenges facing the field. PDIs represent a realistic approach to these problems, specially if some upgrades are brought about that can enhance some of their characteristics. Two of the major aspects needing improvement are their chromatic bandwidth and their dynamic range. The following chapter presents this thesis' proposal to achieve these improvements: the m-PDI. Here, the working principle of this new concept is introduced, and its performance is examined both analytically and in simulation.

The pupil-modulated Point-Diffraction Interferometer

This chapter further develops the theory for a PDI first presented by Bharmal et al. (2012). This approach is an improved version of the phase-shifting PDI (PS-PDI) presented by Medeckı et al. (1996). This technique is not to be confused with the phase-shifting PDI later introduced by Wallace et al. (2011). The latter is similar to a Zernike sensor in that it adds a piston phase to the core of the PSF in order to produce the reference beam. In Medeckı's approach, the phase-shift is introduced as a tilt at the pupil with a small-angle beam-splitter. Because the tilt produces a modulation of the electric field in the pupil, we have decided instead to refer to our method as the pupil-modulated PDI (m-PDI). An introduction to the working principle of the modified m-PDI is then followed by the mathematical formalism. The mathematical model is used to derive the propagation of sources of error, such as photon noise, read-out noise and chromatic effect. The linearity and the dynamic range is then tested against variations in multiple parameters.

In Section 3.1 we present the concept and mathematical formalism of the improved m-PDI, later deriving from it the propagation of photon-shot noise and read-out noise in Section 3.3. The effects of Strehl and of chromatic bandwidth on the visibility of fringes is later studied in Section 3.4. A sensitivity analysis of the

sensor's linearity and dynamic range are carried out in Section 3.5. Finally, in Section 3.6, the sensor is tested in a simulation of a real system and the sensitivity is presented as the exposure time required to achieve a given signal-to-noise ratio.

3.1 Concept

3.1.1 Principle

The scheme of the m-PDI is given in Fig. 3.1. At the entrance pupil, an aberrated wavefront goes through a small-angle beam splitter, which in this case is a grating. The beam is then split into modes, some of which will later be interfered at the exit pupil. The central mode (i.e. mode 0) goes through a narrow pinhole in the focal plane filter mask. This filters out higher frequencies leaving a flat beam which will be used as the reference beam. As shown in Section 3.5, the size of the pinhole and, thus, the frequencies left in the reference beam have an important effect on the sensor's linearity, sensitivity and dynamic range. Mode +1 goes through a larger aperture which filters enough frequencies to avoid aliasing. This is important to produce an unambiguous test beam, as will be explained in Section 3.1.2. Both beams later interfere at the exit pupil producing fringes. The interference fringes are an image of the line-pairs in the grating, which, as it will be shown later, are modulated by both the phase of the electric field and its amplitude.

In Medeck's PDI the apertures in the focal plane filter mask are swapped. Mode 0 goes through the large aperture while mode +1 goes through the pinhole. In this configuration, since mode +1 carries less light than the central mode and is being filtered, the reference beam has less light than the test beam. This reduces the visibility of the fringes. Another disadvantage is that the pinhole filters light by wavelength, letting through only a narrow bandwidth. By letting mode +1 go through a larger aperture, more wavelengths can go through. This is a necessary feature in astronomical AO because of the limited amount of light available.




Figure 3.1: Layout of the pupil-modulated point-diffraction interferometer. The grating splits the beam into modes. Mode 0, shown in a solid red line, goes through a point-diffraction pinhole in the focal plane. Mode +1 goes through a larger aperture.

3.1.2 Formalism

For an aberrated wavefront, the electric field

$$\Psi_0 = P e^{i\phi} = P_0 (1 - \epsilon) e^{i\phi}, \quad (3.1)$$

where P is the amplitude of the electric field, P_0 is its average across the pupil, ϵ is a zero-mean function describing the local scintillation and ϕ is the phase. The function P is considered to be 0 outside of the telescope pupil, hence defining the telescope's aperture shape. In this case, the small-angle beam splitter will be considered to be a grating. The grating is located on Plane A, as defined in Fig. 3.1, and is described by a square wave function G_T of period T . This function alternates between the value 0 representing an obstruction and 1 representing full transmission. For simplicity, the function will be considered to be even, so its description in the Fourier domain is also even and real. The wavefront after the grating is described by

$$\Psi_A = \Psi_0 \cdot G_T. \quad (3.2)$$

Just before the focal plane filter mask on Plane B, the wavefront is described by

$$\hat{\Psi}_A = \hat{\Psi}_0 \otimes \hat{G}_T, \quad (3.3)$$

where \widehat{A} is the Fourier transform $\mathcal{F}[A]$ of a A and the \otimes symbol represents the convolution operation. The hat notation and \mathcal{F} will also include the Fourier optics scaling factor $1/\lambda f$, where λ is the wavelength and f is the lens' focal length. With these considerations taken into account, the Fourier transform of G_T is

$$\begin{aligned} \widehat{G}_T(k, \zeta) = & \frac{1}{2} \delta(k, \zeta) \\ & + \frac{1}{\pi} \sum_{m=1,2,\dots}^{\infty} \frac{1}{2m-1} \left[\delta \left(k - \frac{\lambda f(2m-1)}{T}, \zeta \right) + \delta \left(k + \frac{\lambda f(2m-1)}{T}, \zeta \right) \right], \end{aligned} \quad (3.4)$$

where $\delta(k, \zeta)$ is a Dirac delta function, k and ζ are Cartesian position coordinates in the focal plane in meters and the indices $m = 1, 2, \dots$ represent the modes produced by the grating. Since the dispersion modes produced by the grating spread parallel to the axis described by k , ζ will not be included in further equations unless necessary.

At Plane B the wavefront is multiplied by the focal plane filter mask

$$M = M_0 + M_{+1}, \quad (3.5)$$

where M_0 is a circular top-hat functions equal to 1 for $|(k, \zeta)| < D_{B,0}/2$, M_{+1} is a square top-hat function equal to 1 for $|k - \lambda_0 f/T| < D_{B,+1}/2 \cap |\zeta| < D_{B,+1}/2$, 0 elsewhere, $D_{B,0}$ and $D_{B,+1}$ denote their respective diameter and side, and λ_0 is the central wavelength around which the instrument is designed. The aperture's geometry is presented in Fig. 3.2. The central wavelength λ_0 is the wavelength that goes right through the center of the large aperture M_{+1} . When the instrument is fed with polychromatic light, λ_0 may not be the same as λ_c , which is the wavelength at the center of the light's spectrum. Contrary to M_0 , M_{+1} can be either a circle, a rectangle or any other shape. The square aperture configuration lets through the same spatial frequencies in both the X and the Y direction, allowing more spatial frequencies through than a circular aperture. The axes k and ζ in the focal plane are respectively parallel to X and Y in the pupil planes.




Figure 3.2: Focal plane filter mask, in Plane B and with a square large aperture M_{+1} .

After the mask, the electric field is

$$\Psi_B = \widehat{\Psi}_A M = \frac{1}{2} \widehat{\Psi}_0(k) M_0 + \frac{1}{\pi} \widehat{\Psi}_0 \left(k - \frac{\lambda f}{T} \right) M_{+1}. \quad (3.6)$$

When $D_{B,0}$ is in the vicinity of $\lambda f / D_A$ or smaller, where D_A is the diameter of the entrance pupil, then the first term of Eq. (3.6) can be considered a point-diffraction source

$$\Psi_{B,0}(k) = \frac{1}{2} \widehat{\Psi}_0(k) M_0 \simeq b \delta(k) \quad (3.7)$$

of amplitude b . This assumption, that the pinhole can be approximated to an infinitesimally small point-diffraction source, is the only assumption of this monochromatic model. The model will be later compared to a simulation which is not built on this or any other assumptions. Regarding the second term of Eq. 3.6, M_{+1} acts as a low-pass filter on $\widehat{\Psi}_0$, leaving the equation as

$$\Psi_B \simeq \Psi_{B,0}(k) + \frac{1}{\pi} \widehat{\Psi}_{LP} \left(k - \frac{\lambda f}{T} \right), \quad (3.8)$$

where Ψ_{LP} has been low-pass filtered by M_{+1} .

The light is then propagated into Plane C, where the detector reads the intensity function

$$I_C = \Psi_C \Psi_C^* = \widehat{\Psi}_B \left(\widehat{\Psi}_B \right)^*, \quad (3.9)$$

where $*$ notes the complex conjugate of a function. In order to retrieve the phase of the original wavefront a Fourier transform is applied to I_C , resulting in Eq. (3.10),

the terms of which are graphically represented in Fig. 3.3.

$$\begin{aligned} \widehat{I}_C &= \widehat{\Psi}_B \widehat{\Psi}_B^* = \Psi_B(k) \otimes \Psi_B(-k) \\ &= \left(\Psi_{B,0}(k) + \frac{1}{\pi} \widehat{\Psi}_{LP} \left(k - \frac{\lambda f}{T} \right) \right) \otimes \left(\Psi_{B,0}(-k) + \frac{1}{\pi} \widehat{\Psi}_{LP} \left(-k - \frac{\lambda f}{T} \right) \right) \end{aligned} \quad (3.10)$$

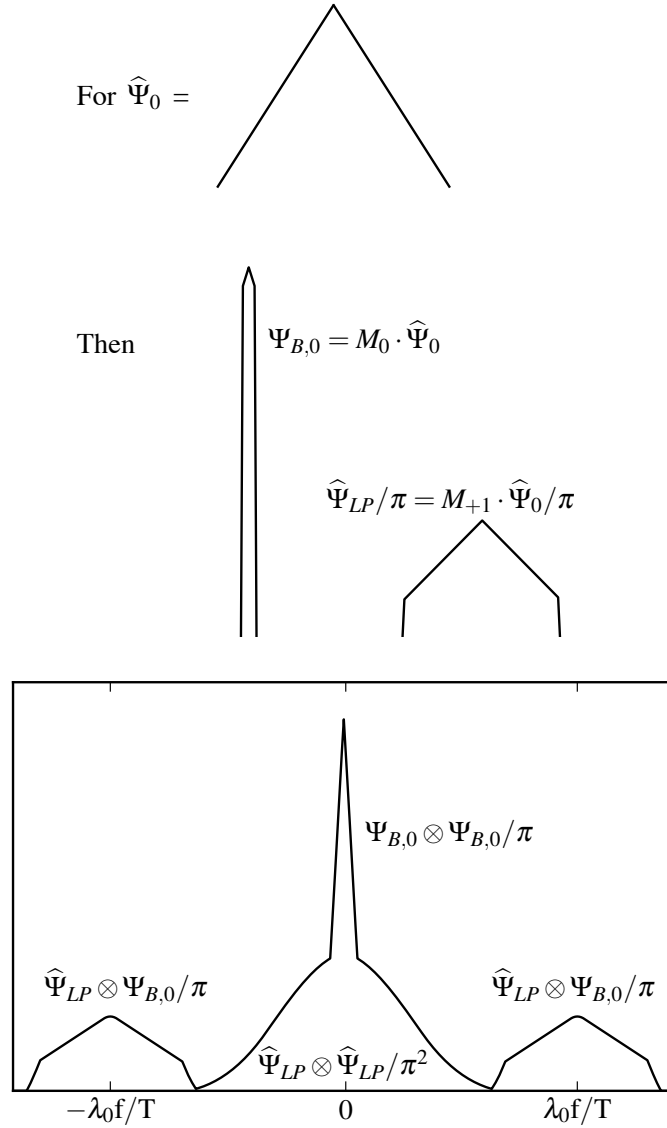


Figure 3.3: Graphical representation of the Fourier transform of the intensity pattern I_C on the detector plane (*bottom*), using an electric field Ψ_0 for which the Fourier transform $\widehat{\Psi}_0$ has an arbitrary triangular shape (*top*). Filtering $\widehat{\Psi}_0$ by apertures M_0 and M_{+1} produces $\Psi_{B,0}$ and Ψ_{LP} respectively (*middle*). The width of filter M_{+1} was chosen to be as wide as possible, without incurring in aliasing between the central lobe $\Psi_{B,0} \otimes \Psi_{B,0}$ and the side lobes. This figure shows the absolute value of all represented terms.

In the figure, the lateral sidebands contain the frequency modulated spectrum of the original phase, convolved with the point-diffraction source $\Psi_{B,0}$. As a consequence, the smaller the point-diffraction source, the higher the fidelity between the sideband and the original wavefront. It is also important to notice that, since Ψ_{LP} has a finite bandwidth, the bandwidth of the central signal $\Psi_{LP} \otimes \Psi_{LP}$ is twice as large as the those of the lateral ones. Therefore, in order to avoid aliasing

$$\frac{3}{2}D_{B,+1} < \frac{\lambda_0 f}{T}, \quad (3.11)$$

since the bandwidth of Ψ_{LP} is set by the aperture M_{+1} of diameter $D_{B,+1}$ in the focal plane filter mask. Note that this means the highest spatial frequency that can be sampled in the X direction is defined by the period of the grating T as

$$k_{max} < 1/3T. \quad (3.12)$$

As a comparison, the period of the grating sets the maximum spatial frequency in the X direction in a similar way the size of subapertures do in a SH. For simplicity in this case M_{+1} has been set to be a square so that k_{max} equals the maximum frequency in the Y direction ζ_{max} . But in principle, for an arbitrary size M_{+1} in the ζ direction, ζ_{max} is independent of the grating's size and only depends on the size μ_p of the pixels sampling I_C . In this case the relationship would be

$$\zeta_{max} < 1/2\mu_p. \quad (3.13)$$

After calculating the Fourier transform of I_C , the next steps to demodulate the wavefront are to filter out the undesirable terms multiplying by M_{+1} , scaling by the amplitude of the point-diffraction source $\Psi_{B,0}$, shifting by $\lambda_0 f/T$ and inverting the Fourier transform. Applying these operations gives

$$\begin{aligned} & \mathcal{F}^{-1} \left[\frac{\pi}{b} \left(\widehat{I}_C M_{+1} \right) \otimes \delta \left(k + \frac{\lambda_0 f}{T} \right) \right] \\ &= \mathcal{F}^{-1} \left[\frac{\pi}{b} \left(\Psi_{B,0} \otimes \frac{1}{\pi} \widehat{\Psi}_{LP} \left(k - \frac{\lambda_0 f}{T} \right) \right) \otimes \delta \left(k + \frac{\lambda_0 f}{T} \right) \right] \\ &= \mathcal{F}^{-1} \left[\frac{\pi}{b} \Psi_{B,0} \otimes \frac{1}{\pi} \widehat{\Psi}_{LP}(k) \right] \simeq \mathcal{F}^{-1} \left[\widehat{\Psi}_{LP}(k) \right] = \widetilde{\Psi}_0, \end{aligned} \quad (3.14)$$

where $\tilde{\Psi}_0$ is the estimate of Ψ_0 . Since the electric field Ψ_0 was defined in Eq. (3.1) as a function of amplitude and phase, it is in principle possible to retrieve both from the estimate $\tilde{\Psi}_0$. In this case, the estimated phase is not a function of the amplitude as it is in some PDIs such as Zernike sensors. Because of this, a non-homogeneous illumination of the pupil does not lead to errors in the estimation of the phase. This algorithm of phase retrieval that consists of using the Fourier transform of the interferogram to retrieve the electric-field was introduced by Takeda et al. (1982). The focus of this thesis will be the estimation of the wavefront, rather than the estimation of the amplitude.

3.2 Simulated model & phase demodulation

In order to test the validity of the analytical model just presented, an end-to-end simulation of the m-PDI is built which does not rely on any of the assumptions made to build the analytical model. So far, the assumption the monochromatic model is built on is that the pinhole is an infinitesimally small point source. As will be seen later, the polychromatic model is built on the further assumption that the SR is constant at different wavelengths within a bandwidth $\Delta\lambda$. Instead, as this section will show, the simulation samples the pinhole using a sufficiently large number of pixels and, for polychromatic simulations, the spectrum is quantised into several monochromatic segments.

A flow diagram of the monochromatic optical simulation of the m-PDI is shown in Figure 3.4. The simulation starts at the entrance pupil by converting an input wavefront ϕ_z in meters to a phase ϕ_r in radians as $\phi_r = (2\pi/\lambda)\phi_z$, and then using Eq. 3.1 to compute the electric field Ψ_0 . This electric field is then multiplied by a binary pupil mask, where a 1 and a 0 represent a transparent and an opaque region respectively, in order to obtain Ψ_A . In the binary pupil mask, which includes both the pupil and the grating, N_P is the number of pixels sampling a line-pair and N_A is the number of pixels across the diameter D_A of the pupil. Both these

numbers remain the same at the back end of the simulation, the detector plane at the exit pupil. N_T , equal to the total number of array elements after zero-padding, is important to determine the scale of an element at the focal plane. Note that the specific values N_P , N_A and N_T take in the figure were chosen for display purposes and are not necessarily representative of the values used in actual simulations.

The electric field Ψ_A is propagated to the focal plane by computing its Discrete Fourier Transform (DFT). The resulting scale of an array element in meters at the focal plane is $\delta_k = f\lambda/N_T\delta_x$, where δ_x is the scale of an element at the pupil plane, also in meters. Since $\delta_x = D_A/N_A$, the scale of an element at the focal plane becomes

$$\delta_k = \frac{\lambda f}{D_A} \frac{N_A}{N_T}. \quad (3.15)$$

The number of elements sampling the pinhole is $N_{B,0} = D_{B,0}/\delta_k$. For a diameter $D_{B,0} = \lambda/D_A$, then $N_{B,0} = N_T/N_A$. The ratio N_T/N_A , which will be referred to as the padding factor, determines the number of elements sampling a λ/D_A length on the focal plane.

For the simulation not to rely on the assumption that the pinhole can be approximated by a delta function, the padding factor needs to be sufficiently large that the error of the simulation is small, but not so large that the simulation becomes too slow, as the total number of array elements increases to the square of the factor. In order to determine what constitutes a sufficiently large padding factor, every time a test is performed, the factor is increased until the results of the simulation converge. In concrete terms, for all the simulations presented in this Chapter it was found that for a padding factor equal to 5, arbitrarily increasing the factor beyond this point produces a maximum error below 1% in the results. This value was found to yield a good compromise between simulation accuracy and usability for the purposes of testing the assumptions of our models against simulated results.

Going back to Figure 3.4, the electric field at the focal plane is multiplied by a binary focal plane filter mask before being propagated to the exit pupil plane by




Figure 3.4: Flow diagram of the monochromatic optical simulation of the m-PDI. This diagram is simplified and does not show all minor operations. The input of the simulation is a wavefront (panel *a*), which in this case follows a Kolmogorov law. The wavefront is turned into an electric field before being multiplied by a pupil plane mask (*b*), which includes the pupil aperture and the Ronchi ruling. The result is then zero-padded in order to control the scale of the sampling at the focal plane. Panel *c* shows the masked angle of the padded electric field that will be propagated to the focal plane. In this case $N_T/N_A = 5$. In all panels black pixels have a value of zero. The padded electric field is propagated to the focal plane (*d*) with a DFT. The electric field is then multiplied by a binary focal plane filter mask (*e*), before being propagated to the exit pupil by another DFT and then multiplied by its complex conjugate to obtain the intensity pattern on the detector plane (*f*).

a second DFT and then divided by N_T^2 for conservation of energy, obtaining Ψ_C . Once at the exit pupil plane, the intensity pattern I_C on the detector is computed using Equation 3.9. At this point, Poisson noise can be added to simulate photon-shot and read-out noise (Janesick, 2001).

The next step after producing the exit pupil intensity pattern is to demodulate it, hence obtaining an estimation $\tilde{\Psi}_0$ of the incoming electric field. The demodulation algorithm is independent of the physical simulation, meaning that with the right parameters it could work on both simulated and non simulated images. Nonetheless the demodulation process is included in this section as it is a necessary component of a full simulation.

Figure 3.5 presents a flow diagram of the demodulation process. The first step is to perform a DFT on the intensity pattern I_C . In the Fourier plane, each sideband contains the result of the convolution $\Psi_{LP} \otimes \Psi_{B,0}$. In order to demodulate Ψ_{LP} , a square of side $2\lambda f/3T$, or $2N_T/3N_P$ in terms of array elements, is cropped off and then an inverse DFT is performed on it. Since the cropped sideband has fewer pixels than the original image, the resulting estimation $\tilde{\Psi}_0$ also does. This does not mean that the cropping sacrifices spatial resolution or that $\tilde{\Psi}_0$ covers a smaller surface than the intensity pattern, but rather than the spatial scale of array elements in $\tilde{\Psi}_0$ is different and reflects the spatial resolution of the instrument, which in turn depends on the number of line-pairs across the pupil. Finally, the wavefront is derived from computing the phase or angle of the complex electric field $\tilde{\Psi}_0$.

Once both the physical simulation and the demodulation algorithm have been produced, they can be tested against known wavefronts. Figure 3.6 shows the output of the physical simulation and of the demodulation algorithm to a few known inputs. In all cases, it is possible to see that the system successfully retrieves all four input wavefronts with a RMS error between 5% and 7%, confirming its validity.

Finally, the model can also simulate polychromatic light of bandwidth $\Delta\lambda$ and centered around a wavelength λ_0 . In order to do so, the physical simulation discre-




Figure 3.5: Flow diagram of the demodulation process, showing actual results from the optical simulation and the demodulation. After a wavefront (a) is fed into the m-PDI (or in this case a simulation thereof) and the corresponding intensity pattern is obtained on the focal plane (b), the demodulation starts by computing the DFT of this pattern. On the Fourier plane (c), a sideband is cropped, as shown by an orange square and an inverse DFT is used to retrieve an estimation of the electric field at the entrance pupil. Panel d shows the phase of the estimated electric field. Although the resulting phase estimation has less spatial resolution than the input wavefront, in this figure it has been resampled in the Fourier plane to have the same number of pixels N_A as that input.

tises a continuum spectrum into several individual wavelengths. Each individual wavelength λ is used to produce an intensity pattern on the detector plane. The input phase $\phi(\lambda)$ is scaled as $\phi(\lambda) = (\lambda_0/\lambda)\phi(\lambda_0)$, where $\phi(\lambda_0)$ is the phase for central wavelength λ_0 , so that the wavefront is the same for all wavelengths. The final intensity pattern is the sum of all the individual intensity patterns, weighted by the light's spectrum. All polychromatic simulations in this chapter will assume a flat spectrum, meaning all wavelengths have the same weight.

The step $\delta\lambda$ between wavelengths is determined using the Rayleigh criterion at




Figure 3.6: Demodulation outputs to three different known phase aberrations: (a) defocus, (b) astigmatism and (c) coma. Simulation performed with 30 line-pairs across the pupil, a padding factor of 5, and a pinhole of diameter $D_{B,0} = \lambda/D_A$. The RMS error between the input and the demodulated aberrations is between 5% and 7% for all three aberrations.

the focal plane, which states that the minimum angular separation the system can resolve in the diffraction limit is

$$\Delta\theta = 1.22 \frac{\lambda}{D_A}. \quad (3.16)$$

Because of the chromatic dispersion of the grating, the angle covered by a step $\delta\lambda$ is

$$\Delta\theta = \frac{\delta\lambda}{T}. \quad (3.17)$$

Equalising both equations and solving for $\delta\lambda$ yields

$$\delta\lambda(\lambda) = 1.22T \frac{\lambda}{D_A}. \quad (3.18)$$

As this number is not constant and is a function of λ , the iterative equation $\lambda_{i+1} = \lambda_i + \delta\lambda_i(\lambda_i)$ is used to determine all the wavelengths sampling the bandwidth $\Delta\lambda$. The last detail to consider is that for all wavelengths to be multiplied by the same focal plane filter mask, they all must have the same spatial scale $\delta_k(\lambda)$. This is controlled by changing the amount of zero-padding for every discrete instance of λ so that

$$N_T(\lambda) = \frac{\lambda}{\lambda_0} N_T(\lambda_0), \quad (3.19)$$

where $N_T(\lambda_0)$ is the value of N_T for λ_0 .

3.3 Noise propagation

Introducing the simplification in Eq. (3.7) into Eq. (3.10), the intensity in the exit pupil becomes

$$I_C = b^2 + \frac{1}{\pi^2} P^2 + \frac{2b}{\pi} P \cos\left(\frac{2\pi}{T} x - \phi_{LP}\right), \quad (3.20)$$

where ϕ_{LP} is ϕ low-pass filtered by M_{+1} . On one hand this simplified expression will prove useful to study the effects of Strehl and chromaticity in Section 3.4. On the other hand, it is the starting point for the analytical derivation of noise propagation. The details of the derivation can be found in Appendix A. It is emphasized that, as the main purpose of the m-PDI is to act as a WFS, only the effects of noise on the retrieved phase are studied.

If the illumination of the pupil is considered to be homogeneous so that $P(x, y) = P_0$, where P_0 is a constant value, and the intensity at the entrance pupil is $I_0 = P_0^2$, the error contribution of the read-out noise and the photon-shot noise is

$$\sigma_R^2 + \sigma_P^2 = \frac{2\pi^2}{9N_P^2 b^2 I_0} \delta S_R^2 + \frac{2\pi^2}{9N_P^2 b^2 I_0} \left(b^2 + \frac{1}{\pi^2} I_0\right), \quad (3.21)$$

where N_P is the number of pixels sampling a line-pair and δS_R^2 is the power of the read-out noise.

For the case of maximum fringe visibility, i.e. when $b = P_0/\pi$, and for a small line-pair sampling with $N_P = 4$, the noise is

$$\sigma_R^2 + \sigma_P^2 \simeq \frac{1.4}{I_0^2} \delta S_R^2 + \frac{0.3}{I_0}. \quad (3.22)$$

Figure 3.7 shows a comparison between simulation results and the model described by Eq. (3.21). The simulation is done with a pinhole diameter $D_{B,0} = \lambda_0/D_A$, which translates into $b \simeq 0.22$. By using two different spatial frequencies, simulation results confirm noise propagation is flat across the frequency spectrum and show good agreement with the theory until the unraveling limit. The simplest form of wavefront reconstruction for this sensor consists in reconstructing the electric field and then retrieving the phase, as stated in Section 3.1.2. Because the phase produced by noise can not be unraveled, the noise reaches a limit. After this limit reconstructed signals are meaningless.

For design purposes, it is interesting to study the dependency of read-out and photon-shot noise to the reference beam's amplitude b . Figure 3.8 shows the propagation of noise as a function of the reference beam's normalized amplitude b/P_0 , for a given sensor flux and as described by Eq. (3.21). Since the noise decreases as b/P_0 increases, there is an incentive to maximize the intensity of the reference beam by having the largest possible pinhole. But as it has already been stated in this section, there is a pinhole size which maximizes visibility. As it will later be shown in Section 3.5 a larger pinhole leads to a reduction in dynamic range and in accuracy for low-order aberrations.

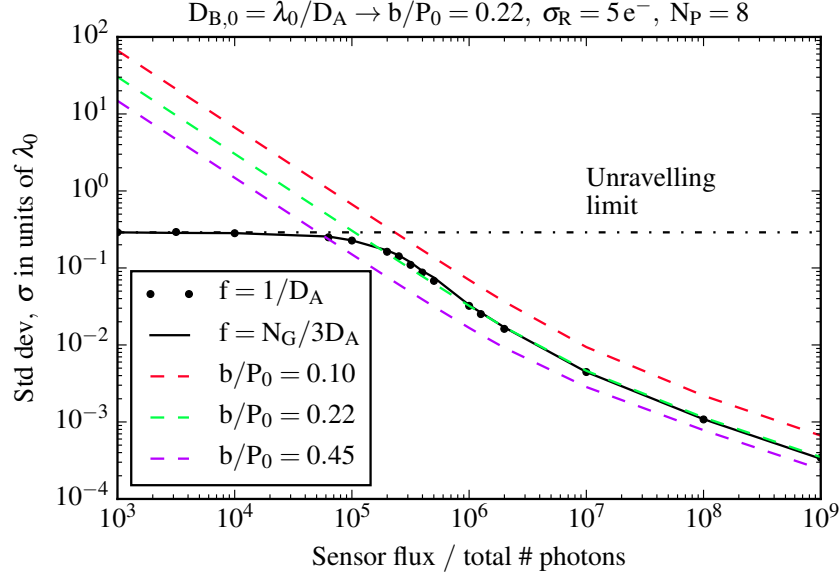


Figure 3.7: Propagation of read-out and photon-shot noise as a function of sensor flux available before the entrance pupil. Simulated data is presented for two sinusoidal wavefronts fed into the WFS, with frequencies $f = 1/D_A$ and $f = N_G/3D_A$ (*black dotted solid*). Simulations are compared to estimations using Eq. (3.21) for 3 values of b normalized by P_0 . The unravelling limit (*black dash-dotted*) is a limit on the standard deviation of noise, as its contribution to phase can not be unravelled beyond $\pm 0.5\lambda_0$. Consequently, as the noise increases the resulting phase contribution converges to a uniform distribution with a standard deviation of $1/\sqrt{12} \simeq 0.29$ in units of λ_0 .

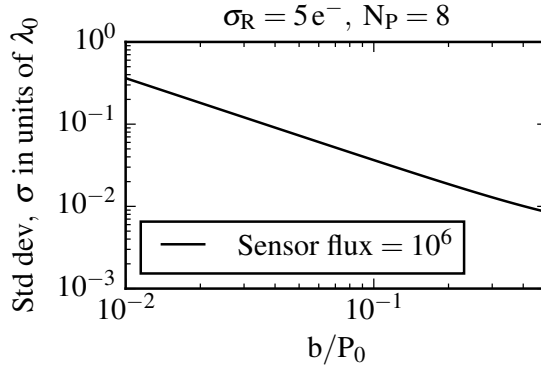


Figure 3.8: Propagation of read-out and photon-shot noise as a function of normalized reference beam amplitude, for a given sensor flux.

3.4 Visibility loss from chromatic bandwidth and a decrease in Strehl

A simple way of including the effects of Strehl into the analytical model is to consider $b = b_0\sqrt{S}$ where b_0 is the highest possible value of b produced with a flat non aberrated wavefront and S is Strehl.

For a rectangular chromatic bandwidth $\Delta\lambda$ around a central wavelength λ_0 and $P = P_0$, the intensity in the exit pupil is (Bharmal, 2005)

$$I_C(x, y, \Delta\lambda, \lambda_0) = \frac{1}{\Delta\lambda} \int_{\lambda_0 - \Delta\lambda/2}^{\lambda_0 + \Delta\lambda/2} I_C(x, y, \lambda) d\lambda, \quad (3.23)$$

where I_C is as described in Eq. 3.20, which in turn relies on the assumption stated in Eq. 3.7 that the pinhole can be approximated to a Dirac delta function. In order to develop this integral, it was also assumed that for a sufficiently small bandwidth $\Delta\lambda$ 1) $S(\lambda) = S(\lambda_0)$, i.e. the SR is constant across the bandwidth, 2) the power spectrum is flat and 3) $\Delta\lambda^2/4 \ll \lambda_0^2$. With these assumptions we have

$$I_C(x, y, \Delta\lambda, \lambda_0) \simeq S b_0^2 + \frac{P_0^2}{\pi^2} + \frac{2\sqrt{S}b_0P_0}{\pi} \text{sinc}\left(\frac{\Delta\lambda\phi_{0,LP}}{2\lambda_0}\right) \cos\left(\frac{2\pi}{T}x - \phi_{0,LP}\right), \quad (3.24)$$

where $\text{sinc}(x) = \sin(x)/x$ and $\phi_{0,LP}$ is the low-pass filtered phase for λ_0 . The first thing to notice about the equation is the modulation of the fringes only depends on the phase of the central wavelength. The effects of the chromatic bandwidth are all wrapped inside the sinc function. Combined, the effects of the Strehl and of the chromatic bandwidth are given by

$$V(x, y) = \sqrt{S} \text{sinc}\left(\frac{\Delta\lambda\phi_{0,LP}(x, y)}{2\lambda_0}\right), \quad (3.25)$$

where $0 < V < 1$ will be considered a proxy of visibility.

It is important to note that V depends both in general on S and locally on $\phi_{0,LP}(x, y)$. In other words, on the one hand a decrease in Strehl produces a

general visibility loss across the pupil, and on the other aberrations also introduce local visibility losses proportional to the bandwidth. But local effects are so small they can be neglected. For example, for $\Delta\lambda = 0.2$ and $\phi_{0,LP} = 0.25$ both in units of λ_0 , then $\text{sinc}(\Delta\lambda\phi_{0,LP}(x,y)/2\lambda_0) = 0.996$. In contrast, for a sinusoidal aberration with an amplitude equal to the previous value, then the general term is $\sqrt{S} = 0.11$, which produces a strong loss of visibility.

To test Eq. (3.25), a sinusoidal aberration perpendicular to the fringes is put across the input pupil as shown in Fig. 3.9. This input allows the measurement of the




Figure 3.9: (a) Input wavefront at entrance pupil A. The sinusoidal input is perpendicular to the grating's line-pairs and to the resulting fringes on the exit pupil. (b) Intensity at exit pupil C

visibility V for known values of $\phi_{0,LP}$. Figure 3.10 shows a comparison between the theoretical visibility loss in the monochromatic case against three polychromatic simulated cases. As expected from the model, the effect of the chromatic bandwidth is negligible when contained under 50%. The biggest mismatch between the simulation and the analytical model is close to 15% and takes place in the low Strehl regime around $S \simeq 0.16$, as shown in Fig. 3.10. This disagreement lies outside of most of the scientific cases for this WFS.

3.5 Accuracy and dynamic range

The design of the m-PDI involves several parameters of which are here considered the number of line-pairs or grooves in the pupil N_G , the number of pixels




Figure 3.10: Loss of visibility as a function of wavefront phase RMS in units of wavelength, for a sinusoidal aberration, perpendicular to the grating’s line-pairs. Solid lines represent simulated polychromatic cases and the dashed line represents the theoretical monochromatic prediction. Visibility is measured in the place of maximum aberration $\phi_{0,LP}$. The vertical line marks the largest disagreement between the theoretical model and the simulation.

sampling the interference fringes N_P , the chromatic bandwidth $\Delta\lambda$ in units of the central wavelength λ_0 and the focal plane pinhole’s diameter $D_{B,0}$. The values these parameters take will determine the accuracy, linearity and dynamic range of the resulting WFS. Since quasi-static aberrations originate in optical misalignments and surface polishing errors (Dohlen et al., 2011; Hugot et al., 2012), they have both low and high order components. Figure 3.11 shows how the response of the WFS concept to low order aberrations changes when these parameters take on different values. Results for different central wavelengths are not presented since all plots remain the same as long as the axes are in units of λ_0 .

In the initial configuration, as well as in most configurations, our sensor remains linear for twice the range of the Zernike sensor (N’Diaye et al., 2013). The figure shows the linearity and the dynamic range are very insensitive to changes in the number of pixels per interference fringe N_P and to the chromatic bandwidth. As a consequence the m-PDI has a large chromatic bandwidth which compensates for the loss of light on the grating and on the spatial filter. The sensor’s throughput,

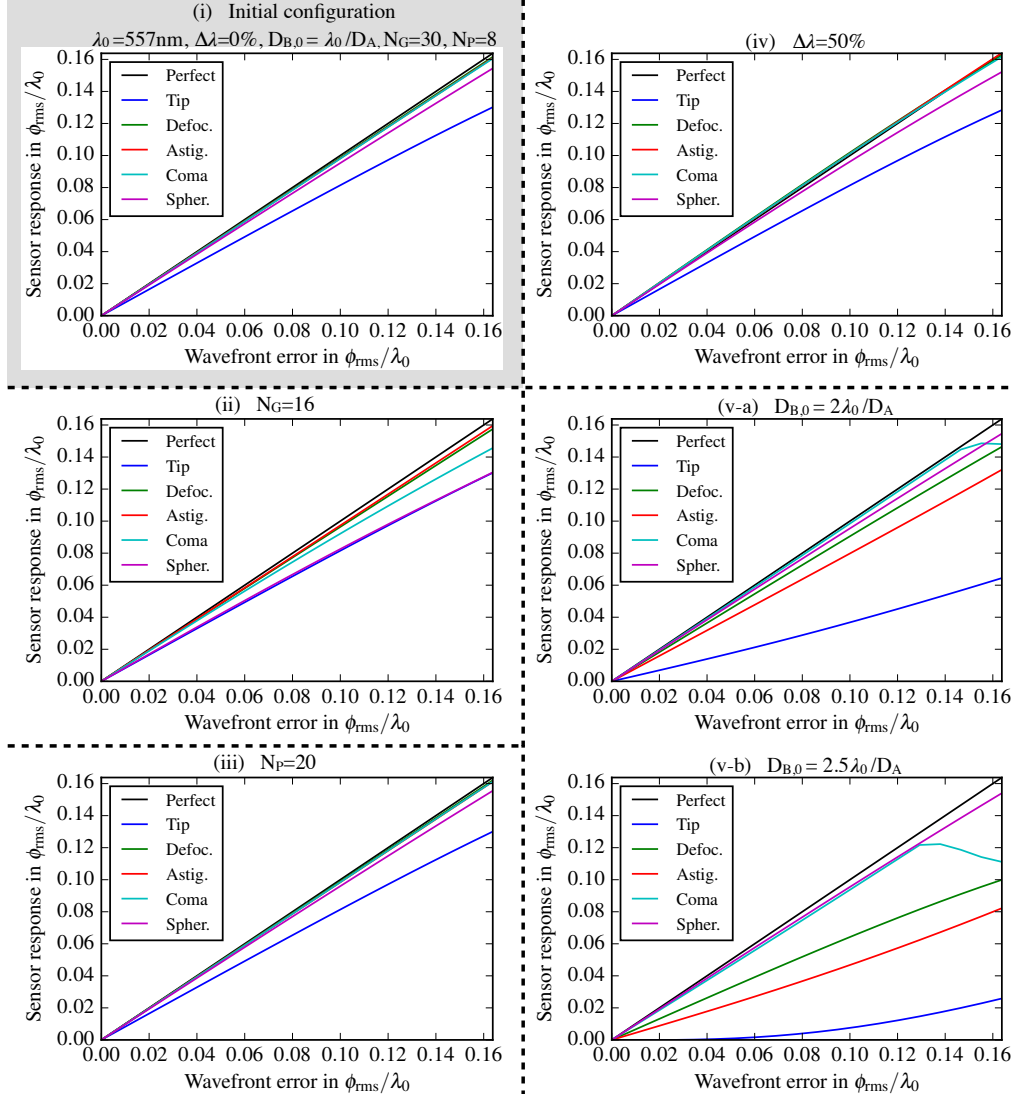


Figure 3.11: Response of the m-PDI to wavefront errors for different low-order aberrations and in different configurations. Each configuration is produced by changing one parameter from (i) the initial configuration. The 4 parameters that are changed are: (ii) the number of line-pairs in the pupil N_G , (iii) the number of pixels sampling an interference fringe N_P , (iv) the chromatic bandwidth $\Delta\lambda$ in units of λ_0 , and (v-a) and (v-b) the focal plane pinhole's diameter $D_{B,0}$.

given perfect optics, is given by

$$\eta = \frac{b^2}{P_0^2} + \frac{1}{\pi^2}, \quad (3.26)$$

which for $D_{B,0} = \lambda/D_A$ gives $\eta = 0.15$. This means that with a bandwidth of $\Delta\lambda = 50\%$, the sensor takes in as much light as another sensor with a throughput of 1 and a bandwidth $\Delta\lambda = 7.5\%$

In panel (ii), a reduction to the number of line-pairs across the pupil N_G produce a small reduction of 16% in the response to spherical aberrations. Since the maximum spatial resolution is proportional to N_G and that both N_G and N_P can easily be changed without much effect on the linearity and dynamic range, the number of pixels can be kept low, all while having a high resolution.

Finally, the size of the pinhole is the variable with the greatest effect on the accuracy. Indeed, in Fig. 3.11 panel (v-a) and for $D_{B,0} = 2\lambda_0/D_A$ the response to tip (or tilt) is reduced by 50% as most of the PSF's core is well inside the pinhole. This is balanced by an increase in the amplitude of the reference beam b and a subsequent decrease in photon-shot and read-out noise as shown in Fig.3.8. Doubling the size of the pinhole increases b by a factor of 1.9 up to 0.41, which in turn decreases noise by 42%. This configuration could be desirable in cases where the WFS does not have to measure tip and tilt but the greatest possible value for b is required to minimize photon-shot and read-out noise. As the pinhole's diameter becomes greater than $2.5\lambda_0/D_A$, one by one and almost from the lowest to the highest order, the responses to different aberrations start to decrease. This can be seen in Fig. 3.11 panel (v-b). Here $b = 0.42$, leading to a reduction of 43% on the noise compared with the initial configuration. It is also possible to notice there is a sharp decline in accuracy to coma, abruptly restricting the dynamic range for this mode. The origin of this phenomenon as well as its observability in the fringe pattern is to be the subject of further study. Larger pinholes are not desirable as all the accuracy is lost with no noise reduction. For increasingly higher pinhole diameters, b converges to 0.45, which at most produces a decrease of 46% with respect to the initial configuration.

The effects of spatial filtering on the WFS's accuracy to different spatial modes are better explored in Fig. 3.12. The figure shows the transfer function relative to spatial frequency for different pinhole sizes and different pre-existent tilts. Tilts can be caused by pointing errors or other sources of misalignments. The transfer function spans all the spatial frequencies that can be measured with $N_G = 30$, as

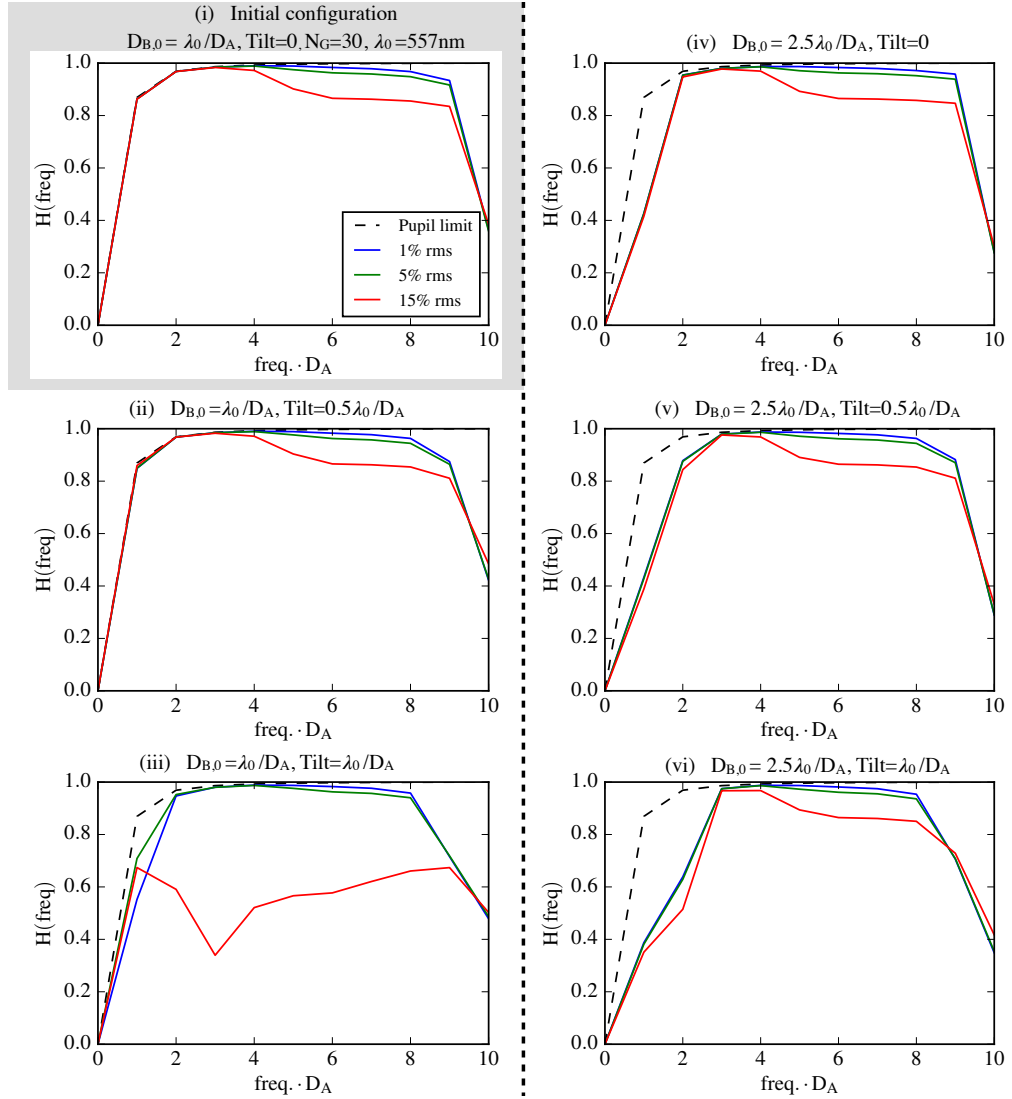


Figure 3.12: Transfer function with respect to spatial frequency for different aberration amplitude and in different configurations. Columns have different pinhole sizes and rows have different preexisting tilts.

described by Eq. (3.12). As can be observed in all the panels, the transfer functions suffers from a sharp decline as it approaches the maximum frequency k_{max} . It is also important to notice that the maximum sensitivity to low order aberrations is given by the shape of the pupil. In this case, a circular pupil has a poor transfer function for frequencies with a spatial scale in the vicinity of its diameter. This is referred to in the figure as the pupil limit and is represented by a dashed line.

The first thing to notice is the difference between the left panels, which all have

a small pinhole size $D_{B,0} = \lambda_0/D_A$, and the right panels, with a larger pinhole size $D_{B,0} = 2.5\lambda_0/D_A$. With a bigger pinhole size the WFS loses sensitivity to small aberrations. This happens because all aberrations that are not filtered by the pinhole are present in the test beam and therefore become invisible to the sensor.

The disadvantage of a small pinhole though are aberrations can more easily drain it of the light necessary to produce the reference beam. Figure 3.12 shows, from top to bottom, how the transfer function changes with an increasingly larger preexisting tilt. At first, the effects go unnoticed, but in panel (iii) there is clear loss of dynamic range. This doesn't happen in the panel to the right, panel (vi), where the same tilt is applied but the pinhole is larger. The advantage of the larger pinhole is to keep more light into the reference beam despite aberrations.

Finally, the transfer function acts as a good indicator of the shape of the SNR as a function of spatial frequency, since noise is flat across the spectrum. The sensitivity of the m-PDI in terms of SNR is better discussed in the following section.

3.6 Application to a real system

Here, the application of the sensor to a real system is considered. The sensor will get a 5% share of light from a beam-splitter, so other operations can be performed in parallel (like science observations). It is worth noticing the sensor could also be fed the entirety of an unused chromatic band by a dichroic. This alternative is not presented here. The system will be similar to other XAO systems that have been developed (Macintosh et al., 2008; Sauvage et al., 2010), but with a lower read-out noise to consider newer detector technology. These systems are optimized to work around $\lambda = 1.6 \mu m$. In the case of GPI, for its first light the system was able to close the AO loop for stars $I < 8 mag$, but was expected to operate down to $I \simeq 10 mag$ under better seeing conditions (Macintosh et al., 2014). Table 3.1 shows the parameters of the simulation.

Table 3.1: Parameters used for simulating the exposure time.

Parameters	Values
Central wavelength λ_0	$1.625 \mu\text{m}$
Bandwidth $\Delta\lambda$	20%, 50%
Zero mag. flux density	1080 Jy
Apparent star magnitude	10
Telescope diameter	8 m
Telescope transmission T_{tel}	40%
Beamsplitter transmission T_{BS}	5%
Line-pairs across pupil N_G	60, 185
Pixels per line-pair N_P	4
Read-out noise	$1 e^-$

Figure 3.13 shows the sensor's sensitivity to the order and amplitude of aberrations, represented as the exposure time needed to achieve $SNR = 1$. The simulation is

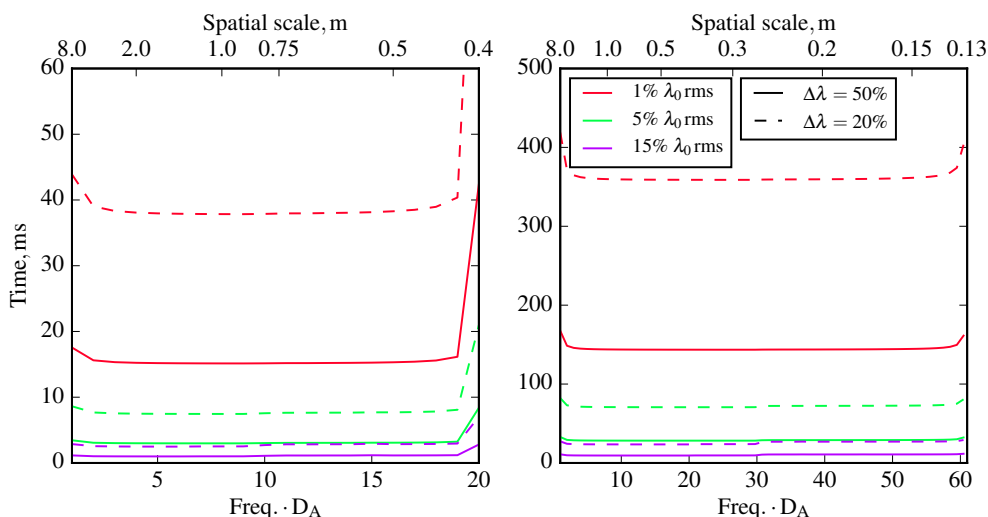


Figure 3.13: Exposure time required to achieve $SNR = 1$ as a function of aberration spatial scale. (left) $N_G = 60$, (right) $N_G = 185$. The exposure time is presented for different amplitudes of aberration and for different chromatic bandwidths. The parameters of the simulation can be found in Table 3.1.

performed for two different bandwidths, $\Delta\lambda = 20\%$ and $\Delta\lambda = 50\%$. In the results, the exposure time is inversely proportional to the chromatic bandwidth. This expands upon previous results showing the sensor's achromaticity for all orders.

As mentioned, the transfer function shapes the curves of exposure time. The exposure time is also inversely proportional to the transfer function. This means

higher amplitude aberrations will present curves that are less flat around the center than lower amplitude ones.

The simulation was also performed for a low-order case with $N_G = 60$, and a high-order case, with $N_G = 185$. In the low-order case, the exposure times required are in the order of milliseconds and tens of milliseconds, meaning the sensor can operate at frequencies of up to hundreds of Hz. This is not fast enough to act as the main sensor of an XAO loop operating at frequencies that exceed 1 kHz, but would be enough for real-time quasi-static aberrations measurements. In the high-order case the sensor operate at frequencies of a few Hz. This is still enough to measure quasi-static aberrations down to the deformable mirror's scale, with actuators with a 13 cm separation.

3.7 Conclusions

The principle for a m-PDI has been presented. Later an analytical model was developed. The model allows for electric field reconstruction, i.e. to retrieve both the phase and the amplitude of an incoming electric field. The propagation of photon-shot noise and detector read-out were derived. The model was also extended to quantify chromatic effects and the influence of Strehl on the visibility of fringes. Both sets of predictions, on the propagation of noise and on the loss of visibility were tested against simulations built on first principles. The analytical models show good agreement with the simulations, showing the WFS is well understood.

An important result shows this interferometer has a wide chromatic bandwidth. Not only is the visibility of fringes unaffected by a wider bandwidth, but neither is the accuracy of the WFS to low nor high order aberrations. When $\Delta\lambda = 50\%$ the high accuracy of the monochromatic case is conserved. A wide chromatic bandwidth allows one to compensate the loss of light at the grating and at the spatial filter. The trade-off of a wider chromatic bandwidth is a reduced range of measurable aberrations.

The WFS model is also shown to be sensitive to the size of the pinhole. While doubling the pinhole's diameter from λ_0/D_A to $2\lambda_0/D_A$ reduces the sensitivity to tip and tilt by 50%, it also decreases photon-shot and read-out noise by 42%. Increasing the size of the pinhole beyond this point is ill-advised as the noise can only decrease a further 4% while losing all the sensitivity to low-order aberrations. Finally, the exposure times required are in the order of milliseconds in a low-order case, meaning in this system the sensor can operate at frequencies of up to hundreds of Hz, with the beamsplitter and star magnitudes that were tested. This is not fast enough to act as the main sensor of an XAO loop operating at frequencies that exceed 1 kHz, but would be enough for real-time quasi-static aberrations measurements.

The exposure time required to achieve a given SNR was shown to be mostly flat across the spatial spectrum. This time is inversely proportional to the chromatic bandwidth and the aberration's amplitude, and is in the order of milliseconds for low-order aberrations, and tenths of a second for high-order aberrations. Such timescales are reasonable for the real-time measurement of quasi-static aberrations.

Design & implementation

After developing the analytical models for the m-PDI and studying some of its properties in simulation, the next step in the concept validation path is to build a physical implementation in the lab and test it experimentally. The first implementation of this concept, dubbed the Calibration and Alignment Wavefront Sensor (CAWS), was presented and explored by Bharmal et al. (2012). This chapter presents a new iteration of this instrument, with a design that integrates the lessons from its first instance and that is geared towards being tested on a high-order AO testbed, Canary Hosted-Upgrade for High-Order Adaptive Optics (CHOUGH) (Bharmal et al., 2014, 2018). Because of this, the instrumental requirements and the restrictions imposed by CHOUGH need to be translated into design choices. Once the optical design has been determined and verified by using ray tracing software, the mechanical design can be produced. This design must comply with stringent tolerances, the most demanding of them being the precise and stable alignment of the focal plane filter mask and of its pinhole. Consequently, the mechanical layout is closely linked to the alignment procedure, which has to be planned for simultaneously. Finally, once the CAWS has been designed and built, an initial round of testing is conducted to confirm the instrument's correct functioning and to characterise it for future experiments.

This chapter presents the two main stages required to go from the initial concept to

a fully working instrument in the lab. Section 4.1 goes through the process of taking system restrictions and requirements and translating them into an optical and a mechanical design. Then, Section 4.2 presents the experimental results for the initial round of testing and characterisation. This characterisation will then allow us to move into more sophisticated experiments, that further integrate the CAWS into the AO system and examine their interaction with each other.

4.1 From requirements to optical design

The optical design of the CAWS is straightforward. At its core, the instrument is an optical relay which can be made of just two lenses. At the entrance pupil, a small angle beamsplitter, in this case a diffraction grating, produces the beams that will become the reference and test beams respectively. These beams get filtered by a focal plane mask at the relay's focal plane. Finally, the beams get recombined in an exit conjugate pupil, where the detector samples the resulting interferogram. With these elements set, the bulk of the optical design effort consists in tuning several degrees of freedom, such as the magnification of the relay, the period of the grating and dimensions of the apertures on the focal plane filter mask. As this is the first time this is attempted, this section focuses on establishing a systematic approach to sequentially translate restrictions and requirements into parametric design constraints. Once the parameters surrounding any given component, such as a lens, have been sufficiently constrained, then that component can either be selected from a commercially available stock or be manufactured according to a custom design. Either way, the selection of a component crystallises the values for the parameters involved, helping to tighten the constraints around other parameters and leading to the determination of other components.

Given that this implementation of the CAWS is primarily oriented towards lab testing, the majority of the constraints are imposed by CHOUGH, the AO testbed that will host it. It is worth dedicating a short introduction to this bench, as it

will be central to the experiments presented in this chapter and in the following one. Figure 4.1 presents CHOUGH’s layout. Despite it being originally designed

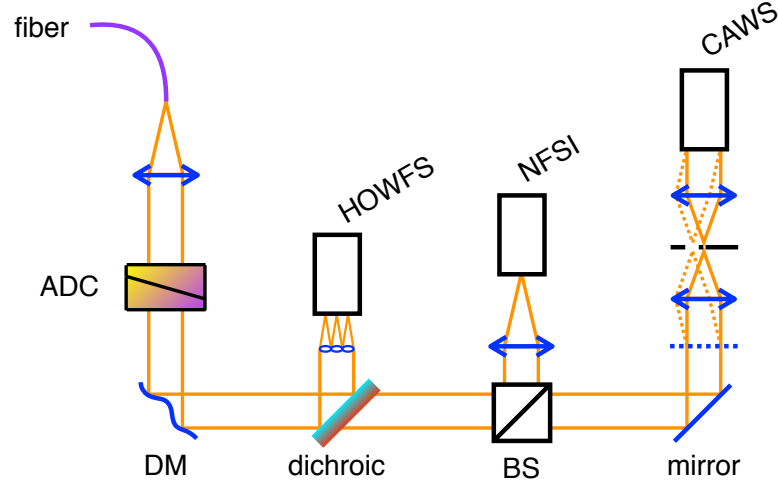


Figure 4.1: Layout of the high-order AO testbed CHOUGH. The HOWFS is a 31-by-31 SH, the NFSI is a narrow FOV imaging camera, the DM is a 32-by-32 Boston Micromachines Kilo DM, the dichroic is a 647 nm longpass one, and light is provided by a fibre-fed illumination interface. For all tests presented in this thesis the Atmospheric Dispersion Corrector (ADC) has been removed.

as an add-on to enhance other lower order AO systems, it has all the necessary components to act as a fully independent AO bench, provided an adequate illumination interface such as IRIS is used. The system’s 31-by-31 SH, the HOWFS, the narrow FOV imaging camera, the NFSI, and the host port are all in closed-loop configuration with the DM. In addition, a dichroic splits the light’s spectrum, sending all of the shorter wavelengths to the SH, while leaving longer wavelengths to pass through to the NFSI for higher SRs. This not only enables closed-loop capabilities, but it also allows one to perform a detailed diagnosis on any visiting instrument located on the host port. Because of this, the CAWS will be designed to be compatible with this port, which will limit the instrument’s size and give us the starting point for the design process.

4.1.1 Optical relay

In order to begin, it is important to identify the most relevant restrictions on the system. In this case, the CAWS can only go on CHOUGH's host port, at the beginning of which is a conjugated pupil plane. This port has a small footprint, of only 380 mm long, which is not long enough to build a separate relay, before the CAWS, to resize the pupil being delivered to it. Consequently, the CAWS needs to be designed to use an entrance pupil of diameter $D_A = 10.54$ mm, conjugated to a 32-by-32 actuators Boston Kilo DM. The spatial resolution of the CAWS, which is determined by the number of diffraction grating line-pairs inside the pupil, needs to be sufficiently high that it can properly sample the DM. On the one hand, the maximum possible spatial frequency that the DM can produce is

$$k_{control} = \frac{N_{act}}{2D_A} \text{ m}^{-1}, \quad (4.1)$$

where N_{act} is the number of actuators across the pupil. On the other hand, the maximum spatial frequency that can be observed by the WFS is

$$k_{obs.} = \frac{1}{3T_G} \text{ m}^{-1}, \quad (4.2)$$

where T_G is the period of the diffraction grating. For the CAWS to sample all of the DM's possible spatial frequencies, the condition

$$k_{obs.} > k_{control} \quad (4.3)$$

needs to be met. By developing, this becomes

$$T_G < \frac{2D_A}{3N_{act}} = 0.22 \text{ mm}. \quad (4.4)$$

This in turn results in $N_G > 48$ line-pairs of the diffraction grating across the pupil.

There are several benefits to meeting this constraint, maximising T_G (or minimising N_G). The first is that fewer line-pairs in the pupil will translate into a smaller requirement on the number of detector pixels. In turn, having fewer pixels increases the flux per pixel and potentially reduces the read-out time and the computational

load for phase demodulation. It is important to remember that the resulting WFS would have enough spatial resolution to sample most high-order DMs and could therefore be transported and tested on most low-order AO systems. The last advantage is that, since fewer pixels are needed, then the footprint of the exit pupil can be smaller, which reduces the focal length of the collimating lens leading to such a pupil. This is important to meet the tight space constraints imposed by CHOUGH. Based on this, a Ronchi ruling with a period $T_G = 0.2$ mm is selected, which delivers $N_G = 52.7$ line-pairs. In comparison to other dispersive elements such as *échelles*, Ronchi rulings have the advantage of being cheap, compact and easy to replace. Their disadvantage is that they are not very efficient, having a throughput of only 50%. More expensive and better performing elements can be considered in future upgrades, when the instrument is better understood and is ready to be taken to more advanced stages of development. Regarding the period of the selected grating, it is the longest we were able to find in commercially available stocks. The next option in decreasing period was half as short, resulting in many more line-pairs than are needed.

The next restrictions to take into account are the sampling of interference fringes on the detector and the availability of cameras. By knowing the pixel sizes of the cameras and the required number of pixels per line-pair of fringes it is possible to determine the magnification of the optical relay. Regarding the sampling, Figure 4.2 illustrates the relationship between the total number of pixels N_T and the number of grating line-pairs N_G , both across the pupil. From the figure it can be concluded that for the sideband M_{+1} to be completely within the range of detectable spatial frequencies, the condition

$$\frac{N_T}{2} > \frac{4}{3}N_G \quad (4.5)$$

must be met. By rearranging this, the number of pixels per line-pair is

$$N_P = \frac{N_T}{N_G} > \frac{8}{3} \simeq 2.66 \text{ pixels/lp.} \quad (4.6)$$

Since the detector is sampling the pupil at the exit plane C, the number of pixels

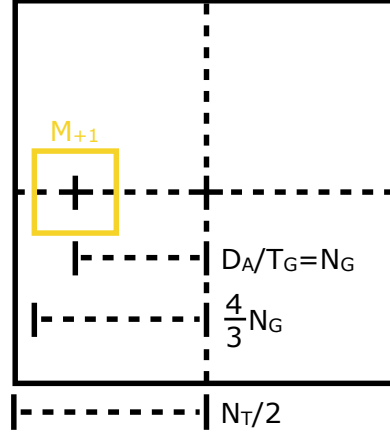


Figure 4.2: M_{+1} sideband region on DFT of the interferogram. N_T is the total number of pixels across the pupil.

can be calculated as $N_T = D_C/\mu_P$, where D_C is the diameter of the exit pupil and μ_P is the pixel size. In turn, $D_C = D_A/M$, where M is the demagnification of the optical relay and $D_A = T_G \cdot N_G$. Combining all of the above into Eq. 4.6 gives

$$\frac{T_G}{M \cdot \mu_P} > \frac{8}{3} \quad (4.7)$$

$$\Rightarrow M < \frac{3T_G}{8\mu_P}. \quad (4.8)$$

This upper limit on the magnification depends upon the pixel sizes. The two available cameras are an Imperx Bobcat B0620M and an AVT Manta G-145B NIR. Both camera have relatively similar pixel sizes, with $7.4 \mu\text{m}$ and $6.45 \mu\text{m}$ respectively. The resulting constraints on the magnification are $M < 10.14$ and $M < 11.63$. The lower limit on the magnification is determined by the cameras' number of pixels. While the Bobcat has a 640×480 detector, the Manta has a large chip with 1388×1038 . For N_T to be smaller than the shorter side of each camera, the lower limits on the magnification are $M > 2.97$ and $M > 1.57$. Within these ranges lower values of M are preferred in order to avoid any potentially unknown effect that might emerge from sampling the fringes with too few pixels.

Another restriction on the relay is that it should be based on a 4f telecentric design. Although one of the advantages of this design is that it remains compact by using only two lenses, its main purpose is that the chief rays of the test and the reference

beams are parallel to each other. As shown in Figure 4.3, if the rays were not parallel to each other, then any defocus of the focal plane filter mask would lead to an asymmetric filtering of the test beam through aperture M_{+1} and could also potentially register as a tip in phase. This happens because the converging test




Figure 4.3: Diagrams of the first half of the CAWS going from the diffraction grating to the first lens' focal plane, with a defocused focal plane filter mask, with a (a) a 4f telecentric and (b) a non-telecentric optical relay design.

beam no longer travels well centred through the aperture. In this sense, a telecentric design is preferred to avoid this and other unforeseen effects, making for a more robust implementation.

The last restriction before selecting a pair of lenses is the length of the available space. For two lenses ℓ_1 and ℓ_2 of focal lengths f_1 and f_2 respectively, the distance from the entrance pupil to the exit pupil in a 4f system is $2f_1 + 2f_2$. This length must be shorter than the available space $L_{avail.} \simeq 300$ mm, which is the difference between the length of the space envelope provided by CHOUGH and the length of the longest camera (to be conservative). Given these limitations and the commercially available lenses, two achromatic doublets with focal lengths $f_1 = 100$ mm and

$f_2 = 25$ mm were selected. With these lenses, the relay measures 250 mm and the magnification is $M = f_1/f_2 = 4$, which is near the lower limit but still compliant with the restrictions set by both cameras.

Indeed, this design can host either of the two cameras, in terms of available space and sampling. Furthermore, both cameras are comparable in other respects such as read-out noise, pixel depth, data transmission rate (GigE interfaces), having a C-mount and features such as binning and windowing or Region of Interest. This means that immediately selecting a camera is not a crucial decision leading to an irreversible situation. For now, the Bobcat will be used due mainly to two practical reasons. The first one is that the research group has a great deal of experience integrating this camera to the Durham Adaptive Optics Real-time Controller (DARC) (Basden and Myers, 2012). This real-time controller can provide many services of a standard AO system, such as processing SH data and controlling DMs. Furthermore, configuration files already exist to run this camera with DARC, and this real-time controller is already being used in CHOUGH. The second reason is that with a smaller frame, higher frame rates are achievable without having to bin or window pixels. With the Bobcat's pixel size the sampling of fringes is $N_P = 6.8$, which is on the lower end of the allowed range, as was intended.

The next aspect to be considered into the optical design is the chromatic range of the instrument. This will affect the central wavelength λ_0 , the spectral bandwidth $\Delta\lambda$ and in consequence the coating of the lenses and the chromatic aberrations. The biggest restriction in this regard is the FF596-Di01 Semrock (Semrock, 2020) longpass dichroic sitting at the beginning of the port. This dichroic, with a cutoff wavelength of $\lambda_c = 647$ nm at an incidence angle of 13 degrees, sends the bluer end of the spectrum to CHOUGH's SH and leaves the redder part for the host port assigned to the CAWS. Since, as can be seen on Figure 4.4, the quantum efficiency (QE) of the camera is almost negligible for wavelengths beyond 900 nm, this could be set to be the CAWS' maximum wavelength λ_{max} . In a system using a dispersive element as its small-angle beamsplitter, such as a Ronchi ruling, the test beam has




Figure 4.4: Quantum Efficiency of the Imperx Bobcat B0620M, based on data found in the camera’s manual.

chromatic dispersion. As a result, the maximum and minimum wavelengths are set by the outer and inner edges of aperture M_{+1} respectively, and can be computed as $\lambda_{max} = \lambda_0 4/3$ and $\lambda_{min} = \lambda_0 2/3$. Having set the maximum wavelength, the central and minimum wavelengths become $\lambda_0 = 675$ nm and $\lambda_{min} = 450$ nm. This range exceeds the limits imposed by the dichroic and allows in commonly used alignment sources such as Helium-Neon (HeNe) and red diode lasers, respectively at 633 nm and 635 nm. On that note, a 2 mW HeNe laser is powerful enough that a substantial amount of light still makes it through the dichroic and into the host port.

The spectral limits of the instrument can then be used to select the lenses’ coating. Given that the system is designed to work centered around the red part of the spectrum, the most suitable coating is Edmund Optics’ VIS-NIR anti-reflection coating, which reflects less than 1% of light in the 400-1000 nm range. With this, the resulting optical design produced with Zemax’s OpticStudio design suite is as shown in Figure 4.5.

4.1.2 Bias

Two main characteristics of this design must be examined in order to check its performance: the chromatic aberrations around the focal plane filter mask introduced




Figure 4.5: Optical model of the CAWS using OpticStudio. D_A and D_C are the diameters of the entrance and exit pupils respectively.

by the lens ℓ_1 and the static aberrations. Starting with the chromatic aberrations, they are presented in Figure 4.6. Since this design was focused around its central




Figure 4.6: Estimation of the differential chromatic defocus term in the CAWS' focal plane, between wavelengths spanning its chromatic range and the central wavelength $\lambda_0 = 675$ nm. Curve was produced with OpticStudio

wavelength $\lambda_0 = 675$ nm, the figure shows the differential defocus term on the focal plane of ℓ_1 , where the filter mask is. In the case where the instrument was used with broadband light, this chromatic defocus would introduce a bias in the demodulated wavefront. The first thing to notice is that the defocus term between the central wavelength and the wavelength of an alignment source, such as a HeNe laser, is only 20.2 nm. This aberration is small enough that the focal plane filter mask could be focused using such a source without introducing significant defocus bias. What is more problematic is the aberrations at the edge of the CAWS' chromatic

bandwidth. At the extremes, the defocus error goes from 48.6 nm at $\lambda = 528$ nm to 130.7 nm at $\lambda = 900$ nm. At the edge of the pupil, if broadband light was used and both these wavelengths were present, their difference in radians would be 0.82π . This being less than π means that the peaks of the interference fringes produced by one wavelength, would not cross over the valleys of the fringes produced by the other. Because of this, fringes would still be visible, albeit with a reduced visibility, and a phase measurement could be derived from them. This effect is partially mitigated by the low QE of the camera at longer wavelengths and could be further mitigated by using a bandpass filter, restricting the chromatic bandwidth. With respect to the added bias, a pair of complementary and independent methods are presented in Chapter 6 that allow one to calibrate for this chromatic bias.

The second characteristic that must be examined in order to check the instrument's performance is the bias added by the first and second lenses, ℓ_1 and ℓ_2 respectively, and which will register as a static aberration. Figure 4.7 shows these aberrations, as estimated from the optical model. The aberrations introduced by ℓ_1 are filtered out of the reference beam by the pinhole, but are still left on the test beam. Since what appears as a phase measurement on the WFS is the phase difference between the reference and the test beams, essentially all the aberrations added by ℓ_1 onto the test beam will become a bias. This is what is shown on the top left of the figure. Considering the effects of ℓ_2 , if the instrument was completely common path, then the footprints and angles of incidence of both beams would be the same, resulting in them also having the same aberrations. As a result, these aberrations would become invisible, therefore not adding any bias. But because one beam is slightly offset and tilted with respect to the other, it gets slightly different aberrations. This difference is plotted on the top right of the figure. Finally, the sum of both differential aberrations is added to produce the total narrowband bias, shown at the bottom of the figure. The total narrowband bias is just 4.8 nm rms, with a PtV of $\sim \lambda_0/38$. This is small enough that if an otherwise perfect AO loop was closed, the maximum achievable SR would be 99.8%.




Figure 4.7: Biases introduced by (*top left*) ℓ_1 , (*top right*) ℓ_2 and (*bottom*) $\ell_1 + \ell_2$, for $\lambda_0 = 675$ nm.

4.1.3 Focal plane filter masks

Many characteristics of the focal plane filter masks, such as the diameter of M_0 , the geometry of M_{+1} or the distance between both apertures, depend on the rest of the design. Just these few variables make it so that masks with the desired characteristics do not exist *off the shelf*, requiring them to be manufactured with great precision and accuracy. For example, a pinhole of diameter $D_{B,0} = 2\lambda/D$ would be $14\ \mu\text{m}$ across in the current design. A good option is to manufacture the masks using laser machining, which allows one to produce apertures of arbitrary geometry and with a precision of about $\pm 2\ \mu\text{m}$ (14%). This means their design can be left last, without influencing back on the rest of the optical design.

Coming back to a critical aspect of this mask, the pinhole diameter $D_{B,0}$ affects




Figure 4.8: b_0 as a function of pinhole diameter $D_{B,0}$, for a circular pupil.

both the visibility of fringes and the accuracy of the instrument. With respect to the visibility of fringes, it will depend on the ratio between the amplitudes of the test beam and the reference beam. In a low-aberration regime, the visibility depends upon the variable b , which is the amplitude of the reference beam. In turn, this variable can be described as $b = P_0 b_0$, where P_0 is the average amplitude at the entrance pupil and b_0 is an adimensional number between 0 and 1. From Eq. 3.20 in the previous chapter, it is possible to deduce that the visibility in a low-aberration regime is

$$V = \frac{2b_0}{\pi b_0^2 + 1/\pi}. \quad (4.9)$$

Calculating b_0 can be done by knowing the geometry of the pupil. Because of this, it can be difficult to obtain an analytical result in most cases. Instead, it is easier to compute this value numerically. Figure 4.8 shows the relationship between b_0 and the pinhole diameter $D_{B,0}$, in units of λ_0/D_A . Two special cases are highlighted in the figure. The first one, in red, is the case of maximum visibility $V = 1$, which occurs when $b_0 = 1/\pi$ and requires a relatively small pinhole, with $D_{B,0} = 1.33\lambda_0/D_A$. This pinhole size is near $D_{B,0} = \lambda_0/D_A$, which was studied in the previous chapter, and it is expected to yield similarly high accuracies. For this reason, pinholes with diameters in this vicinity are preferable, to get both high accuracy and high visibility. But due to practical reasons related to the difficulties of aligning such small pinholes, it was decided to go for larger diameters.

In the second case, shown in dashed blue, the pinhole has a diameter of $D_{B,0} = 2.5\lambda_0/D_A$. As shown previously in Chapter 3, with this diameter only about 50% of accuracy is lost for astigmatism and defocus, while being left intact for all other higher aberration modes. Furthermore, the visibility also remains high with $V = 0.96$. In this case, the actual diameter of the pinhole on the mask would be $16 \pm 2 \mu\text{m}$, with a $\pm 13\%$ error. From a previously built prototype, it was discovered that aligning a pinhole of a comparable size and getting interference fringes was possible in a similar setup and with the tools available in the lab. For example, some of these tools include translation stages that can move $0.7 \mu\text{m}/\text{degree}$ of adjuster knob rotation. Under the pessimistic assumption that the knob can be manually set with a precision of $\pm 2^\circ$, then the pinhole can be aligned with a precision of $\pm 1.4 \mu\text{m}$ or $\pm 8.8\%$. Smaller pinholes, such as with $D_{B,0} = \lambda_0/D_A$, could prove to be more difficult to align as the precision becomes $\pm 22\%$, but provide better accuracy when aberrations are small. In addition, a smaller pinhole would drift more relative to its size due to thermal expansions and other yet uncharacterised effects, making the final alignment less stable.

In order to test different designs, several masks were initially commissioned to the company Laser Micromachining Ltd. (Laser Micromachining Ltd., 2017) This came with the added bonus that the elevated fixed costs of making a single mask are offset by making many. The standard mask layout is presented in Figure 4.9 In this




Figure 4.9: Layout of focal plane filter masks.

layout, the dimensions B and a are fixed in order to prevent aliasing, whereas H and d_0 are independent. The advantage of this is that by having a larger H , higher

frequencies could be sampled in the direction parallel to the fringes, or Y direction, which is something that could be explored in the future. The dimensions of two of the manufactured masks are shown in Table 4.1. Due to their larger pinholes,

Table 4.1: List of focal plane filter masks.

	Mask 1	Mask 2
$D_{B,0}$	$2.5\lambda/D_A \rightarrow 16 \mu\text{m}$	$2.5\lambda/D_A \rightarrow 16 \mu\text{m}$
H	225 μm	1350 μm
B	225 μm	225 μm
a	338 μm	338 μm

these masks will be the first to be tested as they are the easiest to align, and are therefore the most likely to produce stable interference fringes. If these experiments are successful, then other masks with smaller pinholes can be tested as well.

4.1.4 Mechanical implementation

Having determined the optical design of the instrument, the next step is to establish the mechanical support for these optics. The optical design produced in the previous section presents several implementation challenges. First, from the preliminary work done aligning an early prototype, it was discovered that finding the pinhole during the alignment process is very hard, and that once found, its position relative to ℓ_1 should be kept as fixed as possible. This requires these elements to be mechanically coupled. Second, later movements of ℓ_2 and of the camera tend to misalign the pinhole if in turn it happens to be mechanically coupled to them. It is then preferable to reduce the coupling between the pinhole and ℓ_2 or the camera. Third, the layout is very compact, leaving very little room between certain elements such as those just mentioned. This makes decoupling them for stability more difficult. And lastly, it would be an advantage if the whole WFS was easy to align and to reconfigure, as this design is oriented towards lab testing.

The last feature is the easiest to address and will provide the basic structure for the rest of the build. As shown in Figure 4.10, the entire WFS is mounted on

a rail that adds strong constraints in the X and Y axes (transverse and height respectively) and allows easy movement along the Z (longitudinal) axis. The same rail can then help fix the position of elements in the Z direction once in place, but it does not stop the optics from rotation around their vertical axis. This can be fixed by adding a cage, which grants the additional advantage of imposing the same height on all the optical elements. Unfortunately, a single overarching cage would mechanically couple all optical elements.




Figure 4.10: Mechanical assembly of the CAWS made with Inventor.

As seen in the figure, this is solved by using two coaxial cages. The inner cage, which holds ℓ_1 and the pinhole, is only coupled to the second cage, which holds ℓ_2 and the camera, through the rail. This way, pushing and pulling the camera along the Z axis does not affect the position of the pinhole. As will be discussed in the following section, this will be often the case when turning the WFS into its spectrograph mode and back.

Lastly, there is the challenge of mounting the ℓ_2 lens to just 25 mm from both the camera's detector and the focal plane filter mask on either side, without their mounts having a conflict. For this, a 17.8 mm in diameter tube mount is used to hold ℓ_2 through the 48 mm in diameter clearance left by the 60 mm cage plate attaching the camera to the outer cage. But since this tube is ultimately also

attached to the outer cage, the cage plate holding it also leaves a clearance, this time of 26.3 mm in diameter, for another 17.8 mm tube mount to hold the focal plane filter mask inside that clearance and at the right distance from ℓ_2 . This design has the added benefit of leaving very little room for light to sneak into the camera and onto the detector. The instrument, as built on the lab, is shown in Figure 4.11.




Figure 4.11: Image of the CAWS.

4.2 Testing & characterisation

With the CAWS built and placed on the host port, CHOUGH's capabilities are used to test and characterise the instrument. The purpose of these initial tests is to corroborate that the instrument's physical behaviour and the demodulation algorithms are well understood and are working properly. For this, the tests will be divided into two sets. The first set uses monochromatic light and the DM to measure the CAWS' spatial transfer function. This transfer function describes the WFS's gain against the whole range of spatial frequencies in its range, and can be used for calibration purposes. The second set uses polychromatic broadband light to determine the chromatic spectrum as seen by the CAWS's detector. This

measurement is a prerequisite to performing polychromatic wavefront demodulations, as it enables the estimation of an equivalent central wavelength, which is a necessary value to convert the retrieved phases into wavefronts.

4.2.1 Control software & data pipeline

The software to run all experiments in this and the next chapter is divided into two main layers. These software layers, as well as the physical layer which includes all the hardware, are presented in Figure 4.12. The first software layer is the interface layer containing the real-time controller (RTC) which, as mentioned earlier, is implemented using DARC. This particular RTC can be configured to acquire data from multiple cameras, control the DM, and provide some basic data processing, such as computing centroids from SH spots.

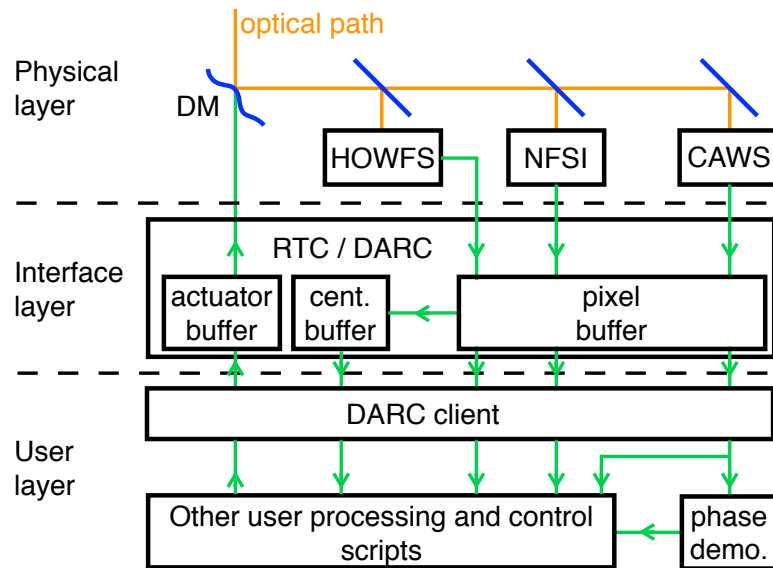


Figure 4.12: Layout of the data acquisition, processing and control software for experiments of CAWS on the CHOUGH high-order AO bench. In the interface layer, the real-time controller DARC interfaces with the DM, the HOWFS, the NFSI and the CAWS. DARC loads pixel data from the cameras of these instruments into its pixel buffer and pushes commands found in its actuators buffer onto the DM. DARC also computes HOWFS centroids and loads them into its centroid buffer. All these buffers are accessible from the user layer by a Python DARC client. This client then delivers this data to a phase demodulation algorithm or to any other processing and control script defined by the user.

The RTC is then controlled from the second software layer, the user layer, by a DARC client written in Python and provided by the standard DARC installation. This client allows retrieving and pushing data from and into DARC's pixel, centroid and actuator buffers. Besides the DARC client and all other standard libraries such as Matplotlib (Hunter, 2007) and NumPy (Oliphant, 2006), all of the software in the user layer is original. This includes the demodulation algorithm presented in Section 3.2, which as mentioned previously works on both simulated and real m-PDI images, as well as any other processing and control scripts defined by the user.

4.2.2 Spatial transfer function

In principle, the measurement of an instrument's spatial transfer function is simple: a pure tone, i.e. a single spatial frequency, of known amplitude is produced, in this case by the DM, and then measured by the instrument. This spatial frequency, which has the shape of a sinusoid across the pupil, is held static for the duration of the detector's exposure. The transfer function for any given frequency is the ratio between the amplitude of the measurement and that of the input generated by the DM.

Before performing this experiment, the first step is to measure the static aberrations of the system, so they can be calibrated out of the data. Figure 4.13 shows monochromatic interference fringes and the corresponding demodulated phase. The data were obtained with the HeNe light source ($\lambda = 633\text{ nm}$) and with a DM sent to its zero position. Note this position is a flat command on the DM at the middle of its range, which does not guarantee a flat shape. On a small side note, all interferometers have at most a $\pm\pi$ phase limit, after which, if the aberrations push the phase beyond π or below $-\pi$, then the phase is wrapped back to $-\pi$ or π respectively. In order to solve this, a fast algorithm developed by Herráez et al. (2002) is used to unwrap phase discontinuities.

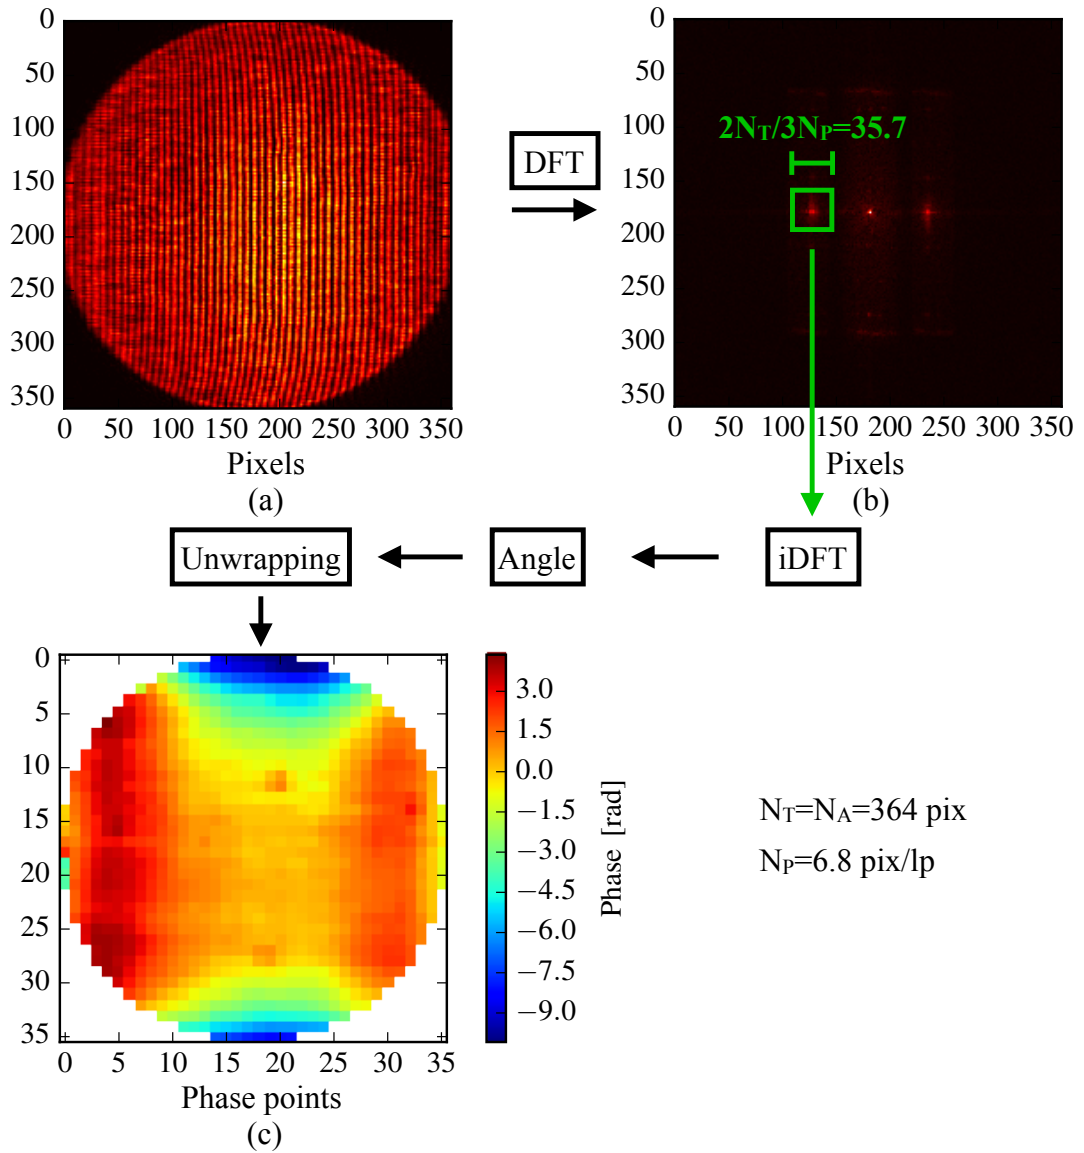


Figure 4.13: (a) Interference fringes produced with a DM in its zero position, (b) the DFT of the fringes and (c) the corresponding demodulated reference phase. Panel b shows the square root of the absolute value of the DFT for display purposes. The light source is a HeNe laser at $\lambda = 633$ nm. The physical scale of the pupils in the fringe image and the demodulated phase is the same as they span the length of the pupil, but the sampling is lower on the demodulated phase. The reason for this is a line-pair is being sampled by $N_P = 6.8$ pixels and the width of 3 line-pairs equals the width of 2 phase data points. The figure also shows the demodulation process, including the cropping of the sideband in the Fourier plane, the inverse DFT, the computing of the angle of the resulting electric field and the unwrapping of the phase. The cropped region in the Fourier plane is rounded down to 35 pixels in order to avoid aliasing.

Now that the static aberrations have been calibrated, it is possible to proceed with the experiment of feeding the CAWS sinusoidal aberrations, which will be produced up to the maximum spatial frequency the DM can produce, equal to 16 cycles/pupil. Figure 4.14 presents the amplitudes measured by the CAWS, after being presented several static sinusoidal wavefronts produced by the DM at different frequencies and uncalibrated amplitudes. The DM commands in volts were calculated using

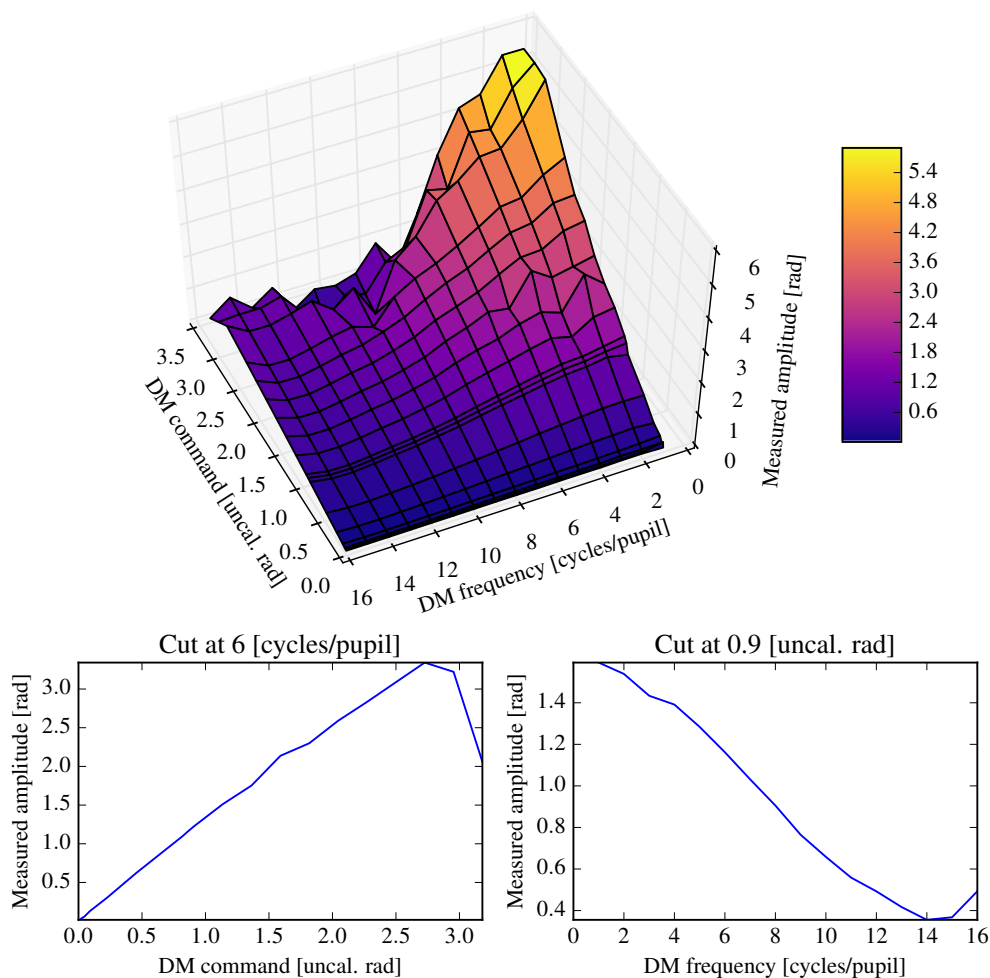


Figure 4.14: (*top*) Amplitude of sinusoidal wavefronts measured by the CAWS. The DM command axis shows the uncalibrated amplitude commanded onto the DM, using a simple volts-to-radians proportionality constant. (*bottom-left*) Cut along the DM command axis. (*bottom-right*) Cut along the DM frequency axis. The light source is a HeNe laser at $\lambda = 633$ nm.

the intended amplitude and frequency of the sinusoidal wavefront, such that

$$u(y) = 2 \frac{\lambda}{2\pi} G_{DM}^{-1} a_{uncal.} \sin(2\pi f y / D), \quad (4.10)$$

where $u(y)$ is the tension of an actuator of coordinate y in the pupil's vertical axis (tests were also carried out in the horizontal direction X with the same results), $a_{uncal.}$ is the intended, uncalibrated amplitude in radians, f is the spatial frequency in cycles/pupil and G_{DM} is a constant value in meters/volts provided by the DM's manual. Note that this G_{DM} is equivalent to the array of the same name presented in Chapter 1, with the difference that by assuming that all actuators are heterogeneous and have no interactions with each other, then it can be substituted by a single number. Several commanded amplitudes were used in case the transfer function changes as a function of them. As shown in the figure's bottom-left panel, the measured amplitudes seem to increase linearly with those of the DM, for any fixed spatial frequency. But from a different point of view as seen in the bottom-right panel, the measured amplitudes decrease as the frequencies increase, for any given commanded amplitude.

This behaviour, which is not predicted by the simulation results of Figure 3.12, seems to be caused by the DM, rather than being a feature of the CAWS. Whereas there is no effect that could explain this on the side of the WFS, a potential explanation on the side of the DM is that at higher frequencies, the inter-actuator coupling (in a continuous DM, contiguous actuators are coupled by the mirror sheet) stops them from reaching high amplitudes. Subsequent investigations, described below, support the conclusion that this effect is due to the DM

A way to resolve this issue is to characterise the DM by reproducing the sinusoidal wavefront experiment, but with another independent WFS, and use these results to calibrate the CAWS' results. The accuracy of this characterisation will, in part, determine the accuracy with which the CAWS' transfer function is ultimately calibrated and estimated. In the case of CHOUGH, the DM's characterisation can be done with the 31-by-31 SH, HOWFS. Figure 4.15 presents the results of this

experiment performed with the HOWFS. Note that the sinusoidal wavefront ex-

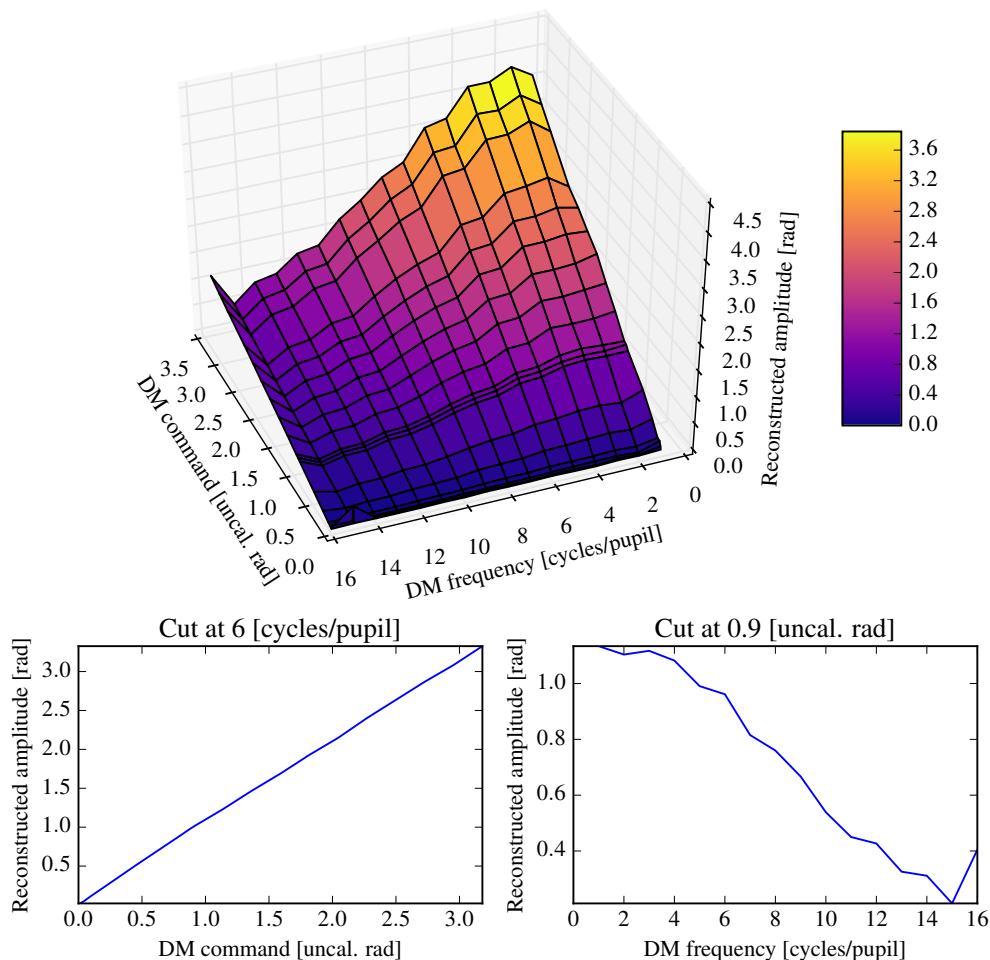


Figure 4.15: (top) Amplitude of reconstructed sinusoidal wavefront estimated by the HOWFS. The DM command axis shows the uncalibrated amplitude commanded onto the DM, using a simple volts-to-radians proportionality constant. (bottom-left) Cut along the DM command axis. (bottom-right) Cut along the DM frequency axis. The light source is a HeNe laser at $\lambda = 633$ nm.

periment was performed simultaneously with both the CAWS and the HOWFS, to make sure that both WFSs were receiving the same signal, in case of DM hysteresis. As shown in the figure, an effect similar to that observed with the CAWS is present here, where the reconstructed amplitude decreases at higher frequencies. This supports the hypothesis that the effect is produced by the DM, rather than by the WFS. Furthermore, since the commanded amplitudes span the entirety of

the DM's command range, at the end of which the maximum allowed voltage is applied to the actuators, then the HOWFS measurements show us the DM's actual dynamic range as a function of frequency. Amplitudes above the DM's dynamic range are inaccessible in this setup and can not be used to test the CAWS' transfer function.

These results can nonetheless be used to calibrate the CAWS' measurements against that of the HOWFS, thereby removing the effects of the DM. The disadvantage of this approach is that the characterisation of the CAWS' transfer function will be relative to the HOWFS rather than to the DM. Consequently, this indicates that it is impossible to characterise any single component in absolute terms, and that instead components can only be characterised in relation to others.

Since there is reason to believe that the behaviour of the DM in open-loop is highly non-linear, the CAWS' measurements need to be calibrated against those of the HOWFS, in order to obtain the CAWS' transfer function. For this, the amplitudes measured with the first are divided by those reconstructed with the second. Figure 4.16 shows a contour plot of the ratio of both sets of data. If

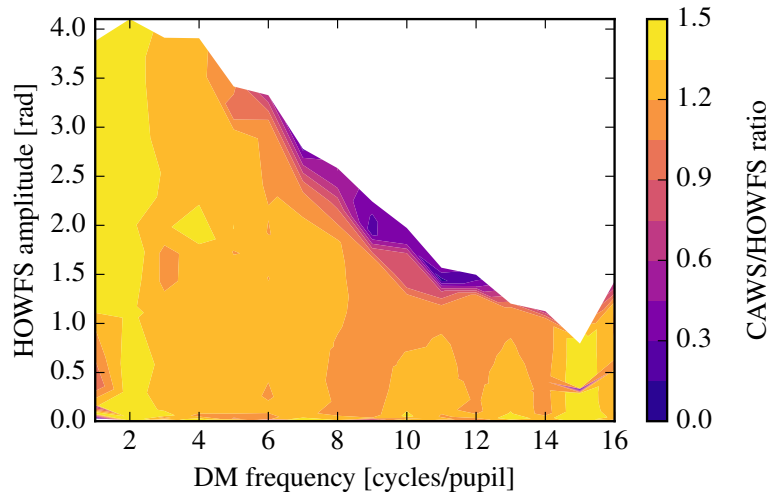


Figure 4.16: CAWS' monochromatic transfer function, calibrated by the HOWFS for $\lambda = 633$ nm.

the amplitudes estimated from the reconstructed HOWFS data are considered to

be the best estimations of the actual amplitudes, then the ratios presented in the figure can be taken as the CAWS' transfer function, calibrated by the HOWFS. For the axes in the figure to be consistent with this approach, the Y axis represents the amplitudes read by the HOWFS rather than the one commanded onto the DM.

The first thing to note on the figure is that there are no data points at high frequencies and high amplitudes, hence, the top-right half of the figure is empty. This happens because, as mentioned for the previous figure, the DM can not access these amplitudes for those frequencies. Therefore, in the contour plot the frontier between the void and non-void part of the figure marks the maximum amplitude the DM can deliver. As explained earlier, the decrease in maximum amplitude at higher frequencies is due to the inter-actuator coupling at high spatial frequencies, which prevents the DM from actually delivering these amplitudes. Also from the figure, above around 3 cycles/pupil, the maximum amplitude produced with the DM decreases proportional to the frequency.

The next thing to note is that, for most of the parameter space, the ratio is above 1. This ratio is especially high below 3 cycles/pupil, reaching values between 1.3 and 1.5. By comparing the bottom-right panels of Figures 4.14 and 4.15, it can be observed that this is due to a turnover or flattening of the HOWFS' reconstructed amplitudes at these frequencies. No intrinsic reason regarding the nature of the WFSs has been found that can explain this turnover or the overall high ratio. An error could be hiding in the physical implementation of the WFS, **CHOUGH**, the phase demodulation algorithm or the wavefront reconstruction. More experiments are still needed to look for unexpected phenomena, such as looking at focal plane images during the production of sinusoidal wavefronts. For the time being it will be considered that the ratio is close enough to one for the calibrated transfer function to provide an initial understanding of the CAWS' workings.

On this note, it is possible to observe a tendency for the ratio to decrease from values above 1.3 at the lower end of frequencies, down to values between 1 and 1.2 at the higher end. The total end-to-end drop of about 15% to 30%, is the

result of both a gradual decrease across the frequency range and of a step at the middle. Both effects are predicted by the simulated transfer functions presented in Figure 3.12, showing a gradual decrease in sensitivity for high frequencies on the one hand, which is exacerbated at higher amplitudes where it reaches about 20%, and the step drop that appears at the middle of the frequency range on the other. This step, half way through the frequency range of the instrument, is explained by the fact that the Fourier transform of an electric field with constant amplitude but sinusoidal wavefront, will present harmonics of the fundamental tone. As the amplitudes increase, the proportion of power in all the harmonics increases relative to the fundamental. These harmonics are visible in the focal plane, where the CAWS filter mask is located. For frequencies above half the sensor's range, the first and strongest harmonic no longer goes through aperture M_{+1} in the mask. As a result, a fraction of the tone's power is filtered out, producing the sensitivity step drop observed.

A small exception to this tendency seems to take place at a frequency of 15 cycles/pupil, where the ratio increases to above 1.3. From the raw data it is concluded that this small exception is due local aliasing effects beginning to appear in both WFSs as they approach their spatial frequency limit. It is important to remember that aliasing can occur at certain phases despite still satisfying Nyquist's frequency criterion. Since the CAWS can sample slightly higher frequencies than the HOWFS, it is less affected by aliasing at this particular frequency.

The last thing to note is that between 6 and 10 cycles/pupil, the ratio between both instruments drops almost to zero at the higher end of the amplitude range. Examining the CAWS' data presented in Figure 4.14, this is explained by a sharp dip in the amplitudes measured by the CAWS at the same frequencies and for DM commanded amplitudes above 3rad. In turn, the dip itself is explained by the CAWS reaching the limit of its dynamic range. This range is not simply expressed by a constant number as is the case for a SH, where the dynamic range of the instrument is given by the local dynamic range of subapertures, which is generally

the same for all of them. Instead, this range is the result of the interaction between several variables and phenomena, some of which have not yet been modelled by analytical expressions. Nonetheless, it is understood that one of the most important factors limiting the dynamic range is the wavefront's local gradient or tilt. Some, but not all, of the evidence for this is present in Figure 4.16, where the dynamic range decreases at increasing frequencies. It is then possible that the dynamic range is inversely proportional to the wavefront's local slope, although this is not well explored in this data set, only reaching the CAWS limit for some frequencies.

It is believed that the way the dynamic range depends on the wavefront's gradient is that the latter reduces the amount of light going through the pinhole, reaching the region of the pupil in question. This then reduces the visibility of the fringes in the area, blending them with the background and decreasing their local SNR. Below a certain SNR threshold, the noise starts determining the phase, which approaches a random uniform distribution between $-\pi$ and π . Furthermore, in low visibility areas, where the amplitude of the electric field is low and can be read as zero due to noise, phase vortices appear, which can be propagated across regions of the pupil by the unwrapping algorithm. Due to the complexity of these phenomena, modelling them and finding strategies to address the problems they ensue is ongoing work at the time of writing and is not presented in this thesis.

Regardless of this, it is still possible to conclude from the figure that the phase demodulation algorithm works, that the CAWS has an approximately flat response across most of its frequency domain, with a slight attenuation on the second half, and that the dynamic range decreases with frequency. Having performed this initial assessment of the CAWS' monochromatic behaviour, the next step is to perform spectral measurements using the instrument's spectrograph mode. These measurements are essential if the WFS is to be used with broadband light in polychromatic operations.

4.2.3 Spectrograph mode

Since one of the aims of the m-PDI concept is to perform wavefront sensing in an astronomical AO system, it must be able to do so using broadband polychromatic light from a natural guide star. The broader the chromatic range of the instrument, the more light it can accept, allowing it to observe fainter stellar targets. For this reason, it is essential that the polychromatic capabilities of the CAWS are confirmed empirically.

Since the demodulated phase is described in units of radians, before a polychromatic wavefront can be reconstructed in metres and used to command a DM, it is important to determine the equivalent central wavelength to perform a radians to meters conversion. The main reason why the demodulated phase in radians can not be used directly to command the DM is that it is unlikely that the control matrix will have been built with an internal light source with the same spectral properties as the stellar target. If the internal light source and the guide star have different equivalent central wavelengths, then either the demodulated wavefront or the control matrix need to be rescaled accordingly. In this sense, the equivalent central wavelength will be, by definition, that wavelength which when used to rescale the wavefront into units of meters, minimises the error between the estimated and the actual wavefront.

In order to measure the spectrum of the light being used and estimate its equivalent central wavelength, the CAWS is turned into a spectrograph by adding an auxiliary lens ℓ_{aux} . This configuration, which had been originally devised to image the focal plane for alignment purposes, can also measure the chromatic dispersion of light produced in aperture M_{+1} by the diffraction grating. Consequently, this configuration, which is presented in Figure 4.17, will be called the spectrograph mode. An advantage of using this configuration over using another spectrograph, is that the spectral measurements are performed using all the same elements as those in the CAWS and in the same positions, plus one lens. Thanks to this,




Figure 4.17: Layout of the CAWS in spectrograph mode. The detector is pushed back into Plane D and an auxiliary lens is placed after ℓ_2 , in order to re-image the focal plane filter mask, located on Plane B.

they incorporate all of the system's chromatic effects, such as lenses' chromatic transmission curves or the detector's QE, to produce a spectrograph from the perspective of the CAWS fringe creation process.

Once the instrument has been set into the spectrograph mode, broadband measurements can be taken and spectral calibrations performed. Broadband measurements will be produced using a Thorlabs OSL2B halogen lamp, coupled to CHOUGH's fibre feed. Regarding the pixel spectral scale calibration (i.e. pixel scale to wavelength), given the short spectral range of the CAWS, going from $\lambda_{min} = 450$ nm to $\lambda_{max} = 900$ nm, it is not necessary to use traditional calibration sources with many emission lines over a broader range, such as Neon or Argon lamps. Instead, a HeNe laser, which only provides one spectral line at 633 nm is enough for the whole of the range. The HeNe laser is actually preferred over calibration lamps given it can easily be coupled to CHOUGH's fibre-fed illuminator without adding any additional optics and provide sufficient light to amply illuminate the system, even after going through CHOUGH and its 647 nm longpass dichroic. Figure 4.18 shows both a broadband measurement acquired using the halogen lamp and a calibration image produced with the HeNe laser. In the figure, both images show diffraction mode 0 going through the pinhole M_0 on the left, and diffraction mode +1 going through




Figure 4.18: Focal plane images of diffraction modes 0 and +1 with (*top*) OSL2B halogen lamp and (*bottom*) HeNe laser as reference for the spectral calibration. Red lines show, from left to right, the positions of $\lambda = 0 \text{ nm}$, $\lambda = 633 \text{ nm}$, and $\lambda = 647 \text{ nm}$, which correspond to the pinhole, the HeNe laser, and the dichroic's cutoff respectively. Original values have been square-rooted to increase the contrast of the images presented in this figure.

aperture M_{+1} on the right. Note that the HeNe laser is visible because it is powerful enough to leak through the dichroic.

In the images, under the assumption that the system has sufficiently small field distortions, the distance x between mode 0 and any wavelength in mode +1 is

$$x(\lambda) = \frac{f_1 f_{aux}}{T f_2} \lambda, \quad (4.11)$$

where f_{aux} is the focal length of ℓ_{aux} and T is the diffraction grating's period. The compound error in those terms produced due to misalignments can be corrected by using the spectral calibration measurement. This way, all wavelengths can be accurately found on the image plane.

By finding the dichroic's cutoff, which is marked in the figure with one of the red lines, it is possible to note that with the halogen lamp there is light at shorter wavelengths than this filter would allow. There are several reasons for this. First, the light directly left of the cutoff leaks through the dichroic's transition zone and then gets spreaded over by the spectrograph's low spectral resolution, which will be calculated further ahead. Second, the small bump at 550 nm is due to a poor rejection of the dichroic at that wavelength.

The last step before producing a spectrogram is to measure the intensity of each wavelength. In diffraction limited conditions, the radius of the PSF grows proportional to the wavelength, reducing the intensity at the centre. In order to compensate for this, the best way to measure the intensity at a given wavelength is to integrate the light vertically, along the images' Y axes. In non-diffraction limited conditions, such as is the case here due to static aberrations in CHOUGH, the PSF does not necessarily grow proportional to wavelength. Note that this experiment is performed in open-loop, by setting the DM into a position that minimises the static aberrations and that was previously found in closed-loop with monochromatic light, as will be shown in the following chapter. Despite this, some static aberrations still remained that prevented us from reaching the diffraction limit. Fortunately, the method of integration in the Y axis is also useful in the presence of residual static aberrations. Figure 4.19 presents the resulting spectrogram. One of the simplest results that can be obtained from the curve is that, with a full-width at half-maximum (FWHM) of 76.5 nm, the light source's spectrum is well within the CAWS' spectral range, spanning only 18% of it. The left and right margins of the figure are set to the CAWS' spectral limits to illustrate this fact.

In order to set a benchmark for these results, the spectrogram produced by the CAWS is compared to that obtained by another independent spectrograph, the Ocean Insight USB2000+, for the same light source and coupled to the same fibre. Whereas the CAWS' curve does not extend far below the dichroic's cutoff wavelength at 647 nm, both measurements match after about 680 nm. This serves as




Figure 4.19: Normalised chromatic spectrum of the OSL2B halogen light source, (*blue*) as seen through CHOUGH and measured by the CAWS on mode +1, and (*black*) as measured by an independent spectrograph. The CoM and FWHM are presented for the spectrum as sensed on mode +1.

an initial validation of the CAWS' spectrograph mode and of its data reduction pipeline.

But as can be observed around 650 nm, the independent spectrogram presents high frequency variations that the CAWS' spectrogram is not being able to match. This is because the CAWS has a lower spectral resolution which averages out small details. By definition, the spectral resolution of a spectrograph is

$$R = \frac{\lambda}{\Delta\lambda}, \quad (4.12)$$

where $\Delta\lambda$ is the minimum wavelength difference that the spectrograph can resolve around the wavelength λ . On one hand, the angular difference between two wavelengths $\lambda + \Delta\lambda/2$ and $\lambda - \Delta\lambda/2$ is

$$\begin{aligned} \Delta\theta &= \frac{\lambda + \Delta\lambda/2}{T} - \frac{\lambda - \Delta\lambda/2}{T} \\ &= \frac{\Delta\lambda}{T}. \end{aligned} \quad (4.13)$$

On the other hand, according the Rayleigh's criterion the minimum angular separation the optical system can resolve in the diffraction limit is

$$\Delta\theta = \frac{1.22}{D}\lambda. \quad (4.14)$$

Equating Eq. (4.13) to Eq. (4.14) and solving for $\lambda/\Delta\lambda$ gives

$$R = \frac{\lambda}{\Delta\lambda} = \frac{D}{1.22T} = 43.2. \quad (4.15)$$

This is the best possible resolution the instrument can achieve in the diffraction limit, when all static aberrations have been removed. In this case, because of the static residuals, the number should be slightly lower. What is important is this order of magnitude is enough to make sure that the light source is within the spectral range of the CAWS and to calculate the equivalent central wavelength.

It is expected that a valid way to estimate this value is to calculate the centre of mass (CoM) of the curve. This is because the contribution of each wavelength to the final interferogram is proportional to that wavelength's intensity and to its coupling to the pinhole. In this scenario, where lab conditions can be controlled and the SR be made relatively high, it will be assumed that all wavelengths are well coupled to the pinhole. Using the CoM method, the equivalent central wavelength is $\lambda_0 = 667$ nm. Confirming this requires using the CAWS as a WFS to measure the same aberrations with both monochromatic and broadband light, and testing that the equivalent central wavelength found by this method is indeed the one which minimises the error between them. This experiment is presented in Chapter 6.

The current chapter presented the process through which the CAWS was designed, built and preliminarily tested. The design of this WFS required systematically evaluating the different restrictions imposed by the hosting AO testbed CHOUGH, such as available space, pupil size, minimum spatial resolution, sampling, and spectral range. After the instrument was assembled and aligned to CHOUGH, an initial round of testing was performed to confirm that the instrument was working as expected. First, a series of static sinusoidal wavefronts were created using CHOUGH's DM, in order to estimate the CAWS' transfer function, which was then calibrated using the HOWFS. Secondly, the CAWS was set into its spectrograph mode to characterise a halogen lamp's spectrum in the range allowed by the dichroic, successfully matching the results obtained with an independent spectro-

graph. Having confirmed that the instrument behaves as expected in open-loop and that the demodulation and spectral estimation algorithms work, the following chapter introduces closed-loop experiments.

Closing the control loop

One of the main purposes of a WFS, especially in the context of AO, is to inform the actions of wavefront correction devices such as DMs. As mentioned in a previous chapter, placing the WFS and the DM in a closed-loop configuration is useful in several scenarios to consecutively minimise residual aberrations, whether they are the product of systematic biases, DM non-linearities or actively changing atmospheric turbulence. It is therefore important to test our novel WFS, the CAWS, in closed-loop with a DM to, first of all, confirm that obtaining a stable loop is indeed possible and, second, that there are not any unforeseen phenomena arising from the interaction between them.

As this is the first time a control loop is closed using the CAWS, this will be done with static aberrations already present in CHOUGH, rather than with atmospheric turbulence or a lab simulation thereof. Demonstrating that the CAWS can be used to deal with this type of aberrations is not only necessary as, like other PDIs, it is in part intended to do this, but it is also simpler as their amplitude is limited, the loop's progress is easy to observe from one iteration to another, and they are already present in the AO system. Hence, this chapter develops some of the necessary tools to close an AO control loop and to assess its performance, and then presents the results of doing so. Section 5.1 introduces a method for estimating the system's SR, which will be used to independently quantify residual

aberrations. Then, Section 5.2 presents a particular approach used to acquire one of the essential elements of a closed loop: the control matrix. Finally, Section 5.3 scans the space of parameters that determine the performance of the control loop, in order to determine the values that minimise the final residual aberrations, and then presents the convergence of those aberrations and of the SR.

5.1 Strehl ratio measurements

The ultimate goal of AO is to maximise the angular resolution of an optical system. Because of this, and since WFSs have biases, the final metric of performance should always observe that system's PSF. Multiple metrics exist that can assess the quality of a PSF, such as encircled energy or FWHM, but out of these, the SR tends to be a favourite as it can easily be translated into a RMS phase value using Maréchal's approximation (Rigaut et al., 1991; Roddier, 1994). This value can then be compared against the one found by the WFS to estimate the NCPAs. Unfortunately, accurately measuring the SR is difficult and many methods exist which do not necessarily agree with each other (Roberts Jr et al., 2004). In the worst case scenario, for Nyquist sampled data, the fractional error between different methods is contained within 20%. But when the sampling is doubled, the error drops to less than 5%, making the estimation much more accurate. A useful aspect of this ratio is that the methods to calculate it are self consistent, meaning values are meaningful relative to each other, i.e. if a SR value is larger than another obtained with the same method and for the same system, then that PSF has fewer aberrations. This allows the user to unequivocally assess the progress of an AO control loop, simply by comparing the SR at different times. Because of these advantages, an approach was developed to retrieve the SR from CHOUGH's NFSI camera, in order to independently measure the performance of the control loop. This section presents this approach, which will be used in a later section to optimise the implementation of the closed loop.

The first step in estimating the SR is to build a synthetic PSF, with no aberrations. This perfect PSF will serve to compute the denominator in the SR calculation, i.e. the maximum intensity value of the PSF with no aberrations. In order to build this PSF, a zero-padded discrete model of the pupil is made. This requires knowledge of the pupil's geometry and of the sampling on the imaging camera, i.e. the pixel scale in units of λ/D . This model is presented in Figure 5.1. The PSF is obtained

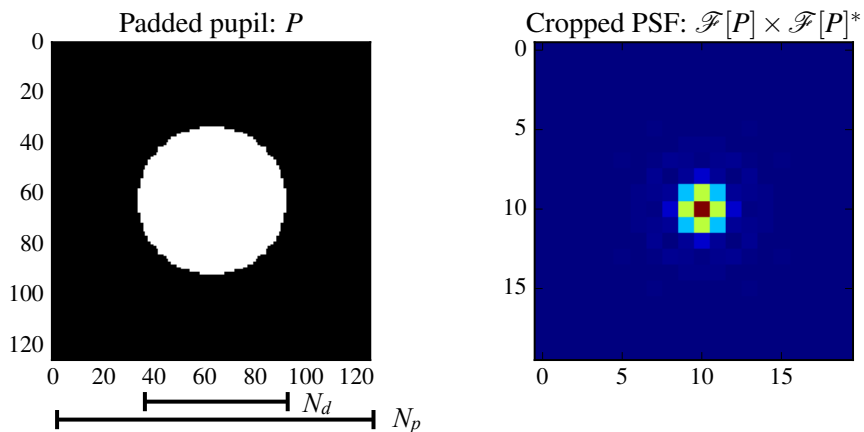


Figure 5.1: (*left*) Zero-padded discrete model of the pupil and (*right*) a cropped image of the resulting PSF, with values in arbitrary units. The binary model of the pupil has ones inside of it and zeros outside. N_d and N_p are the number of array cells across the diameter of the pupil and across the entire array once zero-padded, respectively. On the PSF image, array elements have the same plate scale as pixels in the NFSI camera.

by computing the Fourier transform of the discrete model of the pupil and then by multiplying the result with its complex conjugate.

In the resulting PSF, the pixel scale is controlled with the pupil's zero padding. This is because zero padding a function changes the spatial frequency sampling of its Fourier transform. Let $p = N_p/N_d$ be the pupil's padding factor. For the synthetic PSF to have the same pixel scale as real images, the padding factor must be

$$p = \frac{\lambda f}{D \mu_p}, \quad (5.1)$$

where f is the imaging system's effective focal length, D is the diameter of its entrance pupil, and μ_p is the camera's pixel size. For the NFSI $p = 2.095$. In

the figure $N_d = 60$, which results in the total number of array cells to be rounded to $N_p \simeq 126$.

Once the synthetic PSF has been produced, the next steps are to acquire an image with the camera, subtract the background, equalise the total energy of the synthetic image with that of the real one, and then finally centre the camera image. This last step of centring the image guarantees that the peak value of the PSF is centred around a single pixel, rather than around the junction between 4 of them. The way the centring algorithm works is first by cropping the image around its brightest pixels. For a cropped image of size $M \times M$, M needs to be large enough that most of the energy of the PSF will be included in it, for example $M > 40p$. After cropping the image, a ramp is fitted to the phase of its Fourier transform and the result is subtracted from the phase. The slope of the fitted ramp is bound to within $\pm\pi/N_c$, where N_c is the number of pixels in the cropped image. This ensures that when the ramp is subtracted the PSF only shifts by ± 0.5 pixel. After the ramp has been subtracted, then the inverse Fourier transform is applied to return to the image plane.

The reason for performing this centring is to not underestimate the SR. This is specially important at lower PSF sampling (i.e. larger pixel scales), where fewer pixels are used to sample the PSF. In this case, if the PSF is centred around a junction between pixels, then the observed peak value is lower than if it was centred on top of a pixel. If the peak value of an image drops and this value is used to calculate the SR, then the SR is underestimated by a proportional amount. This phenomenon is illustrated in Figure 5.2, where a synthetic image is intentionally uncentered, and, inversely, a real NFSI image is centred using the centring algorithm. In the figure, the centred synthetic PSF has been normalised to have a peak value of 1 and the uncentered version, which has the same total energy, has a peak value 27% lower. Since the peak value drops by 27% at most for the sharpest possible PSF, then in real images without centring it should drop by anywhere between 0% and 27%. Indeed, in the two bottom panels of the figure which show an example

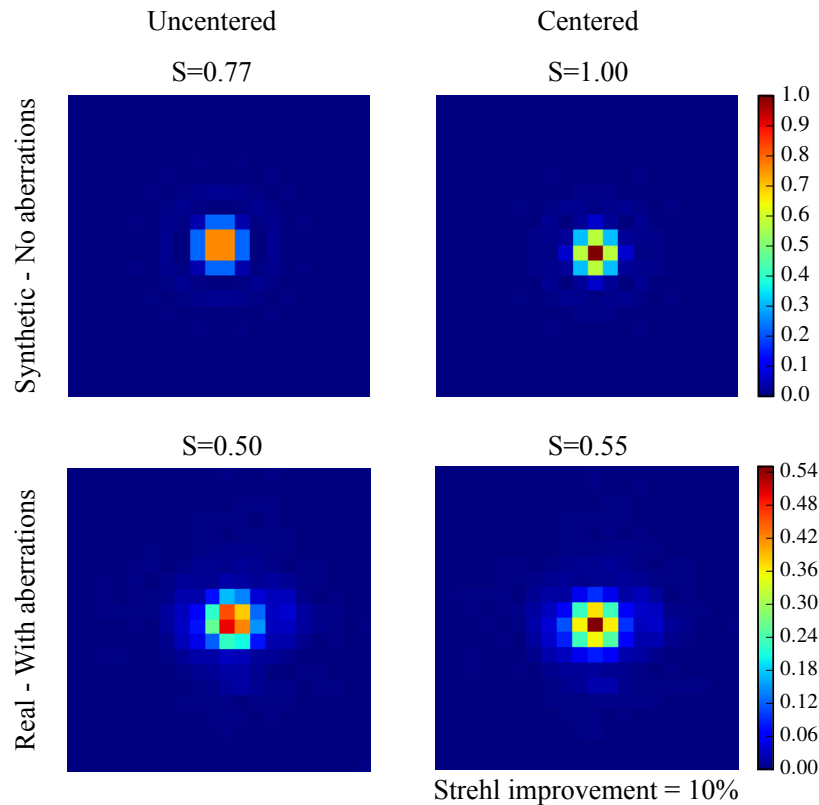


Figure 5.2: (*top-right*) Synthetic PSF with no aberrations, (*top-left*) same PSF intentionally centred around the junction between four pixels, (*bottom-left*) real image from the NFSI, and (*bottom-right*) real image, numerically processed to have its peak value centred around a pixel. All images have been normalised so that the synthetic centred PSF has a peak value of 1, and also so they all have the same total energy, i.e. the sum of the intensity.

with real images, the SR estimation increases by 10% after centring, demonstrating the correcting potential of the centring algorithm. Once the the image has been centred and its total energy has been made equal to that of the synthetic PSF (or vice versa), the ratio between their brightest pixels directly yields the SR.

Estimating the SR is not necessary to running an AO control loop, but is instead a useful diagnosis tool to measure its performance. The following section presents a method used to retrieve a component which is essential to closing the loop: the control matrix.

5.2 Control matrix

In order to run an AO control loop that corrects for wavefront aberrations, DM actuator commands are computed from WFS measurements, whether it is an open or a closed loop. One of the fastest ways to compute actuator commands is by using a linear operator known as the control matrix. As shown in Eq. (5.2)

$$C_{\text{mat}} \cdot \vec{s} = \vec{a}, \quad (5.2)$$

the control matrix C_{mat} is multiplied directly by the wavefront sensor data in vector \vec{s} to obtain the actuator commands in vector \vec{a} . This method, which works under the assumption that within a given range the interaction between the DM and the WFS can be mapped linearly, has the advantage of being simple to implement and fast to compute.

A simple way to acquire the necessary data for the production of a control matrix is to poke every actuator individually and to record the corresponding WFS response. All actuator commands are horizontally concatenated into an array A so that

$$A_{L \times M} = [\vec{a}_1, \vec{a}_2, \dots, \vec{a}_M], \quad (5.3)$$

where L is the number of actuators and M the number of commands sent. In this case $L = M$. For every command sent, the WFS measurements are recorded and concatenated in an array S so that

$$S_{N \times M} = [\vec{s}_1, \vec{s}_2, \dots, \vec{s}_M], \quad (5.4)$$

where N is the length of a WFS measurement vector \vec{s} . In the case of a SH, N equals twice the number of valid subapertures. This is because every subaperture slope has a horizontal and a vertical component. For CAWS, the equivalent of a subaperture, which is called a phase point, only delivers one number: the average phase at that point. As a result, for CAWS N equals the number of phase points. Having defined arrays A and S , it becomes possible to extend Eq. (5.2) into

$$C_{\text{mat}} \cdot S_{N \times M} = A_{L \times M}. \quad (5.5)$$

Given that every actuator has been poked, there is sufficient information to calculate C_{mat} by minimising the least-squares error

$$\epsilon = \left\langle \|C_{\text{mat}} \cdot S - A\|^2 \right\rangle, \quad (5.6)$$

leading to

$$C_{\text{mat}} \simeq A \cdot S^t \cdot (S \cdot S^t)^{-1}, \quad (5.7)$$

where $S^t \cdot (S \cdot S^t)^{-1} = S^\dagger$ is considered the *pseudo-inverse* of S (Herrmann, 1980; Gendron and Léna, 1994; Hardy, 1998).

When only individual actuators are poked one after another, the result is a sparse array S , which only contains information relating to a single actuator and to how many phase points were activated by it per column (usually 4 phase points per actuator). If instead, all actuators are poked for every command \vec{a}_k , then in turn every measurement \vec{s}_k contains information relating to all of them. In this case, actuators are not idly waiting for their turn to be poked, but are instead constantly contributing to a fuller array S and to a final control matrix with a higher SNR. A method that pokes all actuators simultaneously is potentially a time-efficient one.

A standard approach to poke all actuators at the same time is to excite pupil modes rather than individual actuators. In the case of circular pupils, the DM can be commanded into individual Zernike modes. Unfortunately this poses a problem when using a Boston Micromachines Kilo DM, as is the case with CHOUGH. In this DM, the edges of the mirror are glued to the mount (Evans et al., 2006). As a result, these edges can not fit the edges of Zernike modes, which is where they tend to reach their higher amplitudes, leading to a misrepresentation of the these modes. Hence, this approach is unsuitable for this system.

The alternative is to poke different temporal modes with each actuator. In this case, every actuator is assigned a specific temporal frequency at which it oscillates sinusoidally (Kellerer et al., 2012). Because every actuator oscillates at a different frequency, the effects of each actuator on the WFS data can be identified by filtering

for that actuator’s frequency. However in practise, explicit filtering is not required as it will happen implicitly when *pseudo-inverting* S .

For the control matrix to capture as many different combinations of interactions between any given actuator and the ones directly around it, temporal frequencies need to be assigned so that adjacent actuators have as widely different frequencies to each other as possible. Note that the interactions occur due to the continuous nature of the mirror’s face sheet, which couples actuators depending on their distance. A way to capture all possible combinations is by assigning frequencies in a checkerboard pattern, where the actuators falling on white spaces have ‘fast’ frequencies, while the actuators in black spaces have ‘slow’ frequencies. The frequencies assigned using this pattern are presented in Figure 5.3. Not all actuators




Figure 5.3: Map of temporal frequencies assigned to active actuators for the construction of the control matrix. There are 812 active actuators, with cycles per period of array A from 1 to 812. Inactive actuators are masked.

are assigned a frequency in order to keep the number of commands M in the poking sequence described by A short. Only the actuators that can influence the wavefront within the circular pupil are considered active and assigned a frequency other than zero. Inactive actuators are masked.

Regarding the length M of the poking sequence, it is proportional to the number of active actuators L_{active} , which in this case equals 812. In the discrete Fourier

domain, for every frequency to be distinguishable from others, they all must fall within distinct array cells. From here it follows that $M \geq 2L_{\text{active}}$. The transpose of the resulting poking sequence A is shown in Appendix B, Figure B.1.

Once the command sequence A has been sent and the corresponding WFS response S has been recorded, the control matrix can be computed as previously described. The control matrix that is later used to close the control loop can be found in Appendix B, Figure B.2. Once built, assessing the quality of the control matrix, i.e. how well it will perform, before actually testing it in a control loop is difficult. Because of this, it is also difficult to estimate and tune in advance the parameters involved in its construction. The following section introduces these parameters as well as those involved in closing the control loop. It then describes the process of scanning the parameter space for the values that minimise the loop's residual error, and finally presents the results obtained when running this optimised system.

5.3 Closing the control loop

As mentioned at the beginning of this chapter, the CAWS will be tested against CHOUGH's static aberrations, providing an initial assessment of this WFS-DM pair in closed-loop. There are two parameters regarding the construction of the interaction matrix that require tuning in order to optimise the performance of this particular control loop. Note that the following list is neither exhaustive nor universal: on the one hand, some parameters, such as the light source's luminosity or the exposure time, do not require fine tuning in the current lab setup, and on the other hand, different types of AO systems will consider other parameters, such as the temporal and spatial domains assigned to different DMs in a system with multiple ones.

The first one of these parameters is the actuator stroke used to poke actuators for the construction of the control matrix. For continuous DMs, inter-actuator coupling increases with the gradient between actuators and reduces their linearity.

This is evermore so for high-order DMs with their edges glued to the frame as is the case with CHOUGH's. In order to avoid this problem and produce a control matrix that describes the system in its linear regime, the stroke can be chosen to be very low. Unfortunately, if the stroke is too low, then this will also reduce the SNR of the WFS and by extension the final accuracy of the control matrix. The challenge is then to select a stroke that will produce an accurate and linear description of the system.

Once the actuators have been poked, the array S needs to be *pseudo-inverted*. A standard numerical method for this is singular value decomposition (SVD). This method is conditioned by a parameter R_{cond} , the second one in the list of parameters to tune, which acts as a threshold under which singular values, normalised to the largest one of them, are ignored. This thresholding filters singular vectors with low SNR. If noisy vectors are allowed, then the control matrix also has a lower SNR, resulting in a less precise AO compensation of aberrations. If too many vectors are filtered out, then information is lost from the control matrix, making it insensitive to some modes. The result is a slow convergence and larger final residuals.

It is difficult to accurately predict the effects the two aforementioned parameters will have on the performance of the control loop. For example, depending on the precise alignment of the WFS with the DM, the last one can produce out of phase near-Nyquist frequency spatial modes that are invisible to the WFS. Usually known as waffle modes, they can appear either in contained areas of the pupil or across the entire pupil. This depends respectively on whether the WFS slightly over samples the DM, as is the case here, or samples it exactly at the Nyquist frequency, as is the case for most SHs. The best way to determine the performance of the control loop in the presence of complex phenomena such as this one, so that parameters such as R_{cond} and the DM's stroke during the construction of the control matrix can then be tuned, is to scan the parameter space producing several matrices, test them in closed-loop, measure the residual aberrations and compare them.

The way the control loop is implemented is by, for every control frame k , computing

actuator commands \vec{a}_k as

$$\vec{a}_k = \vec{s}_{k-1} + \alpha C_{\text{mat}} \vec{s}_k, \quad (5.8)$$

where α , known as the loop's gain, is a constant value between 0 and 1. Similarly to R_{cond} thresholding, when α is near 1, the system converges quickly at the expense of some stability, which is relevant when dealing with quickly evolving atmospheric turbulence. But when dealing with static aberrations, as is the case here, the convergence of the loop can be slowed down by reducing α , thus guaranteeing its stability. In the current setup, the loop's gain does not require a fine degree of tuning, and only a coarse scan of this parameter will be performed. In summary, the parameters to tune are the poking stroke of the DM actuators used to build the control matrix, the conditioning value R_{cond} used for the *pseudo-inversion* of S , and the gain of the control loop α . The criterion used to choose the best tuning is the residual RMS error measured by the CAWS after the loop has converged and averaged over 10 frames.

After closing the control loop with the HeNe laser and performing the parameter space scan, the lowest final residual RMS error was achieved with $\alpha = 0.85$, a poking stroke of 18 nm, and $R_{\text{cond}} = 0.01$. Figure 5.4 shows the singular values of array S obtained with these values, sorted in decreasing order, and the threshold set by the conditioning. As can be seen in the figure, 779 out of 1020 (total number of active and inactive actuators) singular values are kept after thresholding. This means that the majority of singular vectors contribute more useful information than they do noise, which is a sign that the CAWS measures signals from the DM with a high degree of fidelity.

Figure 5.5 shows the phase measured with the CAWS before and after closing the control loop and waiting until it converges. As can be seen in the figure, CHOUGH's initial static aberrations, as measured by the CAWS, have RMS and PtV values of 271.5 nm and 1.4 μm respectively, which in units of wavelength cycles are 0.42 and 2.2, for 633 nm. With respect to the PtV values, they approach the maximum PtV




Figure 5.4: Normalised singular values, sorted from highest to lowest. The *left* panel shows the 15 largest values in detail and the *right* panel shows the rest.

stroke of the DM of $1.5\ \mu\text{m}$. Furthermore, the largest amplitudes are concentrated around the edges of the pupil: those of negative sign are around the top and bottom while those of positive sign are around the left and right sides. Because of this and since the DM is glued around the edges, this is where most residuals are still present after closing the loop and letting it converge. If the entire pupil is considered, then the final RMS value of residual aberrations is $55.3\ \text{nm}$, or 0.08 wavelength cycles. But if instead the edges where the DM could not remove all the aberrations are excluded, then this value drops by more than a factor of 4 down to $12.5\ \text{nm}$, or 0.02 cycles. This is equivalent to sensing and correcting an offset of 2% RMS in the position of the interference fringes, which is very encouraging for a laboratory demonstration setup.

Although positive, these results are incomplete because there is a disadvantage in only using residual aberrations measured with the CAWS to determine the performance of the control loop. The problem is that if there are any aberration modes which are invisible to the CAWS, then the system would only correct the visible ones, leaving behind those that are invisible. Furthermore, the DM could be producing aberrations in these invisible modes, actually feeding into the system's total aberrations. Since the ultimate goal of AO is not to reduce the aberrations mea-

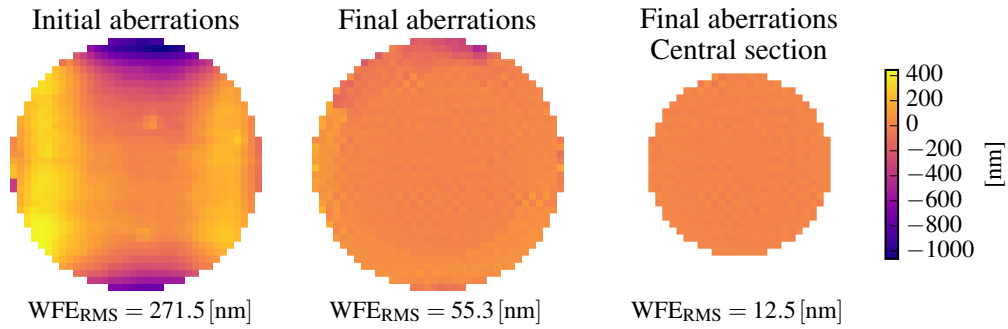


Figure 5.5: CAWS phase measurements of the initial static aberrations (*left*), the final aberrations the system converges to after closing the control loop (*centre*), and a section at the centre of the pupil for those central aberrations (*right*).

sured with a WFS, but rather to increase the angular resolution and brightness of focal plane images, then the best way to assess the performance of the control loop is by directly observing the value of metrics describing the quality of these images. This can be achieved by acquiring images with CHOUGH’s NFSI camera during the operation of the control loop, simultaneously to those of the CAWS, and then estimating their SR using the method described in Section 5.1. Figure 5.6 shows NFSI images from before and after closing the control loop. The SR of 0.66 in the




Figure 5.6: NFSI focal plane images with (*left*) the initial static aberrations, (*centre*) the final aberrations the system converges to after closing the control loop, and (*right*) the synthetic PSF used as reference to calculate SR, for comparison. All three images have been normalised to have the same cumulative intensity, and have been centred around a single pixel.

central panel is the highest one achieved with the monochromatic HeNe source, after tuning the control matrix and letting the loop converge. The increase in SR from the original 0.22 in the left panel of the figure confirms that the control loop

is indeed removing seen aberrations overall.

As described earlier, an advantage of the SR is that it can be used to estimate the RMS aberrations through the use of Maréchal's approximation. This way the aberrations on the NFSI's path can be compared to those on the CAWS'. Figure 5.7 shows the temporal evolution of both the RMS aberrations and the SR after closing the loop, as estimated using both the CAWS and the NFSI. By subtracting the




Figure 5.7: Temporal evolution of (*black*) residual RMS aberrations and (*blue*) SR from the moment the control loop is closed until it converges. Data with the same type of marker are equivalent and were converted from (*black*) RMS aberrations to (*blue*) SR, or *vice versa*, using Maréchal's approximation. The origins of the data are (*solid line*) CAWS phase measurements including tip-tilt, (*dotted line*) the same measurements excluding it, and (*dashed line*) NFSI focal plane images.

tip-tilt mode from the CAWS' phase measurement and then computing the RMS value, we can estimate the SR in its path. After converging, this SR is around 0.81, compared to the 0.66 seen by the NFSI. This is a discrepancy of 0.15. Conversely, the NFSI's SR becomes a RMS aberration of 65 nm, compared to the CAWS' 45 nm after tip-tilt subtraction. The 20 nm difference observed between both paths is relatively constant after 5 iterations. This, plus the fact that the aberrations are greater on the NFSI's path suggest that the difference is mostly due to NCPAs and low-order aberrations that are invisible to the CAWS.

The final detail to note in the figure is that before the loop is closed, there is a difference of about 150 nm between the CAWS's measurements and the NFSI's estimation, instead of the 20 nm observed near the end. Given that the $SR=20\%$, this is due to Maréchal's approximation losing validity. As shown by Mahajan (1983), at low SR Maréchal's approximation underestimates the RMS value of aberrations. The effect is small for primary aberrations, but becomes bigger if they are more complex. The way to confirm this is by using the phase measurements to reconstruct the PSF and then calculate the SR. The resulting SR is around 0.17, which is orders of magnitude away from the almost 0.0006 that was originally estimated from the phase (using the approximation), and much closer to the 0.2 from the NFSI image.

After successfully tuning and closing the loop with monochromatic light, the same task was attempted with polychromatic broadband light produced with a halogen lamp. This is the same halogen lamp that was characterised in Chapter 4. Results are presented in Figure 5.8. Unfortunately it was not possible to feed the exper-




Figure 5.8: Temporal evolution of (*black*) residual RMS aberrations and (*blue*) SR from the moment the control loop is closed until it converges. These results correspond to the same experiment presented in Figure 5.7, but with a broadband halogen lamp rather than a HeNe laser. The equivalent central wavelength is 667 nm and the bandwidth is 76.5 nm.

iment sufficient light using single mode optical fibres. Instead, the fibre that was used had a core just larger than the NFSI's diffraction limit. On the one hand, this reduces the maximum SR attainable on the imaging camera when compared to a diffraction-limited synthetic PSF. As can be seen in the figure, with this fibre the SR converges to 0.35, rather than 0.66 with the monochromatic source and the single mode fibre. On the other hand, using a larger fibre core also reduces the visibility of interference fringes. This happens because a finite size fibre can be modelled as a collection of point sources continuously distributed across the surface of its core, each producing a wavefront with a different tip and tilt. Different tips and tilts result in different sets of fringes with peaks and troughs that coincide less with each other as the fibre becomes larger, leading to a decrease in visibility. This decrease, in addition to other phenomena related to the broadband nature of the light source, could explain the slower convergence of the WFE measured by the CAWS, even after the value of the SR has already converged. At the time of writing this is still not well understood and new experiments with an improved diffraction-limited fibre core and broadband light source are necessary to elucidate whether the slower convergence of the WFE are, indeed, caused only by the larger fibre core, or if there are other reasons behind this. Nonetheless, the experiment is successful at increasing the SR on the focal plane, hence demonstrating that the control loop can be closed with the CAWS using polychromatic light.

This chapter showed that it is possible to close an AO control loop with the CAWS, in order to correct for static aberrations in monochromatic and polychromatic conditions. For this purpose, it was necessary to build and tune a linear control matrix mapping the interactions between the CAWS and CHOUGH's DM. Once the loop was closed, its progress and performance were monitored independently by measuring the resulting SR on the NFSI imaging camera. Increases of the SR from 0.2 to 0.66 and from 0.12 to 0.35, with monochromatic and broadband light respectively, confirm that aberrations were being successfully removed from focal plane images, leaving behind small NCPAs, low-order aberrations, and aberrations

out of the reach of the DM.

On-sky tests

The ultimate validation for any astronomical instrument comes from testing it on a telescope, with real on-sky targets. This is because, despite all the rigorous lab experimentation, there will always be conditions that are either difficult to reproduce or that are just unforeseen. Originally, on-sky tests were not envisaged to be a part of this thesis. But in an unexpected turn of events, the CANARY (Myers et al., 2008) experiment at the William Herschel Telescope (WHT) received funding from the OPTICON program (Gilmore, 2004) to serve as an AO testbed for new instrumentation concepts. We were awarded observation time on this new CANARY run through a joint proposal with the Centre de Recherche Astrophysique de Lyon (CRAL), in order to test CAWS alongside their own novel PDI, the integrated Mach-Zehnder (iMZ) interferometer, also known as the solid Mach-Zehnder (Loupias et al., 2016). The iMZ is a self-referencing Mac-Zehnder, as described in Section 2.1.2, where the splitting of the beams, the pinhole, and the interference of the beams, all occurs inside a transparent, solid optical component. This ensures that both the reference and the test beams experience the same vibrations and no differential atmospheric turbulence, resulting in an enhanced fringe stability and visibility.

This chapter is devoted to presenting the CAWS' telescope run on CANARY, which is divided into two major phases. The first phase consists of off-sky measurements

of quasi-static aberrations and their compensation. This phase is a necessary prerequisite before moving onto the second phase: the measurements of on-sky post AO residual aberrations, which were performed between the nights of the 17th and the 21st of July, 2019. The purpose of this phase is to demonstrate that CAWS is able to handle the dynamic range of these aberrations, all while being fed polychromatic light, and can therefore be implemented on an AO system.

Section 6.1 lists the objectives of the on-sky mission and explains the advantages of acquiring joint measurements with a second PDI. Section 6.2 then goes on to introduce the experimental setup and describe the different parameters and phases of observation. Finally, Sections 6.3 & 6.4 present respectively the results of the quasi-static pre-compensation and those of the on-sky closed-loop measurements.

6.1 Objectives

The purpose of the mission is to expose the two new WFSs to evolving distributed turbulence, to the quasi-static and closed-loop residual aberrations of an actual AO system and to the telescope operational environment. This last factor could include, a variety of phenomena such as telescope vibrations, changes in gravity vectors and flexions, pupil shifts induced by the derotator, temperature variations, dome turbulence, rotating spiders, the low wind effect, atmospheric chromatic dispersion and the stars' spectral profile, amongst others. Simulating comprehensively and precisely all of these phenomena in the lab can be difficult or even impossible. For this reason, submitting the instruments to these conditions is an essential precursor to any planned deployment.

As mentioned earlier, the observations will be used to simultaneously test two new WFS techniques: the iMZ and the m-PDI through its incarnation, the CAWS. Since both these instruments belong to the PDI family and share many similarities, they have comparable requirements for their respective validations. For example, they both need quasi-static aberrations to be compensated and tip-tilt atmospheric

turbulence to be corrected in real-time by a closed-loop system. This means they can be tested at the same time by splitting the light's path, which allows us to streamline many aspects of the preparations and observations. In turn, this saves on installation costs and reduces the necessary off and on-sky time per instrument by about half.

Another advantage of having a joint proposal is it allows us to establish an excellent contemporaneous comparison between the two techniques. Especially on-sky, all observations can be synchronised so they are submitted to the same atmospheric turbulence and downstream AO corrections. Hence, along with the cost and time gains that it brings about, a secondary objective of this double on-sky run is to compare the sensors to each other, in order to highlight the specificities and strengths each of them have. This comparative analysis is still at an early stage, and consequently none of the data or early results are included in this thesis.

Before more specific objectives can be laid out, CANARY's maximum performance and the limitations they bring about for PDIs must be considered. One of the most influential factors determining the proper functioning of a PDI is the SR being delivered to it. The fraction of light going through CAWS' pinhole depends upon how high the SR is. This in turn affects the total throughput and fringe visibility. In the case of CANARY, the values it can reach on-sky and in the visible are very low, of up to 0.5% at 675 nm (Morris et al., 2014). The system's SR is low enough, that it puts into questions whether some goals are even achievable. Some modifications addressing the low SR were contemplated, but as it will be shown, they are either too expensive or involve a complete redesign of the instrument.

The simplest way to get more light through the pinhole is for it to have a larger diameter, relative to λ_0/D , where D is CAWS' entrance pupil diameter. This can be done by either manufacturing a new focal plane mask with a larger pinhole or by enlarging D . Unfortunately, manufacturing a new focal plane mask is too costly for the project's budget. On the other hand, having a larger entrance pupil seems simple at first, requiring a simple redesign of the relay preceding the CAWS, but it

has the undesired consequence of changing the instrument's spatial sampling, which reduces the photon count per line-pair. The later can be fixed by increasing either the period of the diffraction grating or the focal length of CAWS' first lens. The first is difficult because the one being used is already amongst the longest periods commercially available, and the second requires a complete redesign of the optics and focal plane mask. Furthermore, changing the period of the grating would also require manufacturing another focal plane mask, defeating the purpose of simply having a larger pupil.

Another way to get more light through the pinhole is to increase the SR by moving into longer wavelengths at the edge of the visible spectrum and beginning of the NIR. In H band, CANARY can produce SRs of up to 40% (Morris et al., 2014). In order to use a longer wavelength two changes are needed. The first one is to rebuild the focal plane mask which, as stated earlier, is unfeasible due to budget restrictions. The second change is to find an already purchased and available camera that can detect light in the NIR. Buying a NIR camera is even more expensive than making a new focal plane mask, and is therefore discarded as an option. Having found no spare cameras and not wanting to waste the opportunity to go on sky, it is important to determine what can still be achieved by doing so with the existing design as it is.

The objectives stated at the beginning of this section were to submit the CAWS' to the quasi-static aberrations and NCPAs of an actual AO system and also to its closed-loop residual aberrations. Out of these two, the least problematic in terms of SR are the quasi-static and NCPAs. The reason for this is that quasi-static aberrations can be initially mitigated using focal plane wavefront sensing such as Phase Diversity. Other constraints can also be relaxed when working in a controlled laboratory environment without turbulence. For example, if the SR remains relatively low after iteratively optimising the PSF, more light can be made to go through the PDI's pinhole by increasing either the internal light source's brightness or the exposure time of its detector. Since the sensing of quasi-

static aberrations is likely to be successful, we must now turn our attention to the difficulties of on-sky measurements.

Given the low SR delivered by CANARY in the visible, the objective that remains the most ambitious is that of measuring on-sky closed-loop residual aberrations. Since all analytical models developed until this point consider small aberrations to some degree, and by extension a relatively high SR, it is difficult to predict the behaviour of CAWS when faced with strong aberrations and low SR. On the one hand, this means it is risky to go on-sky, as it is possible that not much will be observed. On the other hand, the opportunity is presented to test the CAWS in extreme conditions, exploring the circumstances under which our theoretical models break down and how to extend them to make them more robust. By doing so, new insights could be drawn that could help us more accurately determine the minimum requirements on any future hosting AO system, and in turn to find designs that better suit them.

The following section will present the experimental setup and determine the requirements that best allow us to test the CAWS under the adverse conditions described above. In order to maximise the chances of detecting fringes, the brightest possible targets will be selected and the appropriate exposure times will be estimated, that can achieve a minimum working SNR.

6.2 Experimental setup & requirements

As mentioned earlier, CANARY has an on-axis SH called the truth sensor (TS). This WFS has 14-by-14 subapertures and is designed to work in the visible part of the spectrum. In conjunction with a tip-tilt mirror and a 241-actuator ALPAO DM, these are the three main components to run the natural guide star (NGS) SCAO loop. As shown in Figure 6.1, our experimental setup taps into the TS's path with a 10R/90T beamsplitter; i.e. 10% gets reflected into the TS and 90% gets transmitted into our sensors. Because of this, both the TS and our experiments are in a closed-




Figure 6.1: CANARY’s SCAO loop, with CAWS, an imaging camera and the iMZ.

loop configuration and can therefore receive either uncorrected wavefronts, tip-tilt corrected wavefronts (by only using the tip-tilt mirror) or fully corrected SCAO residuals. This allows us to probe the performance and the utility of these PDIs under all these situations.

After the initial beamsplitter, light gets once again split by another beamsplitter, this time into two equal parts, where the reflected arm goes to CAWS. The transmitted arm is split for the last time, where 10% gets reflected into a PSF imaging camera and the rest goes to the Mach-Zehnder. This configuration allows us to perform simultaneous observations, with all 4 cameras (CAWS, iMZ, PSF imaging camera and TS) working in parallel. The list of cameras is presented in Table 6.1. An image of the setup which includes all of the instruments and their cameras, except for the TS camera, can be seen in Figure 6.2

Table 6.1: List of cameras by instrument.

Instrument	Camera
CAWS	Imperx Bobcat B0620M
iMZ	ANDOR Neo
PSF imaging camera	ANDOR Zyla 4.2
Truth Sensor	ANDOR iXon EMCCD




Figure 6.2: Image of the experimental setup, showing the CAWS, the iMZ, the PSF imaging arm, the beginning of the TS and a part of CANARY.

The first round of experiments, called Phase 0, measures NCPAs between CANARY’s own TS, and our sensors. This calibration round is performed off-sky, using CANARY’s internal sources and without turbulence. Amongst the many internal source provided by the CANARY bench, those used for alignment and calibration purposes are listed in Table 6.2. In the table, all of the fibre-fed sources are delivered to a focal plane within CANARY, previous to the tip-tilt mirror and to the DM, as if replicating a natural on-axis star. This way, they can be used to optimise the PSF on the imaging camera and then record reference slopes with the TS and corresponding reference fringes with the interferometers.

Table 6.2: List of internal sources.

Source	Delivery	Chromatic bandwidth	Angular size
HeNe Laser	Free-space propagation	633 ± 0.001 nm	Pencil beam
Laser diode	Single mode fibre	635 ± 5 nm	Diffraction limited
White light LED	Multimode fibre	Broadband white light	<i>Seeing</i> limited
Halogen lamp	Single mode fibre	Broadband white light	Diffraction limited

Once calibrations have been undergone, Phase 1 focuses on acquiring on-sky data at high speeds, between 2 Hz and around 200 Hz, with broadband light centered around the R band and with multiple WFSs simultaneously. Given the low throughput of CAWS, the high degree of multiplexing in the optical setup (four cameras observing simultaneously in the same band) and the short integration times, the most stringent constraint during this phase is integrating enough light. This in return, is equivalent to selecting a natural star of sufficient magnitude.

In order to calculate the the maximum required magnitude, the first thing is to determine the throughput of CANARY plus the telescope. This can be estimated by using the system’s optical model presented in Figure 6.3. To fit CANARY’s most important components, this image excludes the telescope and the derotator, which are included in the full model. In this figure, all of the surfaces between the telescope’s focal plane and the 10R/90T beamsplitter are mirrors, either flat or in the shape of off-axis parabolas, and with angles of incidence close to either 0° or 45° .

Unfortunately, the optical model does not explicitly state the mirrors’ coating. In this case, it will be assumed that the mirrors, including those of the derotator, use protected aluminium, due to its good reflectance in both the visible and NIR, as shown in Figure 6.4. Having good reflectance in the NIR is important for CANARY’s own PSF imaging camera CAMICAZ (Sivo et al., 2014). This data,




Figure 6.3: CANARY's optical model (CANARY team, 2019), from the telescope's focal plane to the first 10R/90T beamsplitter. The reflected arm containing the TS is included, whereas the transmitted arm alongside the iMZ and CAWS are not.




Figure 6.4: Reflectance of Protected Aluminium, for an angle of incidence of up to 12° , as provided by Thorlabs (2019).

although not presented here, is also offered for an angle of incidence of 45° .

The final assumption, is that the WHT has an overall constant throughput of 75% across the visible. With these approximations, the estimated optical throughput, starting from the telescope's aperture and going all the way through CANARY

until the back of the TS beamsplitter, is as shown in Figure 6.5. This is the




Figure 6.5: Estimated throughput of the WHT + CANARY + 10R/90T beamsplitter

fraction of the light delivered to the joint CAWS and iMZ experiment. Note that at CAWS' central wavelength $\lambda_0 = 675 \text{ nm}$, only about 6% of the total light makes it across, or approximately a 3 magnitudes drop with respect to the telescope's entrance pupil.

By using this estimation and including the throughput of the optical relay leading to CAWS, that of CAWS itself and the QE of its camera, the total sensitivity of this WFS is as shown in Figure 6.6. As can be noted, a very small fraction of the




Figure 6.6: Estimated throughput of the WHT + CANARY + CAWS, including camera's QE. Rapidly decreasing values after 750 nm are explained in part by the approaching sensitivity limit of silicon detectors in the NIR.

light actually makes it through to the detector. Again for the case of the central wavelength, only 0.6% of the light makes it to the detector. This is a factor of 10 from the TS beamsplitter to the camera, or a decrease of about 2.5 magnitudes for a total of 5.5 magnitudes between the telescope’s aperture and the backend of our instrument. It is important to bear in mind that the previous throughput estimation was computed for the case of very small aberrations, when the SR is near 1. In a system designed to have fringe visibility near 1 under no aberrations, when the SR drops to a value around 0, the throughput is halved, adding almost another magnitude to the attenuation (~ 0.75 magnitudes).

The next step in order to calculate the maximum required magnitude is to select the filters to be used. As presented in Figure 6.1, our experiment has a filter wheel with a variety of filters and a dichroic. The components on the wheel are listed in Table 6.3. These are all filters centred around the R band with various

Table 6.3: List of components by position on the filter wheel.

Position Number	Component
1	697 \times 75 nm filter
2	Empty
3	832 \times 37 nm filter
4	Cross target
5	700 \times 10 nm filter
6	> 625 nm dichroic

bandwidths. A larger bandwidth will let more light through, at the cost of decreasing the visibility of fringes. With a few filters rather than just one, it is possible to probe different configurations in order to find that which maximises the PDIs’ performance, depending on the conditions of *seeing*, AO residuals and instrument throughput. The cross target in position 4 is only there for alignment purposes.

The light that makes it through the filters will be sampled by a different number of pixels depending on two factors: the size of the pupil delivered to CAWS and the camera binning. The fewer pixels sample the pupil, the more light each pixel will receive and therefore the fainter the guide star can be. For simplicity of imple-

mentation, the design of CAWS was kept as presented in Chapter 4. So in order to reduce the size of the exit pupil on the detector, a 4F relay was used to reduce the size of the entrance pupil, as shown in Figure 6.7. The tradeoff is that reducing the




Figure 6.7: Optical model of the common-path interface, the optical 4F relay and CAWS.

size of the entrance pupil while keeping the same diffraction grating changes the instrument's spatial sampling. So the pupil size was reduced just enough so that CAWS still had greater resolution than the TS and could measure all the spatial frequencies produced by the 241-actuator DM. With a delivered beam diameter of 5.67 mm, there are 28.4 line-pairs in the pupil, resulting in the equivalent resolution of approximately a 19-by-19 SH (18.9 subapertures across). And since the line-pair detector sampling is still the same (6.9 pixels per line-pair for a $\times 1$ binning), then there are about 30,095 pixels sampling the pupil. These and other configuration parameters for CAWS are listed in Table 6.4.

Table 6.4: Configuration parameters for CAWS

Name	Value
Grid density	5 lp/mm
Entrance pupil diameter	5.67 mm
Line-pairs in pupil	28.37 lp/pupils
Equivalent SH sampling	19 \times 19 subapertures
Central wavelength λ_0	675 nm
Pinhole diameter	16 μm \Leftrightarrow 1.35 λ_0/D
Maximum fringe visibility	0.956 @ λ_0
Fringe sampling	6.9 pixels/lp
Pixels in pupil	30,095 pixels

All of the above allows for the calculation of the average flux per pixel. But, in

order to establish the amount of light that is required, a criterion needs to be defined that sets a minimum detection SNR. The detection SNR refers to that of the detector and not to that of the modulated signal. The detection SNR is preferred over the modulated signal's, specifically for the purpose of determining a minimum amount of light to be integrated. The detection SNR is defined as

$$\text{SNR}_d = \frac{N_\gamma}{\sqrt{N_\gamma + \delta S_R^2}}, \quad (6.1)$$

where N_γ is the number of photons received by a pixel and δS_R is the read-out noise of that pixel in electrons RMS. In this equation the detection or image signal is proportional to the number of photons detected, while the detection noise is a combination of the read-out noise and the shot noise. It will be established that the maximum required magnitude of a guide star will be that which produces a $\text{SNR}_d > 1$. Although at its limit this is still a very poor SNR, it is the bare minimum to unearth the image signal from the noise. For $\delta S_R = 16 e^-$ RMS, the read-out noise of CAWS camera, Figure 6.8 shows the maximum magnitude for different filters and as a function of the exposure time. Again, it is to be noted that the estimations in the figure consider a SR of near 1. For a SR close to 0, all these exposure times are to be doubled. Fortunately, this can be offset by a $\times 2$ pixel binning, which divides the exposure time by 2.9, while still providing enough fringe sampling to avoid incurring in aliasing and without any sacrifice to spatial resolution.

With this compensation at hand, filters 1 and 6 can be used to run CAWS at frame rates above 100 Hz with stars brighter than magnitude 0 and 1 in the R band respectively. This last restriction calls for a careful selection of the guide stars to be observed, since the number of stars with magnitudes around and above 0 are limited. Due to the lack of an atmospheric dispersion corrector, another important criterion is that the targets have a high elevation, $> 70^\circ$. This is in addition to other standard limitations, such as avoiding spectroscopic and highly variable stars. The list of candidates for observation and their elevations are shown respectively in

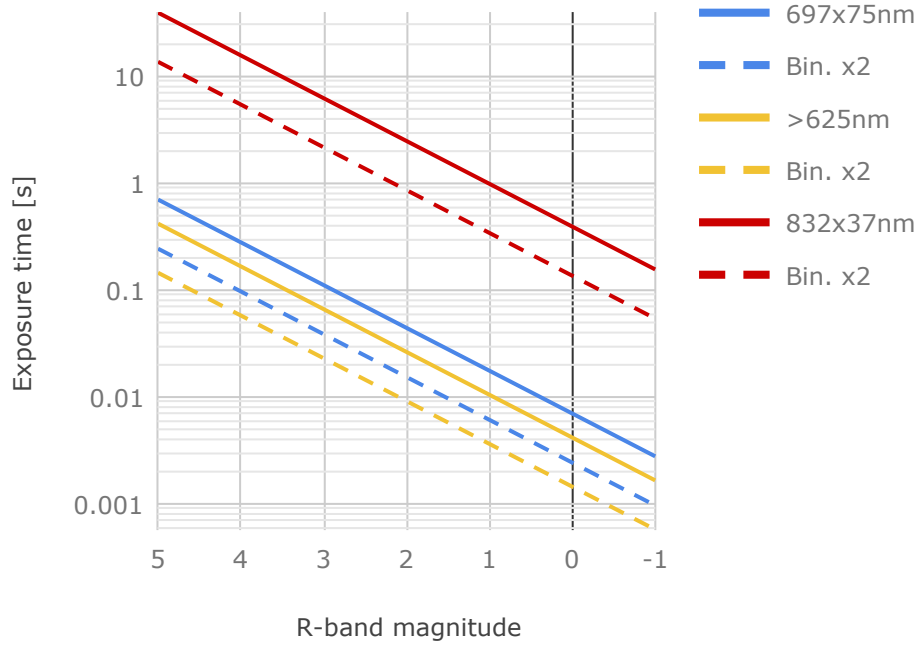


Figure 6.8: Minimum exposure time to achieve $\text{SNR}_d > 1$, for different filters. Solid and dash lines show the exposure time with a respective pixel binning of a $\times 1$ and $\times 2$.

Table 6.5 and Figure 6.9. As shown in the figure, these 4 stars reach altitudes

Table 6.5: List of target guide stars, their coordinates and magnitudes.

Identifier	Coordinates (ICRS, J2000/2000)	R band magnitude
alf Lyr	18 36 56.33635 +38 47 01.2802	0.07
alf Boo	14 15 39.67207 +19 10 56.6730	-1.03
alf Cyg	20 41 25.91514 +45 16 49.2197	1.14
bet Peg	23 03 46.45746 +28 04 58.0336	0.92

above 70° at different periods, allowing us to have a suitable target throughout most of the night time. Furthermore, the closest any of them come to the moon is 57° (bet Peg), which is far enough to have a sufficiently dark background in the R band.

This section presented the experimental setup and requirements to test CAWS on the CANARY system. By estimating the throughput of the WHT and CANARY, and combining it with that of the optical interfaces and CAWS, it was possible to

6.2. Experimental setup & requirements




Figure 6.9: List and elevation of target guide stars, as seen at the Roque de los Muchachos Observatory on the night of July the 17th, 2019. Star elevations processed by the Isaac Newton Group of Telescopes (2019).

determine the necessary integration times for different star magnitudes. This was also done for multiple configurations, using different filters and fringe samplings. Finally, these results were used to design an observation schedule based on 4 stars, which are bright enough to have fast frame rates of about 100 Hz or faster (depending on the configuration). The following section presents on-bench measurements using internal sources. The main purpose of these tests is to calibrate for NCPAs between the TS, the imaging camera and CAWS, and to characterise any residual aberrations. This maximises both interferometers' performance by focusing as much light as possible on their pinholes, and is necessary to distinguish between residual static aberrations and atmospheric turbulence.

6.3 Phase 0: Characterisation of quasi-static aberrations

Before the PDIs can be tested on-sky, the aberrations on the wavefronts reaching them need to be minimised and the reference slopes set on the SH. This will offload most of the NCPAs onto the SH which has a greater dynamic range, in turn enhancing the dynamic range of the PDIs. Since the imaging camera sits between the CAWS and the iMZ, by symmetry, minimising the static aberrations on this camera does not grant an advantage to any one PDI over the other. Therefore, it was agreed that quasi-static aberrations would be calibrated for by sharpening the PSF on that camera with an iterative algorithm provided by CRAL, and then recording references on all other WFSs.

After the PSF was sharpened, reference slopes were recorded on the SH and then used to close the control loop on an internal calibration source. This ensures that the DM keeps a stable shape, while the PDIs can record reference phases of their own. The reference phases on CAWS are as shown on Figure 6.10. The 635 nm




Figure 6.10: (*left*) Residual static aberrations on CAWS in nanometers and (*right*) corresponding PSF image measured on its focal plane mask, mode +1. Reference phases and focal plane measurements were produced using the 635 nm diode laser internal source.

diode laser source was used to record phases because its narrowband spectrum facilitates the productions of fringes in case the aberrations were still too large

to be measured with broadband polychromatic light. Furthermore, all the optics were aligned with this source, as it provides great flux and enough coherence to collimate beams using a shear plate. As a result, both PDIs should see the least amount of defocus at this wavelength, while other wavelengths should suffer from some chromatic aberrations. After acquiring references, an additional positive lens was added to CAWS in order to image a PSF on its focal plane. Since this PSF goes through the aperture on the mask corresponding to the diffraction mode +1, it is only being filtered at very high spatial frequencies, where almost no light remains after the initial sharpening, which has concentrated most light around the PSF's core and first few diffraction rings. As a result, it is almost identical to a PSF not being filtered around the edges. The SR measured on this image can therefore be used as a preliminary test to check that the instrument is receiving a good PSF and that most of the aberrations' power is being sensed by CAWS in its WFS configuration. Indeed, when the standard deviation of the wavefront is converted into a SR using Maréchal's approximation, the resulting value is 0.436, which only has a near 3% discrepancy with the SR calculated from the image, showing a good initial agreement.

Besides providing initial sanity checks and helping to make sure that residual static aberrations are small enough for the CAWS to be well within its dynamic range, the main purpose of these reference phases will be to distinguish atmospheric aberrations from static ones in on-sky measurements. These measurements will mostly be acquired using broadband filters, in order to maximise the amount of light reaching the instrument's detector. As seen earlier on Chapter 3, it is theoretically possible to produce accurate polychromatic phase measurements, given a known central wavelength and a relatively flat chromatic spectrum. But until now, there has been no experimental verification showing that the CAWS can actually do this.

A first validation that accurate broadband measurements are possible is to find a good agreement between the monochromatic and the broadband polychromatic characterisations of the same quasi-static aberrations. For this, the system is illu-

minated using the halogen lamp, delivered by a single mode fibre, and by using the dichroic on position 6 of the filter wheel. This filter is used as it allows through a broadband spectrum, which is also sufficiently confined for wavefront sensing to still be possible. Figure 6.11 shows the interference fringes obtained with both sources. In principle, both sets of fringes, monochromatic and broadband, should




Figure 6.11: Interference fringes obtained using (*left*) the monochromatic diode laser and (*right*) the broadband halogen lamp.

have the same visibility. But as will be explained in more detail later, the drop in fringe visibility when using the halogen lamp is due to chromatic aberrations.

In a perfect system, the phase measurements retrieved from the broadband fringes in radians can directly be scaled into units of length, e.g. nanometers, by using an equivalent central wavelength. Normally, the only other step in this process is to estimate said wavelength, by using the CAWS in its spectrometer configuration, described in Chapter 4, and measuring the light source's chromatic spectrum. Focal plane images for this purpose, with both the halogen lamp and the diode laser, are presented in Figure 6.12. As can be seen in the top panel of the figure, light from diffraction mode +1 has a larger vertical spread than seen in Figure 4.18, due to compound chromatic aberrations present in the optical relay preceding the CAWS and in the CAWS itself. As a consequence, some extra steps are required before the central wavelength can be estimated from these images.

As shown in Figure 6.13, there is a strong chromatic defocus between the wave-




Figure 6.12: Focal plane images of diffraction modes 0 and +1 with (*top*) halogen lamp and (*bottom*) diode laser for reference. Original values have been square-rooted to increase the contrast of the images presented in this figure.

length of the diode laser used to focus the optics (635 nm) and other wavelengths spanning the range of the CAWS. On the redder edge around 900 nm, the differ-




Figure 6.13: Zemax estimation of the differential chromatic defocus term in the CAWS' focal plane, between wavelengths spanning its chromatic range and the diode laser's wavelength (635 nm). The model includes CANARY and the optical interfaces leading to the CAWS.

ential defocus term can reach values around 400 nm rms. Due to this, it should

be expected that longer wavelengths transmit less light through the pinhole of the focal plane mask as they become more defocused. The transmitted light should therefore have a different chromatic spectrum from the one measured on mode +1, as it is weighted by a decreasing fraction of the total light. Since there is no chromatic dispersion on mode 0 (the mode reaching the pinhole), this spectrum can not be directly measured using the CAWS in its spectrometer configuration and, instead, needs to be reconstructed numerically.

In order to reconstruct the spectrum of mode 0, a simulation is performed that is based on the polychromatic simulation presented in Chapter 3 and combines three new elements. The first element is the monochromatic reference wavefront already presented in Figure 6.10, which is assumed to be the same for all wavelengths. The second element is the chromatic defocus obtained from the optical model. These two aberrations are added into the phase of an electric field which is propagated to the focal plane, producing a PSF. For every wavelength, the resulting PSF is multiplied by the binary mask representing the pinhole, allowing us to determine the fraction of total light intensity going through it. Finally, the third element, the spectrum measured by CAWS on mode +1, is weighted by the fractional intensities, hence reconstructing the spectrum of mode 0. Both spectrums, as measured on mode +1 and reconstructed for mode 0, are shown in Figure 6.14.

Also shown in the figure is the normalised amplitude of the interference fringes as a function of wavelength, which is calculated as

$$A_{fringe}(\lambda) = \left(\sqrt{I_0(\lambda)} + \sqrt{I_{+1}(\lambda)} \right)^2 - \left(\sqrt{I_0(\lambda)} - \sqrt{I_{+1}(\lambda)} \right)^2, \quad (6.2)$$

where $I_0(\lambda)$ and $I_{+1}(\lambda)$ are the intensity of light going through apertures M_0 (pinhole) and M_{+1} respectively. In the first place, this curve of amplitudes confirms that the fringes are indeed produced by broadband light, with a FWHM of 101.2 nm, or about a 14% bandwidth with respect to the CoM. Secondly and most importantly, by calculating the CoM, an estimate for the equivalent central wavelength is derived. As the relative contribution of the fringes increases with their amplitude, it

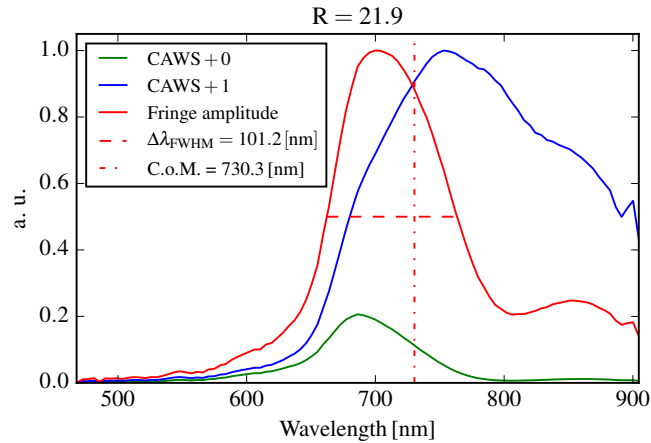


Figure 6.14: Chromatic spectra (*blue*) as measured on mode +1 and (*green*) as reconstructed for mode 0. The values of these curves have been normalised to the same reference so that the maximum common value between them is equal to 1. This normalisation preserves their relative magnitudes. Both spectrums are used to estimate (*red*) the normalised amplitude of interference fringes as a function of wavelength. The CoM and FWHM are presented for the fringes' amplitudes. R is the theoretical spectral resolution.

will be presumed that this value, located at 730.3 nm, is a useful representation of our target parameter.

After scaling the measurements with the estimated central wavelength, it still remains to subtract the differential defocus term. As mentioned earlier, chromatic aberrations add a differential defocus between separate wavelengths that need to be accounted for. The broadband references, corrected for differential defocus, are presented on Figure 6.15, side by side with the monochromatic references for comparison, from Figure 6.10, and the phase error between them. The error between the monochromatic and the broadband references is 23.6 nm rms, or about 26% of the total static aberrations measured with the diode laser. On the one hand, this is not a negligible amount and more study is needed before this error can be systematically accounted for in the instrument's error budget. Further tests could also show if at least part of this error arises from changing light sources on the SCAO loop, leading to a different performance of the loop. On the other hand, these initial results show that the general shape of the static aberrations can be correctly determined with broadband light by simply scaling the phase measurements with

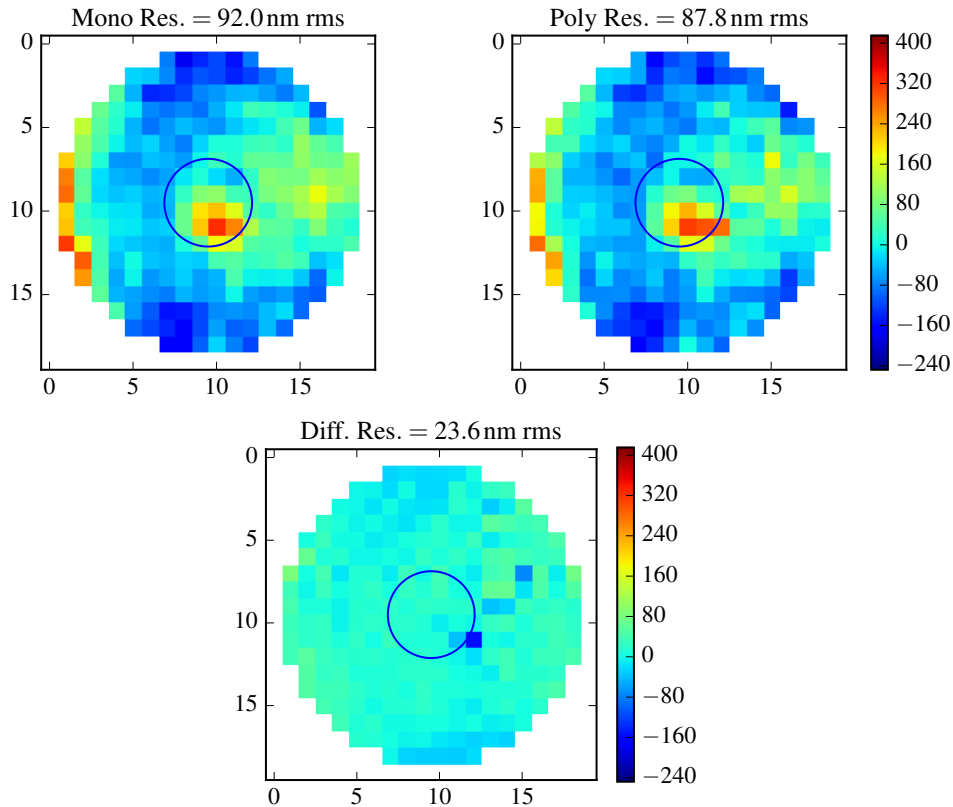


Figure 6.15: (top left) Monochromatic references, (top right) broadband references, corrected for differential defocus, and (bottom) error between them, all in nanometers. Blue central circles mark the location of the telescope’s central obscuration.

the appropriate wavelength and then correcting for the chromatic aberrations.

Despite the success of estimating the central wavelength by reconstructing the spectrum of the light going through the pinhole, it is still not known whether the value produced in this way is the best that could have been achieved. A potential approach to confirm this would be to answer the question of whether other central wavelengths yield a smaller error. In order to test this, the scaling and successive correction of chromatic defocus are performed with a continuum of wavelengths, spanning the whole of the instrument’s chromatic range. Figure 6.16 presents the error between monochromatic and broadband measurements as a function of wavelength. As can be seen in the figure, the curve has a minimum produced at 742 nm, which is only 1.6% away from the central wavelength found through the spectrum reconstruction method. In turn, the difference in RMS error between




Figure 6.16: RMS error between monochromatic and broadband measurements, corrected for chromatic defocus, as a function of central wavelength. Vertical dashed lines show (*blue*) the wavelength for minimum error and (*red*) the central wavelength found through the spectrum reconstruction method.

both wavelengths is less than 0.2%, making the result even more accurate.

The good agreement between the central wavelengths found through two independent methods, by spectrum reconstruction and by error minimisation, supports the hypothesis that the CoM of the interference fringes' spectral power minimises the error between monochromatic and broadband measurements, and is therefore the best possible equivalent wavelength. Although these methods have only been tested in a single scenario and only with one AO system, these initial results are promising and indicative that references measured with an internal light source can be used to calibrate measurements with natural stars, as long as the central wavelength with that star and filter pair can be well determined. Being able to do so is a pre-requisite for the correct scaling and defocus correction of on-sky measurements, such as those presented in the following section.

6.4 Phase 1: On-sky AO residuals

After successfully measuring and calibrating for static aberrations, the experiment is ready to enter Phase 1 and acquire on-sky data. But as mentioned earlier, producing fringes that have a sufficiently high SNR to be detectable will be challenging

even in the best conditions, given CANARY's low SR in the visible (0.5% at 675). As a consequence, the main focus of these observations will be to establish the operational parameters that maximise the appearance of fringes. This will be achieved by controlling and sweeping over two parameters, the chromatic bandwidth and the exposure time, all while keeping the main SCAO loop closed on a bright star. In the case of the chromatic bandwidth, the results will show the optimal solution when broadening the spectrum in the trade-off between allowing more light through to the detector at the cost of a more chromatic defocus and poorer visibility. With respect to exposure time, the trade-off is between the amount of light collected and the loss of visibility due to the averaging out of fringes. But in order to find these optimal values, a metric that is representative of the quality of the interference fringes needs to be defined.

A potential way to determine the quality of the interference fringes is by measuring their amplitude, relative to that of the background noise. This term, which will be considered to be the detection SNR of the fringes, is noted SNR_f and can be efficiently calculated in frequency space. By looking inside the sideband region, defined by the mask aperture M_{+1} , we should find the signal of the interference fringes

$$\begin{aligned}\hat{a}(k) &= \hat{I}_C M_{+1} \\ &= \Psi_{B,0} \otimes \frac{1}{\pi} \hat{\Psi}_{LP} \left(k - \frac{\lambda_0 f}{T} \right) \\ &\simeq \frac{b}{\pi} \hat{\Psi}_{LP} \left(k - \frac{\lambda_0 f}{T} \right),\end{aligned}\tag{6.3}$$

where b is the amplitude of light making it through the pinhole, $\hat{\Psi}_{LP}$ is the Fourier transform of the low-pass filtered electric field, T is the period of the fringes and f is the focal length of the lens preceding the focal plane mask (see Chapter 3). Using Parseval's theorem, the energy of the fringes can be calculated by integrating $\|\hat{a}(k)\|^2$ inside the sideband. In a noiseless scenario, this would be

$$A^2 = \iint_{M_{+1}} \|a(k)\|^2 dk.\tag{6.4}$$

Unfortunately, since the photon-shot noise and the read-out noise should be homogeneously distributed across the frequency space, $a(k)$ is inevitably found added to the Fourier transform of the sum of both noises, $\hat{n}(k)$. The resulting integral is of the form

$$\iint_{M_{+1}} \|a(k) + \hat{n}(k)\|^2 dk, \quad (6.5)$$

where the noise term is inextricably entangled to the fringe signal term in the computation of the absolute value.

This entanglement can be partially mitigated by estimating the noise term elsewhere in the Fourier plane and then subtracting it to the final integral. This is possible because, as mentioned earlier, the noise is homogeneously distributed, which means that integrating its power in two different regions of similar surface should, on average, yield the same value. The only condition is that the region where the noise is estimated does not have any signal coming from the fringes. The power of the noise in the sideband is

$$\delta N_{+1}^2 = \left\langle \iint_{M_{+1}} \|\hat{n}(k)\|^2 dk \right\rangle, \quad (6.6)$$

and the estimated power in a region $M_{\text{estim.}}$ with no signal from the fringes and of equal surface is

$$\delta N_{\text{estim.}}^2 = \left\langle \iint_{M_{\text{estim.}}} \|\hat{n}(k)\|^2 dk \right\rangle. \quad (6.7)$$

Values δN_{+1}^2 and $\delta N_{\text{estim.}}^2$ are in turn random variables that are, in principle, independent and have equal average and equal variance. Consequently, the compound random variable $\delta N_{\text{diff.}}^2 = \delta N_{+1}^2 - \delta N_{\text{estim.}}^2$ has an expectation of zero and a variance $\text{Var}(\delta N_{\text{diff.}}^2) = 2 \text{Var}(\delta N_{+1}^2)$. This new variable adds uncertainty to our further estimation of SNR. As a side note, in the case $M_{\text{estim.}}$ does not have the same surface as the sideband M_{+1} , the estimated noise power should be rescaled by using the ratio between both regions' surfaces, as will generally be the case with real data.

Having estimated the power of the noise, it can be subtracted to the integral of the

sideband to produce a good approximation of the total fringe intensity

$$\begin{aligned}
 A_{\text{estim.}}^2 &= \iint_{M_{+1}} \|a(k) + \hat{n}(k)\|^2 dk - \delta N_{\text{estim.}}^2 \\
 &= \iint_{M_{+1}} \|a(k)\|^2 + \|\hat{n}(k)\|^2 + 2 \operatorname{Re}\{a(k)\hat{n}^*(k)\} dk - \delta N_{\text{estim.}}^2 \\
 &= \iint_{M_{+1}} \|a(k)\|^2 + 2 \operatorname{Re}\{a(k)\hat{n}^*(k)\} dk + \delta N_{\text{diff.}}^2.
 \end{aligned} \tag{6.8}$$

In the previous expression, the conjugate value of the noise, $\hat{n}^*(k)$, can be considered to have a random angle, completely uncorrelated to that of the fringes' signal $a(k)$. As a consequence, $2 \operatorname{Re}\{a(k)\hat{n}^*(k)\}$ has a random sign, leading its integral to have a null average value. This, combined with the fact that $\iint_{M_{+1}} \|a(k)\|^2 dk$ grows with the square of the signal $a(k)$, whereas the standard deviation of $\iint_{M_{+1}} 2 \operatorname{Re}\{a(k)\hat{n}^*(k)\} dk$ can only grow proportionally to $a(k)$, and $\delta N_{\text{diff.}}^2$ has a constant variance, makes both these terms become relatively smaller as a part of $A_{\text{estim.}}^2$ for larger signals. This means there is a point where $A_{\text{estim.}}^2$ can be approximated to

$$A_{\text{estim.}}^2 \simeq \iint_{M_{+1}} \|a(k)\|^2 dk = A^2, \tag{6.9}$$

for a large enough signals.

Now that the intensity of the fringes can be estimated, it can be used to compute the empirical SNR of the fringes

$$\operatorname{SNR}_f^{(\text{empiric})} = \sqrt{A_{\text{estim.}}^2 / \delta N_{\text{estim.}}^2}, \tag{6.10}$$

which is an estimation of the true SNR value

$$\operatorname{SNR}_f^{(\text{true})} = \sqrt{A^2 / \delta N_{+1}^2}. \tag{6.11}$$

This empirical SNR is first tested by using the internal diode laser source and simulating SCAO residuals on the bench. In order to do this, residual slopes were recorded on a previous on-sky SCAO run. After subtracting the tip-tilt term, these aberrations are multiplied by the system's control matrix to produce DM commands. By passing these commands onto the DM, the residual aberrations can be emulated and replayed as they would have been seen at the backend of the

AO system, had the tip-tilt term been perfectly corrected for. In this controlled environment, aberrations have less power due to the subtraction of tip-tilt and the constraints on integration time and flux are relaxed, which allows us to acquire high SNR measurements.

An example of fringes obtained with this *replay* is presented in Figure 6.17, alongside the regions in the Fourier plane being used to estimate the empirical SNR of the fringes. These images were obtained using a 2×2 binning on the detector's




Figure 6.17: Off-sky (*left*) interference fringes with (*right*) their respective Fourier transform (the square-root of the absolute value is displayed for contrast enhancing purposes). Fourier transform panels also show (*blue*) the sideband region M_{+1} , (*green*) the noise estimating region $M_{\text{estim.}}$, and the resulting estimation of $\text{SNR}_f^{(\text{empiric})}$.

camera. As can be seen on the figure, the (*top-left*) image with the sharpest and brightest fringes, also has a brighter sideband (*blue* region), which in turn results

in a higher SNR. These results serve as a visual confirmation that the estimation of the SNR defined above is a good representation of the overall quality of the fringes. As an added remark, the control matrix that was used to create this *replay* was built to be applied on-sky, where the secondary mirror of the telescope produces a central obscuration. Consequently, the control matrix can not control the actuators behind the obscuration, leading to them sometimes wandering aimlessly and introducing the great phase aberrations that are seen in the centre of the bottom-left panel.

In order to further demonstrate the relationship between the empiric SNR of the fringes and their quality, six images with widely different SNRs are selected from the *replay* sequence and presented in Figure 6.18. As can be seen in the figure, images with high SNR have sharp and well defined fringes. This is the case for the images at the top. In the opposite case, (*bottom-right*) images approaching an SNR around 1, where the fringes and the noise have about the same power, present at best blurry fringes in some regions, such as the bottom of the pupil, at worst regions without fringes, such as the top-left of the pupil, and sometimes regions with ambiguous fringes, such as the bottom left. Finally, notice that there are no inconsistent cases, where a higher SNR leads to clearly worst fringes, or vice-versa. This is not just true in the examples presented here, but also in all images that have been visually inspected.

With this tool in hand, it is now possible to acquire and analyse on-sky data. After several nights trying different stellar targets, exposure times and filters, only filters 1 (697×75 nm) and 6 (> 625 nm dichroic) allowed for sufficient light to reach the CAWS and produce visible results on which any meaningful analysis is possible. The best results, yielding the most light and the highest SNRs were obtained during the last night with the star Beta Pegasi of magnitude 0.92 in the R band. Unfortunately, the SNR remains low, just above 1, even in the best images. Figure 6.19 shows a pair of the best on-sky images and their respective SNRs. Just as with the *replay* data, these images were obtained using 2×2 binning. As can




Figure 6.18: Off-sky sequence of images from *replay* data, with decreasing SNRs, from left to right and then from top to bottom.

be seen on the figure, for SNRs close to unity, fringes remain ambiguous on raw images, i.e. it is difficult to distinguish true fringes from what could just be the result of noise. The reason for this is that at these levels, interference fringes have a power which is comparable to that of the noise.




Figure 6.19: On-sky images with the highest $\text{SNR}_f^{(\text{empiric})}$. Images were produced while observing Beta Pegasi, with filter 6 and an integration time of 400 ms.

Most of the images acquired have in fact an empirical SNR below 1. Figure 6.20 shows the average SNR of fringes for multiple exposure times and for two filters, as well as the expected SNR predicted by an unfitted model. This model works

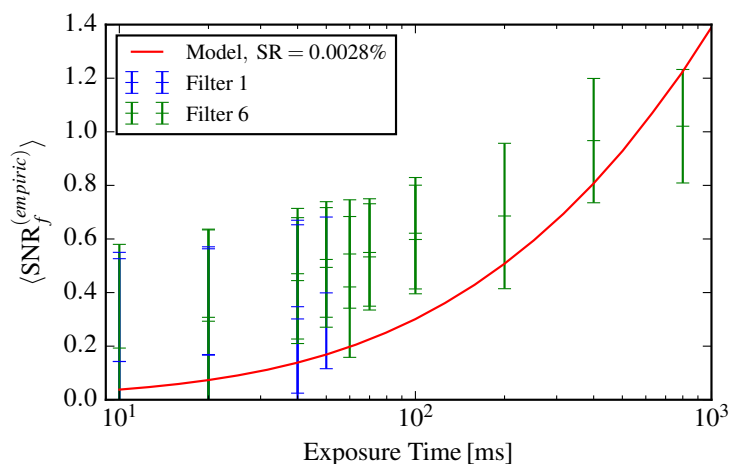


Figure 6.20: SNR of fringes for different exposure times and filters, and the expected value of an unfitted model for $\text{SR}=0.012$. The bars extend $\pm 1.5\sigma$ from the average value, spanning a total of 3σ . Results were produced on-sky while observing Beta Pegasi.

by first computing the amplitude b of the reference beam as described by N'Diaye et al. (2013), such that

$$b = b_0 P_0 \sqrt{S}, \quad (6.12)$$

where P_0 is the average amplitude of the wavefront at the CAWS' entrance pupil,

S is the SR, and b_0 is a fixed value between 0 and 1 determined by the size of the pinhole and the geometry of the pupil. According to this expression low SRs result in reference beams with low amplitudes. This is the result of less light going through the pinhole. Instead, the current design with $b_0 = 0.29$ is such that if $SR = 1$, then the fringe visibility is $\eta = 0.96$, which is almost the highest. The next step in building the model is to replace this expression of b into the simplified formula for intensity of the interferogram described in Eq. 3.20, producing the short exposure amplitude of the fringes

$$A_f = \frac{2b_0\sqrt{S}}{\pi}I_0, \quad (6.13)$$

where $I_0 = P_0^2$ and is a function of the exposure time, the chromatic bandwidth, the system's throughput and the stellar target's magnitude. This expression provides the numerator for the calculation of fringe SNR. Regarding the noise in the image, the total power of read-out and photon-shot noises are

$$\delta S_R = 22.6 \text{ e}^- \text{ RMS} \quad (6.14)$$

when using 2×2 binning, and

$$\begin{aligned} \overline{\delta S_P^2} &= b^2 + \frac{1}{\pi^2}P_0^2 \\ &= (Sb_0^2 + 1/\pi^2)I_0 \end{aligned} \quad (6.15)$$

respectively, as stated in Appendix A. Their contribution to the sideband and by extension to the denominator in the SNR calculation is

$$\delta N_{+1}^2 = \frac{4}{9N_P^2}\delta S_R^2 + \overline{\delta S_P^2}, \quad (6.16)$$

where N_P is the number of pixels across a line-pair.

From the figure, it is possible to note that SNR values obtained with filter 6, the broader of the two, are slightly higher. This is in line with the model's prediction that an increase in overall flux I_0 also leads to an increase in SNR. Note that this model does not include chromatic aberrations, which would have the opposite

action of decreasing the SNR. But given the current conditions of SR, bandwidth and flux, increasing the bandwidth has a net positive effect on the SNR of fringes. The SR parameter of the model was not fitted to the data, but was instead only adjusted manually to approximate it. The resulting average SR is in the vicinity of 0.0028%, which for a central wavelength of 730 nm translates to an average RMS aberration of 376 nm. Given that CANARY in SCAO mode can produce SRs between 0.1 and about 0.45 in the H band (Morris et al., 2014), which translate to between 228 nm and 386 nm RMS, then the variables and the output of the model have the right orders of magnitude. Regardless of this first sign of validity, the model fails to describe the rate at which the SNR increases with integration time. As can be seen in the figure, the average SNR only improves by about 5% between 400 ms and 800 ms, despite a doubling of the integration time. In contrast, between these same intervals the model predicts an increase of more than 50%: 10 times more than the data shows.

This poor fit is due to two reasons. The first one is that at very low SR the instantaneous PSF no longer has a single maximum, but is instead broken down into several speckles. The result is the intensity of the reference beam depends on whether one of these speckles will reach or miss the pinhole, making the concept of SR no longer meaningful. Since the model assumes the SR can be used to compute the intensity of the reference beam, then its validity drops hand in hand with that of this basic assumption.

The second problematic assumption this model is built on is that the integration time is short relative to the evolution rate of the turbulence. In more precise terms, the integration time must be smaller than the coherence time τ_0 , sometimes called the Greenwood time delay or the atmospheric time constant, defined as the time for which the mean square phase error is less than 1 rad^2 (Roddier, 1999). According to Fusco et al. (2004), this value is mostly found between 1 ms and 20 ms, for a wavelength of 500 nm. For 730 nm, τ_0 is between 1.5 ms and 30 ms. When exposure times grow past this delay, as is the case for those on the right half of the figure,

integrated interference fringes become the average of different wavefronts. As a result, fringes become less bright at their peaks and more bright at their valley, bringing an overall decrease in visibility. This decrease in visibility acts against the increase in SNR, causing the slower rate of growth observed in the data. Besides a decrease in visibility, long integration times also average out the aberrations, making phase measurements converge down to zero. In this case, even if fringes had any SNR left, it would be meaningless as the instrument is not perceiving aberrations.

Currently, solving the inaccuracies produced by low SRs, as well as modelling long integration times both remain unsolved problems. Producing accurate predictions of the SNR of fringes, given different atmospheric conditions, guide star magnitudes and AO specifications, is of great importance to understanding the usability and benefits of the CAWS within AO. Analytical models such as the one presented in the figure are, when valid, useful as a first step in this direction, encapsulating and transmitting some fundamental principles, but remain insufficient in some cases. For example, during the design review of complex AO systems, large Montecarlo simulations are preferred over simplified models, in order to thoroughly map the interactions between the many components in the system and making sure that variables and performances meet their requirements. Indeed, if the analytical approach to solving the aforementioned problems proves to be unattainable, then using a numerical model becomes even a larger priority.

This chapter presented the requirements, design choices and final performance of an on-sky experiment aimed at validating and testing the CAWS. The experiment was divided into two main parts. The first part consisted in characterising static aberrations, using both monochromatic and broadband polychromatic light. To do this, it was necessary to develop two independent and consistent ways to compensate for the cumulate chromatic defocus of the system. Results show that measurements with both lights sources are in close agreement. The second part of the experiment exposed the CAWS to AO residuals, both off and on-sky. The

quality of the measurements was assessed using an algorithm capable of measuring the SNR of fringes in real-time. Whereas the SNR was high off-sky, it was insufficient with on-sky wavefronts to produce meaningful demodulated phases, staying generally below 1. Finally, an analytical model that could predict the SNR was compared to the data, finding that despite producing estimations with the right order of magnitude, it failed at describing the rate of change of this value as a function of integration time. Whether or not it is possible to fix this analytically, the most thorough and precise approach to undertake in the future will be to develop a full numerical model, integrated with complete AO simulation.

Discussion

The work presented in this thesis aimed at the development and validation of the m-PDI concept and its incarnation, the CAWS. Although the concept had been proposed previously, it had not been explored to the extent reached in this manuscript. First, the development of the theory was laid out in Chapter 3, where the framework for computing the instrument's error budget was first introduced. This chapter also set the maximum polychromatic bandwidth the m-PDI could work with, showing that without chromatic aberrations, the loss of visibility due to wavefront errors is the same with monochromatic and broadband light.

Later, Chapter 4 presented the design, construction and first characterisation of the CAWS, hosted by the high-order testbed CHOUGH. Here, the transfer function of the CAWS was estimated in open-loop and calibrated against a SH, CHOUGH's HOWFS, showing that the phase demodulation algorithm first tested on simulated images also worked on real ones, that the CAWS has an approximately flat response across most of its frequency domain, with a slight attenuation on the second half, and that the dynamic range decreases with frequency. This chapter also presented a novel configuration for the CAWS: the spectrometer mode. This mode can be used to characterise the spectrum of a broadband source as seen by the instrument, which is important when using the WFS with this type of light.

In Chapter 5 a control loop was closed with the CAWS using both a monochromatic

and a broadband light source. The best results were obtained with the monochromatic source, showing a significant increase in SR on CHOUGH's imaging camera, the NFSI. Indeed, the final residual aberrations seen by the CAWS across the entire pupil and inside a central section were 55.3 nm rms and 12.5 nm rms respectively. These results show that the m-PDI can be used to tackle quasi-static aberrations, one of the three challenges in AO listed in Chapter 1. If using a DM without glued edges can allow us to obtain a 12.5 nm rms error across the entire pupil, this would leave us close to bringing the error below the 10 nm rms threshold, truly achieving nanometer-level correction and enabling the direct detection of exo-planets.

After open-loop and closed-loop characterisation of the CAWS in the lab, the instrument was taken to the WHT for testing on the CANARY system. The design process and the results were presented in Chapter 6. Here, our WFS first measured systematic biases with both monochromatic and polychromatic light, successfully compensating for chromatic aberrations and reaching a good agreement between both. These results further reinforce the idea that the CAWS is able to characterise the quasi-static aberrations that are present in an on-sky AO bench. Later, an algorithm was developed that could estimate the SNR of interference fringes. This algorithm showed that the CAWS could not measure the on-sky residual aberrations at the back of CANARY's SCAO loop, the SNR staying most of the time under 1. These results show that, unlike with quasi-static aberrations, there is still a long way to go before the m-PDI can be used to achieve XAO, the second of the three challenges mentioned before.

Exploring whether the m-PDI is a good alternative to achieving XAO remains a high priority goal. The next step in this direction is to perform an end-to-end Montecarlo simulation of the CAWS operating in closed-loop with different design parameters, such as the number of line-pairs across the pupil, the size of the pinhole, the central wavelength and the chromatic bandwidth, in order to determine the best design and its performance. For example, according to our still perfectible model, the low SNR reached during on-sky testing was most likely due to light not going

through the pinhole. This could be fixed in several ways. The first way is to move to longer wavelengths, hence ensuring better SRs at the back of the AO and with it a better coupling of the pinhole. The second way to improve the coupling is to make the pinhole larger. During the run, the pinhole diameter was $1.35\lambda/D b_0$. By almost doubling this diameter to about $2.5\lambda/D b_0$, the amplitude of the reference beam, and by extension the SNR, also doubles in these low SR conditions, while still keeping good sensitivity to low order aberrations, as shown in Chapter 3.

Simulations could also be performed to test different AO configurations. For instance, with AO systems that include a parallel low-order WFS, even larger pinholes could be tried, increasing the SNR even further. PDIs could indeed find their best utility in such configurations, with two or more hierarchically arranged WFSs, each of which measures aberrations of a particular order. In this setup, the CAWS could be used as a high-order WFS, where the size of its pinhole is determined by the spatial resolution of its low-order counterparts.

If adjusting design parameters and AO configurations is not enough, many other upgrades can be pursued to improve the performance of the CAWS in general. For example, the current throughput of the diffraction grating is only 50%. This could almost be doubled by switching to gratings with a near 100% reflection or transmission, such as reflective échelles, Volume Phase Holographic Gratings (VPHGs), or patterned liquid crystals. All of these alternatives have chromatic effects that still need to be assessed, but bring with them many advantages. In addition to having better throughput, all of these elements could be put in partially blazed configuration, removing light away from diffraction mode -1 and redirecting it towards modes 0 and +1. This in turn, increases the throughput of the focal plane filter mask, which is currently about 40%, depending on the pinhole size. Until the overall low throughput of the CAWS is improved, its relatively large chromatic bandwidth does not grant it an immediate advantage over other PDIs, regarding total flux and resulting SNR.

Finally, besides exploring whether the m-PDI is capable of enabling XAO, the other

important step is to test its capability to cophase segmented mirrors in extremely large telescopes, the third challenge mentioned at the beginning of this thesis. As discussed in previous chapters, other PDIs have shown promising results in this direction, either in simulation or in small telescopes. Preliminary tests, not included in the main body of this thesis, have already been performed in simulation with an arbitrary ELT-like pupil, producing encouraging results which can be found in Appendix C, Figure C.1. If further experiments confirm both this capability and that the m-PDI is capable of achieving XAO, a future version of the CAWS would be a good candidate to simultaneously perform cophasing and fast high-order wavefront sensing at the heart of an advanced AO system on a large segmented telescope.

Analytical derivation for the propagation of photon-shot and read-out noise

As presented in Section 3.3, the intensity in the exit pupil can be described as

$$I_C = b^2 + \frac{1}{\pi^2}P^2 + \frac{2b}{\pi}P \cos\left(\frac{2\pi}{T}x - \phi_{LP}\right), \quad (\text{A1})$$

where ϕ_{LP} is ϕ low-pass filtered by M_{+1} . For a small ϕ_{LP} the expression can be approximated as

$$I_C = b^2 + \frac{1}{\pi^2}P^2 + \frac{2b}{\pi}P \cos\left(\frac{2\pi}{T}x\right) + \frac{2b}{\pi}P \sin\left(\frac{2\pi}{T}x\right) \phi_{LP}. \quad (\text{A2})$$

Back in the Fourier domain and considering $P \simeq P_0$, the demodulated phase is

$$\widehat{\phi_{LP}} = \frac{\pi}{2bP_0} \left(\widehat{I}_C \otimes \left[-i\delta\left(k - \frac{\lambda_0 f}{T}\right) + i\delta\left(k + \frac{\lambda_0 f}{T}\right) \right] \right) M_{+1}\left(k + \frac{\lambda_0 f}{T}\right), \quad (\text{A3})$$

where the mask M_{+1} has been centered around the origin.

When white noise n_w of mean value $\bar{n}_w = 0$ and power δS_w^2 is added to the detected signal I_C and put through the demodulation the result is

$$\frac{\pi}{2bP_0} i \left(-\widehat{n}_w\left(k - \frac{\lambda_0 f}{T}\right) + \widehat{n}_w\left(k + \frac{\lambda_0 f}{T}\right) \right) M_{+1}\left(k + \frac{\lambda_0 f}{T}\right), \quad (\text{A4})$$

where $\hat{n}_w\left(k - \frac{\lambda_0 f}{T}\right)$ and $\hat{n}_w\left(k + \frac{\lambda_0 f}{T}\right)$ are uncorrelated. This means the sum of both results in a white noise $n_{w'}$ where $\delta S_{w'}^2 = 2\delta S_w^2$. The resulting noise is then filtered by M_{+1} . Using Parseval's theorem the power of the filtered noise can be computed as

$$\delta S_{w'}^2 = \left\langle \int_{-\lambda_0 f/2\mu_p}^{\lambda_0 f/2\mu_p} \int_{-\lambda_0 f/2\mu_p}^{\lambda_0 f/2\mu_p} |\hat{n}_{w'}|^2 dk d\zeta \right\rangle = \int_{-\lambda_0 f/2\mu_p}^{\lambda_0 f/2\mu_p} \int_{-\lambda_0 f/2\mu_p}^{\lambda_0 f/2\mu_p} \langle |\hat{n}_{w'}|^2 \rangle dk d\zeta, \quad (\text{A5})$$

For a square filter

$$M_{+1}\left(k + \frac{\lambda_0 f}{T}\right) = \text{rect}\left(\frac{3Tk}{2\lambda_0 f}\right), \quad (\text{A6})$$

the power of the filtered noise is

$$\begin{aligned} \delta S_{w',M_{+1}}^2 &= \int_{-\lambda_0 f/2\mu_p}^{\lambda_0 f/2\mu_p} \int_{-\lambda_0 f/2\mu_p}^{\lambda_0 f/2\mu_p} \langle |\hat{n}_{w'}|^2 \rangle \text{rect}\left(\frac{3Tk}{2\lambda_0 f}\right) dk d\zeta \\ &= \int_{-\lambda_0 f/3T}^{\lambda_0 f/3T} \int_{-\lambda_0 f/3T}^{\lambda_0 f/3T} \langle |\hat{n}_{w'}|^2 \rangle dk d\zeta = \left(\frac{2\mu_p}{3T}\right)^2 \delta S_{w'}^2. \end{aligned} \quad (\text{A7})$$

Let the number of pixels sampling one line-pair of period T be $N_P = T/\mu_p$, then the noise's variance induced by white noise is

$$\sigma_w^2 = \left(\frac{\pi}{2bP_0}\right)^2 \delta S_{w',M_{+1}}^2 = \left(\frac{\pi}{2bP_0}\right)^2 \frac{4}{9N_P^2} \delta S_{w'}^2 = \frac{2\pi^2}{9N_P^2 b^2 P_0^2} \delta S_w^2. \quad (\text{A8})$$

This equation can be directly applied to read-out noise n_R . With respect to photon-shot noise n_P it is important first to notice the noise between two pixels is uncorrelated. Formally, this is

$$\text{Cov}(n_P(x), n_P(x + \xi)) = \begin{cases} I_C(x), & \xi = 0 \\ 0, & \forall \xi \neq 0 \end{cases} \quad (\text{A9})$$

where $\text{Cov}()$ is the covariance between to functions and ξ is the distance between any pair of pixels of coordinates x and $x + \xi$. For such a correlation the power spectral density is flat and equal to the average power of the noise

$$\overline{\delta S_P^2} = \frac{1}{N_G T} \int_0^{N_G T} I_C dx, \quad (\text{A10})$$

where N_G is the number of line-pairs in the pupil. If the average phase $\bar{\phi} = 0$ across the pupil, then the above can be approximated to

$$\overline{\delta S_P^2} = b^2 + \frac{1}{\pi^2} P_0^2. \quad (\text{A11})$$

Then, if the intensity at the entrance pupil is $I_0 = P_0^2$, the error contribution of the read-out noise and the photon-shot noise is

$$\sigma_R^2 + \sigma_P^2 = \frac{2\pi^2}{9N_P^2 b^2 I_0} \delta S_R^2 + \frac{2\pi^2}{9N_P^2 b^2 I_0} \left(b^2 + \frac{1}{\pi^2} I_0 \right). \quad (\text{A12})$$

Control matrix for the CAWS in CHOUGH

Figure B.1 presents the transpose of poking sequence A . Figure B.2 presents the control matrix for CAWS and CHOUGH. As can be seen in the figure, non-null values are distributed along a sigmoid from one corner to another, rather than following a straight diagonal. This happens when the number of phase points across the pupil differs from the number of actuators. This is not a problem, as long as the WFS has enough spatial resolution to sample all actuators, as is the case here by design.




Figure B.1: Transpose of A . This is the shortest possible A for which $M = 2L_{active}$. Green columns show inactive actuators.




Figure B.2: Control matrix mapping the relationship between CAWS' phase points and DM actuators.

Cophasing of segmented apertures

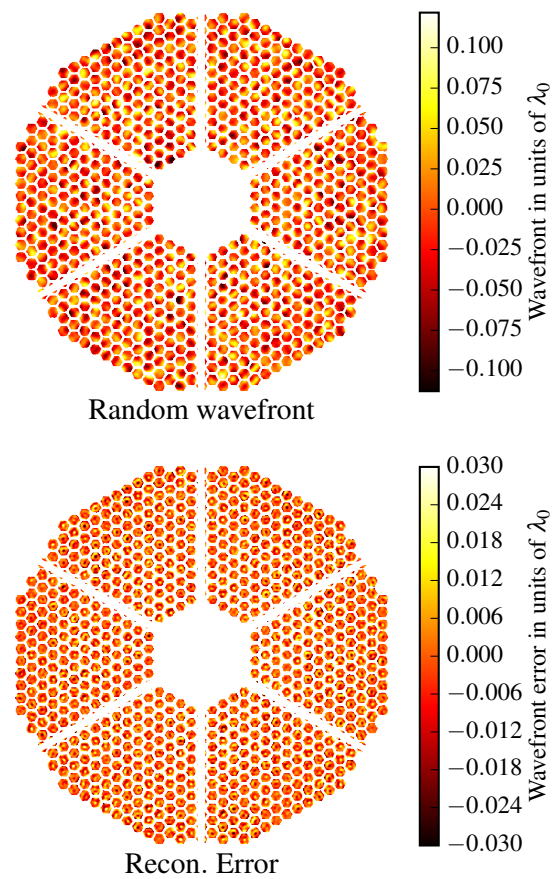


Figure C.1: Preliminary simulation of (*top*) random phase aberrations through an arbitrary ELT-like pupil and (*bottom*) the error of the wavefront reconstructed by the CAWS, in units of λ_0 .

Bibliography

- S. aurantiaca. Mach-Zehnder interferometer. URL https://commons.wikimedia.org/wiki/File:Mach_Zehnder_interferometer.svg.
- E. Azoulay, V. Thiermann, A. Jetter, A. Kohnle, and Z. Azar. Optical measurement of the inner scale of turbulence. *Journal of Physics D: Applied Physics*, 21(10S): S41, 1988.
- A. Basden and R. Myers. The Durham adaptive optics real-time controller: capability and extremely large telescope suitability. *Monthly Notices of the Royal Astronomical Society*, 424(2):1483–1494, 2012.
- J. M. Beckers. Increasing the size of the isoplanatic patch with multiconjugate adaptive optics. In *European Southern Observatory Conference and Workshop Proceedings*, volume 30, page 693, 1988.
- J. M. Beckers. Adaptive optics for astronomy: principles, performance, and applications. *Annual review of astronomy and astrophysics*, 31(1):13–62, 1993.
- J.-L. Beuzit, M. Feldt, K. Dohlen, D. Mouillet, P. Puget, F. Wildi, L. Abe, J. Antichi, A. Baruffolo, P. Baudoz, et al. Sphere: a planet finder instrument for the vlt. In *Ground-based and airborne instrumentation for astronomy II*, volume 7014, page 701418. International Society for Optics and Photonics, 2008.

- N. Bharmal. *Adaptive Optics for Stellar Interferometry*. PhD thesis, University of Cambridge, 2005.
- N. A. Bharmal, R. M. Myers, A. G. Basden, and A. P. Reeves. An interferometric wavefront sensor for high-sensitivity low-amplitude measurements. In *Adaptive Optics Systems III*, volume 8447, page 84472J. International Society for Optics and Photonics, 2012.
- N. A. Bharmal, R. M. Myers, A. G. Basden, D. Hölck, and T. J. Morris. Chough, the canary hosted-upgrade for high-order adaptive optics. In *Adaptive Optics Systems IV*, volume 9148, page 91485V. International Society for Optics and Photonics, 2014.
- N. A. Bharmal, R. M. Myers, D. Hölck, N. S. Dubost, C. M. Dubbeldam, A. G. Basden, M. Black, and D. M. Henry. Chough: current status and future plans. In *Adaptive Optics Systems VI*, volume 10703, page 1070360. International Society for Optics and Photonics, 2018.
- M. Bottom, B. Femenia, E. Huby, D. Mawet, R. Dekany, J. Milburn, and E. Serabyn. Speckle nulling wavefront control for palomar and keck. In *Adaptive Optics Systems V*, volume 9909, page 990955. International Society for Optics and Photonics, 2016.
- CANARY team. OPTICON WP10 AO Networking. <https://sites.google.com/view/opticon-ao/canary-tests/configurations>, 2019.
- G. Chanan, M. Troy, F. Dekens, S. Michaels, J. Nelson, T. Mast, and D. Kirkman. Phasing the mirror segments of the keck telescopes: the broadband phasing algorithm. *Applied Optics*, 37(1):140–155, 1998.
- S. Chueca, M. Reyes, A. Schumacher, and L. Montoya. Dipsi: measure of the tip-tilt with a diffraction image phase sensing instrument. In *Ground-based and Airborne Telescopes II*, volume 7012, page 701213. International Society for Optics and Photonics, 2008.

- T. Currie, O. Guyon, M. Tamura, T. Kudo, N. Jovanovic, J. Lozi, J. E. Schlieder, T. D. Brandt, J. Kuhn, E. Serabyn, et al. Subaru/scexao first-light direct imaging of a young debris disk around hd 36546. *The Astrophysical Journal Letters*, 836(1):L15, 2017.
- R. Dekany, J. Roberts, R. Burruss, A. Bouchez, T. Truong, C. Baranec, S. Guiwits, D. Hale, J. Angione, T. Trinh, et al. Palm-3000: exoplanet adaptive optics for the 5 m hale telescope. *The Astrophysical Journal*, 776(2):130, 2013.
- R. G. Dekany, M. C. Britton, D. T. Gavel, B. L. Ellerbroek, G. Herriot, C. E. Max, and J.-P. Veran. Adaptive optics requirements definition for tmt. In *Advancements in Adaptive Optics*, volume 5490, pages 879–891. International Society for Optics and Photonics, 2004.
- E. Diolaiti, P. Ciliegi, R. Abicca, G. Agapito, C. Arcidiacono, A. Baruffolo, M. Bellazzini, V. Biliotti, M. Bonaglia, G. Bregoli, et al. On the road to the preliminary design review of the maory adaptive optics module for e-elt. In *this conference*, 2017.
- K. Dohlen, M. Langlois, P. Lanzoni, S. Mazzanti, A. Vigan, L. Montoya, E. Hernandez, M. Reyes, I. Surdej, and N. Yaitskova. Zeus: a cophasing sensor based on the zernike phase contrast method. In *Ground-based and Airborne Telescopes*, volume 6267, page 626734. International Society for Optics and Photonics, 2006.
- K. Dohlen, F. Wildi, J. Beuzit, P. Puget, D. Mouillet, A. Baruffolo, A. Boccaletti, J. Charton, R. Claudi, A. Costille, et al. Sphere system analysis predictions. *Adaptive Optics for Extremely Large Telescopes (AO4ELT)*, 2011, 2011.
- Y.-z. Du, G.-y. Feng, H.-r. Li, Z. Cai, H. Zhao, and S.-h. Zhou. Real-time determination of beam propagation factor by mach–zehnder point diffraction interferometer. *Optics Communications*, 287:1–5, 2013.
- S. Esposito, A. Riccardi, L. Fini, A. T. Puglisi, E. Pinna, M. Xompero, R. Briguglio, F. Quirós-Pacheco, P. Stefanini, J. C. Guerra, et al. First light ao (flao) system

- for lbt: final integration, acceptance test in europe, and preliminary on-sky commissioning results. In *Adaptive Optics Systems II*, volume 7736, page 773609. International Society for Optics and Photonics, 2010.
- S. Esposito, A. Riccardi, E. Pinna, A. Puglisi, F. Quirós-Pacheco, C. Arcidiacono, M. Xompero, R. Briguglio, G. Agapito, L. Busoni, et al. Large binocular telescope adaptive optics system: new achievements and perspectives in adaptive optics. In *Astronomical Adaptive Optics Systems and Applications IV*, volume 8149, page 814902. International Society for Optics and Photonics, 2011.
- J. W. Evans, K. Morzinski, S. Severson, L. Poyneer, B. Macintosh, D. Dillon, L. Reza, D. Gavel, D. Palmer, S. Olivier, et al. Extreme adaptive optics testbed: performance and characterization of a 1024-MEMS deformable mirror. In *MEMS/MOEMS Components and their applications III*, volume 6113, page 61130I. International Society for Optics and Photonics, 2006.
- M. Feldman, D. J. Mockler, R. E. English Jr, J. L. Byrd, and J. T. Salmon. Self-referencing mach-zehnder interferometer as a laser system diagnostic. In *Active and Adaptive Optical Systems*, volume 1542, pages 490–501. International Society for Optics and Photonics, 1991.
- J. J. Fortney, K. Lodders, M. S. Marley, and R. S. Freedman. A unified theory for the atmospheres of the hot and very hot jupiters: two classes of irradiated atmospheres. *The Astrophysical Journal*, 678(2):1419, 2008.
- T. Fusco, N. Ageorges, G. Rousset, D. Rabaud, E. Gendron, D. Mouillet, F. Lacombe, G. Zins, J. Charton, C. Lidman, et al. NAOS performance characterization and turbulence parameters estimation using closed-loop data. In *Advancements in Adaptive Optics*, volume 5490, pages 118–129. International Society for Optics and Photonics, 2004.
- T. Fusco, G. Rousset, J.-F. Sauvage, C. Petit, J.-L. Beuzit, K. Dohlen, D. Mouillet, J. Charton, M. Nicolle, M. Kasper, et al. High-order adaptive optics requirements

- for direct detection of extrasolar planets: Application to the sphere instrument. *Optics Express*, 14(17):7515–7534, 2006.
- T. Fusco, J.-F. Sauvage, C. Petit, A. Costille, K. Dohlen, D. Mouillet, J.-L. Beuzit, M. Kasper, M. Suarez, C. Soenke, et al. Final performance and lesson-learned of saxo, the vlt-sphere extreme ao: from early design to on-sky results. In *Adaptive Optics Systems IV*, volume 9148, page 91481U. International Society for Optics and Photonics, 2014.
- E. Gendron and P. Léna. Astronomical adaptive optics. 1: Modal control optimization. *Astronomy and Astrophysics*, 291:337–347, 1994.
- G. F. Gilmore. Opticon: a (small) part of european astronomy. In *Second Backaskog Workshop on Extremely Large Telescopes*, volume 5382, pages 138–142. International Society for Optics and Photonics, 2004.
- J. Glückstad and P. C. Mogensen. Optimal phase contrast in common-path interferometry. *Applied optics*, 40(2):268–282, 2001.
- R. A. Gonsalves. Phase retrieval and diversity in adaptive optics. *Optical Engineering*, 21(5):215829, 1982.
- F. Gonté, C. Araujo, R. Bourtembourg, R. Brast, F. Derie, P. Duhoux, C. Dupuy, C. Frank, R. Karban, R. Mazzoleni, et al. On-sky testing of the active phasing experiment. *The Messenger*, 136:25–31, 2009.
- A. Guesalaga, B. Neichel, A. Cortés, C. Béchet, and D. Guzmán. Using the wind profiler method with wide-field laser-guide-stars adaptive optics to quantify the frozen-flow decay. *Monthly Notices of the Royal Astronomical Society*, 440(3):1925–1933, 2014.
- O. Guyon. Limits of adaptive optics for high-contrast imaging. *The Astrophysical Journal*, 629(1):592, 2005.

- O. Guyon, F. Martinache, V. Garrel, F. Vogt, K. Yokochi, and T. Yoshikawa. The subaru coronagraphic extreme ao (scexao) system: wavefront control and detection of exoplanets with coherent light modulation in the focal plane. In *Adaptive optics systems II*, volume 7736, page 773624. International Society for Optics and Photonics, 2010.
- O. Guyon, F. Martinache, C. Clergeon, R. Russell, T. Groff, and V. Garrel. Wavefront control with the subaru coronagraphic extreme adaptive optics (scexao) system. In *Astronomical Adaptive Optics Systems and Applications IV*, volume 8149, page 814908. International Society for Optics and Photonics, 2011.
- F. Hammer, M. Puech, F. F. Assemat, E. Gendron, F. Sayede, P. Laporte, M. Marteaud, A. Liotard, and F. Zamkotsian. Falcon: a concept to extend adaptive optics corrections to cosmological fields. In *Second Backaskog Workshop on Extremely Large Telescopes*, volume 5382, pages 727–737. International Society for Optics and Photonics, 2004.
- J. W. Hardy. *Adaptive optics for astronomical telescopes*, volume 16. Oxford University Press on Demand, 1998.
- M. Hartung, B. Macintosh, P. Langlois, N. Sadakuni, D. Gavel, J. K. Wallace, D. Palmer, L. Poyneer, D. Savransky, S. Thomas, et al. On-sky low order non-common path correction of the gpi calibration unit. In *Adaptive Optics Systems IV*, volume 9148, page 91485Q. International Society for Optics and Photonics, 2014.
- M. A. Herráez, D. R. Burton, M. J. Lalor, and M. A. Gdeisat. Fast two-dimensional phase-unwrapping algorithm based on sorting by reliability following a noncontinuous path. *Applied optics*, 41(35):7437–7444, 2002.
- J. Herrmann. Least-squares wave front errors of minimum norm. *JOSA*, 70(1): 28–35, 1980.

- E. Hugot, M. Ferrari, K. El Hadi, A. Costille, K. Dohlen, P. Rabou, P. Puget, and J. Beuzit. Active optics methods for exoplanet direct imaging-stress polishing of supersmooth aspherics for vlt-sphere planet finder. *Astronomy & Astrophysics*, 538:A139, 2012.
- J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in science & engineering*, 9(3):90–95, 2007.
- Isaac Newton Group of Telescopes. Staralt. <http://www.ing.iac.es/astrometry/index.html>, 2019.
- J. R. Janesick. *Scientific charge-coupled devices*, volume 83. SPIE press, 2001.
- P. Janin-Potiron, M. N’Diaye, P. Martinez, A. Vigan, K. Dohlen, and M. Carillet. Fine cophasing of segmented aperture telescopes with zelda, a zernike wavefront sensor in the diffraction-limited regime. *Astronomy & Astrophysics*, 603:A23, 2017.
- N. Jovanovic, O. Guyon, J. Lozi, T. Currie, J. Hagelberg, B. Norris, G. Singh, P. Pathak, D. Doughty, S. Goebel, et al. The scexao high contrast imager: transitioning from commissioning to science. In *Adaptive Optics Systems V*, volume 9909, page 99090W. International Society for Optics and Photonics, 2016.
- A. Kellerer, F. Vidal, E. Gendron, Z. Hubert, D. Perret, and G. Rousset. Deformable mirrors for open-loop adaptive optics. In *Adaptive Optics Systems III*, volume 8447, page 844765. International Society for Optics and Photonics, 2012.
- M. Lamb, C. Correia, J.-F. Sauvage, J.-P. Véran, D. Andersen, A. Vigan, P. Wizinowich, M. van Dam, L. Mugnier, and C. Bond. Estimating phase errors from pupil discontinuities from simulated on sky data: examples with vlt and keck. In *Adaptive Optics Systems V*, volume 9909, page 99096D. International Society for Optics and Photonics, 2016.
- M. P. Lamb, C. Correia, J.-F. Sauvage, J.-P. Véran, D. R. Andersen, A. Vigan, P. L. Wizinowich, M. A. Van Dam, L. Mugnier, and C. Bond. Quantifying telescope

- phase discontinuities external to adaptive optics systems by use of phase diversity and focal plane sharpening. *Journal of Astronomical Telescopes, Instruments, and Systems*, 3(3):039001, 2017.
- Laser Micromachining Ltd. LML. <http://www.lasermicromachining.com/>, 2017.
- M. Loupiau, M. Langlois, E. Thiébaud, M. Tallon, and J. Leger. Status of an extreme adaptive optics testbench using a self-referenced mach-zehnder wavefront sensor. In *Adaptive Optics Systems V*, volume 9909, page 99094F. International Society for Optics and Photonics, 2016.
- L. Mach. Ueber einen interferenzrefraktor. *Zeitschrift für Instrumentenkunde*, 12(89):806, 1892.
- B. Macintosh, J. R. Graham, P. Ingraham, Q. Konopacky, C. Marois, M. Perrin, L. Poyneer, B. Bauman, T. Barman, A. S. Burrows, et al. First light of the gemini planet imager. *Proceedings of the National Academy of Sciences*, 111(35):12661–12666, 2014.
- B. A. Macintosh, J. R. Graham, D. W. Palmer, R. Doyon, J. Dunn, D. T. Gavel, J. Larkin, B. Oppenheimer, L. Saddlemyer, A. Sivaramakrishnan, et al. The gemini planet imager: from science to design to construction. In *Adaptive Optics Systems*, volume 7015, page 701518. International Society for Optics and Photonics, 2008.
- V. N. Mahajan. Strehl ratio for primary aberrations: some analytical results for circular and annular pupils. *JOSA*, 72(9):1258–1266, 1982.
- V. N. Mahajan. Strehl ratio for primary aberrations in terms of their aberration variance. *JOSA*, 73(6):860–861, 1983.
- F. Malbet, J. Yu, and M. Shao. High-dynamic-range imaging using a deformable mirror for space coronagraphy. *Publications of the Astronomical Society of the Pacific*, 107(710):386, 1995.

- C. Marois, R. Doyon, R. Racine, and D. Nadeau. Efficient speckle noise attenuation in faint companion imaging. *Publications of the Astronomical Society of the Pacific*, 112(767):91, 2000.
- C. Marois, D. Lafreniere, R. Doyon, B. Macintosh, and D. Nadeau. Angular differential imaging: a powerful high-contrast imaging technique. *The Astrophysical Journal*, 641(1):556, 2006.
- F. Martinache, O. Guyon, N. Jovanovic, C. Clergeon, G. Singh, T. Kudo, T. Currie, C. Thalmann, M. McElwain, and M. Tamura. On-sky speckle nulling demonstration at small angular separation with scexao. *Publications of the Astronomical Society of the Pacific*, 126(940):565, 2014.
- R. Mazzoleni, N. Yaitskova, and L. Noethe. Disentangling between low order telescope aberrations and segmentation errors using a shack-hartmann sensor. In *Ground-based and Airborne Telescopes II*, volume 7012, page 701216. International Society for Optics and Photonics, 2008.
- H. Medecky, E. Tejnly, K. Goldberg, and J. Bokor. Phase-shifting point diffraction interferometer. *Optics letters*, 21(19):1526–1528, 1996.
- T. Morris, E. Gendron, A. Basden, O. Martin, J. Osborn, D. Henry, Z. Hubert, G. Sivo, D. Gratadour, F. Chemla, et al. Canary phase b: on-sky open-loop tomographic lgs ao results. In *Adaptive Optics Systems IV*, volume 9148, page 91481I. International Society for Optics and Photonics, 2014.
- T. Morris, A. Basden, A. Calcines-Rosario, K. Dohlen, C. Dubbeldam, K. El Hadi, E. Fitzsimons, T. Fusco, É. Gendron, F. Hammer, et al. Phase a ao system design and performance for mosaic at the elt. In *Adaptive Optics Systems VI*, volume 10703, page 1070316. International Society for Optics and Photonics, 2018.
- R. M. Myers, Z. Hubert, T. J. Morris, E. Gendron, N. A. Dipper, A. Kellerer, S. J. Goodsell, G. Rousset, E. Younger, M. Marteau, et al. Canary: the on-sky

- ngs/lgs moao demonstrator for eagle. In *Adaptive Optics Systems*, volume 7015, page 70150E. International Society for Optics and Photonics, 2008.
- M. N'Diaye, K. Dohlen, T. Fusco, and B. Paul. Calibration of quasi-static aberrations in exoplanet direct-imaging instruments with a zernike phase-mask sensor. *Astronomy & Astrophysics*, 555:A94, 2013.
- M. N'Diaye, A. Vigan, K. Dohlen, J.-F. Sauvage, A. Caillat, A. Costille, J. Girard, J.-L. Beuzit, T. Fusco, P. Blanchard, et al. Calibration of quasi-static aberrations in exoplanet direct-imaging instruments with a zernike phase-mask sensor-ii. concept validation with zelda on vlt/sphere. *Astronomy & Astrophysics*, 592:A79, 2016.
- M. N'Diaye, F. Martinache, N. Jovanovic, J. Lozi, O. Guyon, B. Norris, A. Ceau, and D. Mary. Calibration of the island effect: Experimental validation of closed-loop focal plane wavefront control on subaru/scexao. *Astronomy & Astrophysics*, 610:A18, 2018.
- R. J. Noll. Zernike polynomials and atmospheric turbulence. *JOsA*, 66(3):207–211, 1976.
- J. Notaras and C. Paterson. Demonstration of closed-loop adaptive optics with a point-diffraction interferometer in strong scintillation with optical vortices. *Optics Express*, 15(21):13745–13756, 2007.
- T. E. Oliphant. *A guide to NumPy*, volume 1. Trelgol Publishing USA, 2006.
- B. Paul, L. Mugnier, J.-F. Sauvage, K. Dohlen, and M. Ferrari. High-order myopic coronagraphic phase diversity (coffee) for wave-front control in high-contrast imaging systems. *Optics Express*, 21(26):31751–31768, 2013.
- B. Paul, J.-F. Sauvage, L. Mugnier, K. Dohlen, C. Petit, T. Fusco, D. Mouillet, J.-L. Beuzit, and M. Ferrari. Compensation of high-order quasi-static aberrations on sphere with the coronagraphic phase diversity (coffee). *Astronomy & Astrophysics*, 572:A32, 2014.

- E. Pinna, F. Quirós-Pacheco, S. Esposito, A. Puglisi, and P. Stefanini. The pyramid phasing sensor (pyps). In *Ground-based and Airborne Telescopes II*, volume 7012, page 70123D. International Society for Optics and Photonics, 2008.
- R. Racine, G. A. Walker, D. Nadeau, R. Doyon, and C. Marois. Speckle noise and the detection of faint companions. *Publications of the Astronomical Society of the Pacific*, 111(759):587, 1999.
- R. Ragazzoni. Pupil plane wavefront sensing with an oscillating prism. *Journal of modern optics*, 43(2):289–293, 1996.
- R. Ragazzoni and J. Farinato. Sensitivity of a pyramidal wave front sensor in closed loop adaptive optics. *Astronomy and Astrophysics*, 350:L23–L26, 1999.
- C. Rao, W. Jiang, and N. Ling. Spatial and temporal characterization of phase fluctuations in non-kolmogorov atmospheric turbulence. *Journal of Modern Optics*, 47(6):1111–1126, 2000.
- F. Rigaut, G. Rousset, P. Kern, J. Fontanella, J. Gaffard, F. Merkle, and P. Léna. Adaptive optics on a 3.6-m telescope-results and performance. *Astronomy and Astrophysics*, 250:280–290, 1991.
- F. Rigaut, E. Vernet, R. Ragazzoni, S. Esposito, and N. Hubin. Beyond conventional adaptive optics. In *ESO Topical Meeting (Ed. E. Vernet, R. Ragazzoni, N. Hubin, S. Esposito, Garching: ESO Conf. and Workshop Proc. 58, 2002)*, page 11, 2002.
- L. C. Roberts Jr, M. D. Perrin, F. Marchis, A. Sivaramakrishnan, R. B. Makidon, J. C. Christou, B. A. Macintosh, L. A. Poyneer, M. A. van Dam, and M. Troy. Is that really your strehl ratio? In *Advancements in Adaptive Optics*, volume 5490, pages 504–515. International Society for Optics and Photonics, 2004.
- F. Roddier. The effects of atmospheric turbulence in optical astronomy. In: *Progress in optics. Volume 19. Amsterdam, North-Holland Publishing Co., 1981, p. 281-376.*, 19:281–376, 1981.

- F. Roddier. Curvature sensing and compensation: a new concept in adaptive optics. *Applied Optics*, 27(7):1223–1225, 1988.
- F. Roddier. Adaptive optics: Performance and limitations. In *Symposium-International Astronomical Union*, volume 158, pages 273–281. Cambridge University Press, 1994.
- F. Roddier. *Adaptive optics in astronomy*. Cambridge university press, 1999.
- J.-F. Sauvage, T. Fusco, C. Petit, S. Meimon, E. Fedrigo, M. S. Valles, M. Kasper, N. Hubin, J.-L. Beuzit, J. Charton, et al. Saxo, the extreme adaptive optics system of sphere: overview and calibration procedure. In *Adaptive Optics Systems II*, volume 7736, page 77360F. International Society for Optics and Photonics, 2010.
- J.-F. Sauvage, L. Mugnier, B. Paul, and R. Villecroze. Coronagraphic phase diversity: a simple focal plane sensor for high-contrast imaging. *Optics Letters*, 37(23):4808–4810, 2012.
- J.-F. Sauvage, T. Fusco, M. Lamb, J. Girard, M. Brinkmann, A. Guesalaga, P. Wizinowich, J. O’Neal, M. N’Diaye, A. Vigan, et al. Tackling down the low wind effect on sphere instrument. In *Adaptive Optics Systems V*, volume 9909, page 990916. International Society for Optics and Photonics, 2016.
- N. Schwartz, J.-F. Sauvage, C. Correia, C. Petit, F. Quiros-Pacheco, T. Fusco, K. Dohlen, K. E. Hadi, N. Thatte, F. Clarke, et al. Sensing and control of segmented mirrors with a pyramid wavefront sensor in the presence of spiders. *arXiv preprint arXiv:1809.08839*, 2018.
- Semrock. 596 nm edge BrightLine single-edge standard epi-fluorescence dichroic beamsplitter. <https://www.semrock.com/FilterDetails.aspx?id=FF596-Di01-25x36>, 2020.
- G. Sivo, C. Kulcsár, J.-M. Conan, H.-F. Raynaud, É. Gendron, A. Basden, F. Vidal, T. Morris, S. Meimon, C. Petit, et al. First on-sky scao validation of full lqg

- control with vibration mitigation on the canary pathfinder. *Optics express*, 22(19):23565–23591, 2014.
- A. J. Skemer, P. M. Hinz, S. Esposito, A. Burrows, J. Leisenring, M. Skrutskie, S. Desidera, D. Mesa, C. Arcidiacono, F. Mannucci, et al. First light lbt ao images of hr 8799 bcde at 1.6 and 3.3 μm : New discrepancies between young planets and old brown dwarfs. *The Astrophysical Journal*, 753(1):14, 2012.
- R. Smartt and W. Steel. Theory and application of point-diffraction interferometers. *Japanese Journal of Applied Physics*, 14(S1):351, 1975.
- I. Surdej, B. Luong, A. Vigan, C. Araujo, R. Bourtembourg, R. Brast, P. Duhoux, C. Dupuy, C. Frank, F. Gonte, et al. Preliminary results obtained with the zeus phasing sensor within the ape experiment. In *Ground-based and Airborne Telescopes II*, volume 7012, page 701218. International Society for Optics and Photonics, 2008.
- I. Surdej, N. Yaitskova, and F. Gonte. On-sky performance of the zernike phase contrast sensor for the phasing of segmented telescopes. *Applied optics*, 49(21):4052–4062, 2010.
- M. Takeda, H. Ina, and S. Kobayashi. Fourier-transform method of fringe-pattern analysis for computer-based topography and interferometry. *JosA*, 72(1):156–160, 1982.
- Thorlabs. Protected Aluminium Mirrors. https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=264, 2019.
- A. Tokovinin. Seeing improvement with ground-layer adaptive optics. *Publications of the Astronomical Society of the Pacific*, 116(824):941, 2004.
- A. Vigan, K. Dohlen, and S. Mazzanti. On-sky multiwavelength phasing of segmented telescopes with the zernike phase contrast sensor. *Applied optics*, 50(17):2708–2718, 2011.

- J. K. Wallace, R. S. Burruss, R. D. Bartos, T. Q. Trinh, L. A. Pueyo, S. F. Fregoso, J. R. Angione, and J. C. Shelton. The gemini planet imager calibration wavefront sensor instrument. In *Adaptive Optics Systems II*, volume 7736, page 77365D. International Society for Optics and Photonics, 2010.
- J. K. Wallace, S. Rao, R. M. Jensen-Clem, and G. Serabyn. Phase-shifting zernike interferometer wavefront sensor. In *Optical Manufacturing and Testing IX*, volume 8126, page 81260F. International Society for Optics and Photonics, 2011.
- N. Yaitskova, K. Dohlen, P. Dierickx, and L. Montoya. Mach–zehnder interferometer for piston and tip–tilt sensing in segmented telescopes: theory and analytical treatment. *JOSA A*, 22(6):1093–1105, 2005.
- L. Zehnder. Ein neuer interferenzrefraktor, 1891.
- F. Zernike. Diffraction theory of the knife-edge test and its improved form, the phase-contrast method. *Monthly Notices of the Royal Astronomical Society*, 94: 377–384, 1934.
- A. Zurlo, A. Vigan, D. Mesa, R. Gratton, C. Moutou, M. Langlois, R. Claudi, L. Pueyo, A. Boccaletti, A. Baruffolo, et al. Performance of the vlt planet finder sphere-i. photometry and astrometry precision with irdis and ifs in laboratory. *Astronomy & Astrophysics*, 572:A85, 2014.

Colophon

This thesis was typeset with L^AT_EX 2_ε. It was created using the *memoir* package, maintained by Lars Madsen, with the *madsen* chapter style. The font used is Latin Modern, derived from fonts designed by Donald E. Kunith.