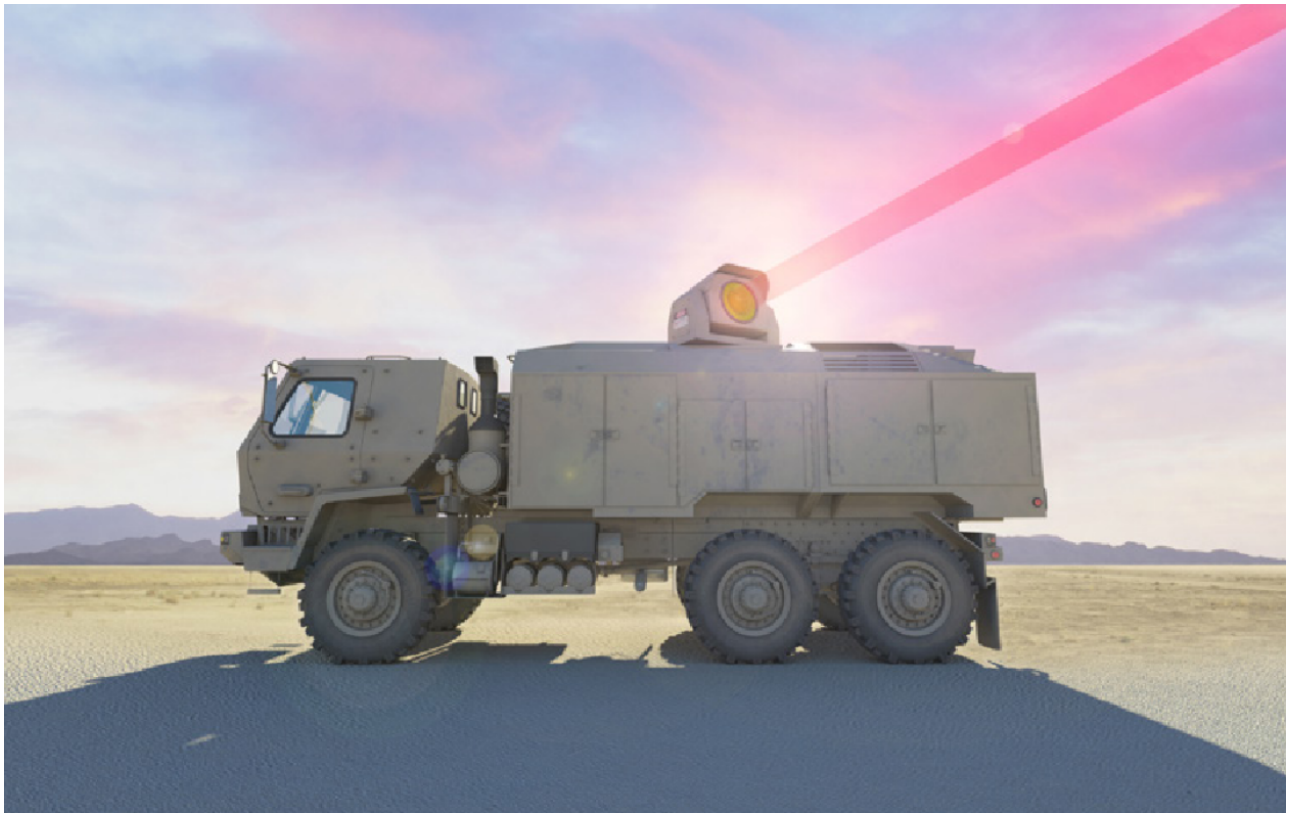

Adaptive Optics-Based Beam Projection for Target-in-the-Loop Scenarios under Atmospheric Turbulence

Adrián Bradosche Pallares
Student Number: 5016339
Date: 06/05/2025

TU Delft

TU Delft Faculty of Aerospace Engineering

Adaptive Optics-Based Beam Projection for Target-in-the-Loop Scenarios under Atmospheric Turbulence

by

Adrián Bradosche Pallares

to obtain the degree of Master of Science
at the Delft University of Technology
to be defended publicly on May 6, 2025 at 13:00

Thesis committee:	
Chair:	Prof.dr. L.L.A. Bert Vermeersen
Supervisors:	Dr.ir. Rudolf Saathof
	Sven van Binsbergen
External examiner:	Dr.ir. M.M. Rene van Paassen
Place:	Faculty of Aerospace Engineering, Delft
Project Duration:	June, 2024 - May, 2025
Student number:	5016339

An electronic version of this thesis is available at http://repository.tudelft.nl/.

Abstract

Target-in-the-loop (TIL) systems are key to engaging aerial threats and enabling laser processing. This work simulates a TIL system and evaluates its performance under varying atmospheric turbulence and target distances. The goal is to assess how these factors affect energy delivery and real-world viability. A computational model was developed using standard wavefront propagation to account for turbulence and target surface roughness. The simulation allows controlled changes in turbulence strength and range. An experimental setup was also created to validate the results. Findings show that compensating for surface roughness is crucial, and system performance improves significantly with a controller—boosting peak intensity by up to 60 % and power in the bucket by 10 %. A tradeoff exists between distance and turbulence strength, but adaptive optics can help extend operational range.

Contents

List of Figures

List of Tables

List of Abbreviations and Acronyms

UAS Unmanned Aerial System

HEL High Energy Laser

TIL Target-in-the-loop

AO Adaptive Optics

DM Deformable Mirror

FWHM Full Width at Half Maximum

CCD Charge-Coupled Device

IR Infrared

COTS Commercial-Off-The-Shelf

MLA Micro-Lens Array

QWLSI Quadriwave Lateral Shearing Interferometer

SLM Spatial Light Modulator

MEMS Micro-Electromechanical Systems

S/N Signal to Noise ratio

SHWFS Shack-Hartmann Wavefront Sensor

BE/R Beam Expander/Reducer

BD Beam Dump

BS Beam splitter

PIB Power-In-a-Bucket

PSD Power Spectral Density

Chapter 1

Introduction

1.1 Target in the Loop

High powered lasers, or High-energy lasers (HELs), have revolutionized multiple industries, offering unprecedented precision, efficiency, and effectiveness in various applications. One major challenge with HELs is how to best focus their energy such that the laser is used to its maximum potential. Specifically, target-in-the-loop (TIL), is the subject of study when one cannot place sensors at a target, and must rely on the return signal for closed loop control. Additionally, a distorting medium may adversely affect the laser beam. The use of TIL spans many fields, including defense, medicine and industrial settings. In defense, these systems can be used to neutralize the increased prevalence of unmanned aerial systems (UAS) on the modern battlefield, operating at a fraction of the cost of the target they are used to eliminate [1, 2, 3]. Meanwhile, in the medical field, laser-assisted in situ keratomileusis (LASIK, or laser eye surgery) can make use of TIL algorithms to permanently remove optical aberrations in the eye of a patient [4]. Adaptive optics is also used in laser processing, such as laser cutting, to optimize the beam spot size on the workpiece [5].

1.2 High Energy Laser and Target In the Loop Challenges

Despite their appeal, HEL systems face significant engineering challenges that must be addressed. Two of the most critical challenges are the effects of the distorting medium and thermal blooming. The distorting medium, for example atmospheric turbulence, distorts the laser beam as it propagates through air, reducing its focusing potential and hence effectiveness at longer ranges. Thermal blooming, on the other hand, occurs when the atmosphere partially absorbs some of the laser beam, causing very local heating and changes in air density that further degrade the beam quality on a target. A short description of these two effects follows.

Atmospheric turbulence arises from non-homogeneous variations in pressure, wind, and temperature along the laser beam's path, leading to fluctuations in the refractive index of the air. These variations cause the beam to spread and distort as it propagates over long distances, reducing the efficiency of energy delivery to the target. Ideally, the exit aperture of a HEL system is kept small, minimizing the beam's interaction with local refractive index variations, leading the system to focus the maximum amount of energy on the target. However, as will be discussed later, a small exit aperture is not always feasible or sufficient to meet operational

requirements, such as achieving sufficient cooling for a fixed laser optical power rating. Additionally, even in the absence of an atmosphere, a larger aperture results in a lower diffractive divergence due to the aperture, leading to a smaller spot size at the target. When the aperture cannot be made small enough to mitigate the effects of atmospheric turbulence, an alternative solution is to output a pre-compensated beam. This approach uses adaptive optics, which introduce corrections that, when combined with the inverse of the distortions caused by atmospheric turbulence, will hopefully produce an optimal, diffraction-limited focal spot at the target. In practice, the spot size will be reduced but not reach the diffraction-limited size.

Thermal blooming is another critical challenge in HEL systems. It occurs when a laser beam heats the medium, air in this context, along its path due to absorption. This heating causes the air to expand, reducing its refractive index and effectively creating a divergent lens within the beam path [6]. As described by [6], "thermal blooming or defocusing occurs when, because of absorption, a laser beam heats the liquid in its path. As the liquid expands, the refractive index decreases; thus the beam creates for itself a diverging lens." This diverging effect significantly hampers energy delivery, reducing the laser's effectiveness over long distances. To mitigate thermal blooming, several strategies can be employed: decreasing the optical power, increasing the beam radius, selecting a laser wavelength with minimal atmospheric absorption or shortening the propagation distance [7]. Among these, for HEL systems with fixed optical power, only increasing the beam radius or adjusting the laser wavelength are practical options. Even so, the laser wavelength is often already chosen to, to some extent, minimize atmospheric absorption. Consequently, increasing the beam radius becomes the primary means to minimize thermal blooming. This approach, however, necessitates a larger exit aperture, which conflicts with the requirement to minimize the effects of atmospheric turbulence—a trade-off that must be evaluated during the design stage.

The trade-off between atmospheric turbulence and thermal blooming is exemplified by studies that incorporate the effects of both, for example [8]. [8] studied the feasibility of engagements of ground-based laser systems on satellites, however the basic idea is comparable to target-in-the-loop (TIL) systems, where flying targets in the atmosphere must be engaged. Figure 5 of [8] shows beam spot radius versus aperture diameter, from which it is clear that both effects work to constrain the optimal aperture diameter for minimum beam spot size. Hence, any sophisticated HEL system will greatly benefit from the pre-compensation of atmospheric turbulence.

1.3 Beam Pre-compensation

It has previously been established that some form of pre-compensation of the laser beam must be done to correct for errors arising from atmospheric turbulence. Because atmospheric turbulence involves some randomness, there must be some element in a HEL system that measures the degree of atmospheric turbulence, and applies a corresponding amount of pre-compensation. This fits the definition of a system that requires closed loop control: a sensor, actuator and controller (that transforms sensor readings and possibly user inputs into actuator commands). The graphical representation of closed loop control is shown in fig. 1.1, which clearly illustrates a loop for constant correction. To further aid the design process, TIL applications typically draw inspiration from the established field of adaptive optics (AO) for high resolution imaging from both the hardware and software.

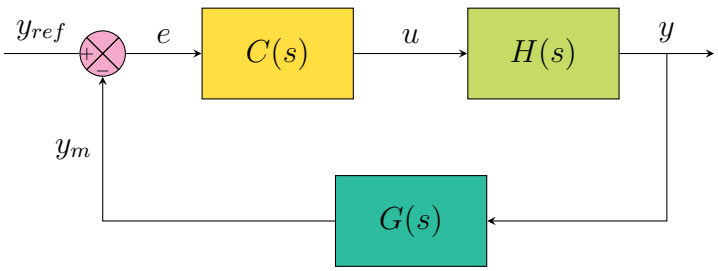

Figure 1.1: General block diagram of a closed loop controller. A user specifies the desired level y_{ref}. The difference between y_{ref} and the measured value y_m, e, forms the input to the controller $C(s)$, which produces a control signal u. u is the input to the process block, $H(s)$, which can be a compound block that incorporates the effects of the deformable mirror, atmospheric turbulence, thermal blooming and any other disturbance, causing a wavefront aberration y. y is measured by $G(s)$ and becomes y_m, closing the loop.

1.4 Adaptive Optics

Adaptive optics is a technology originally developed to correct the blurring effects caused by atmospheric turbulence in ground-based telescopes [9]. It has been instrumental in advancing astronomical observations, including the Nobel Prize-winning discovery of a supermassive compact object at the center of our galaxy [9]. In essence, AO corrects distortions in the wavefront of light, a concept fundamental to understanding how atmospheric turbulence affects optical systems.

1.4.1 Wavefront

A wavefront represents a surface of constant phase in a monochromatic propagating wave, such as single wavelength light. Wavefronts can be drawn as in fig. 1.2. In the ideal case of light traveling through a vacuum wavefronts are uniform and flat (collimated source, fig. 1.2a) or spherical (point source, fig. 1.2b). However, in the presence of atmospheric turbulence, non-homogenous refractive indices cause the wavefront to deviate and become non-flat, as in fig. 1.2c. Such distortions lead to a blurring of stars even in theoretically diffraction-limited setups.

Adaptive optics works by dynamically correcting these distortions in the wavefront. A key component of AO is the deformable mirror (DM, the actuator) which can change its shape microscopically and with great accuracy. The process relies on a wavefront sensor, which detects irregularities in the incoming light and translates them into actuator commands. A computer connects the two components, leading to a closed loop system.

In AO applications that observe the sky, a guide star is used, which in theory is a point source, but because of atmospheric turbulence it is no longer so. By recording the change from a point source using the wavefront sensor, the deformable mirror can be used to correct for atmospheric turbulence.

1.4.2 Target-in-the-loop

Target-in-the-loop represents a similar but significantly more challenging environment that typical adaptive optics [10]. As the target cannot be assumed to be cooperative, one cannot assume a point source present for the wavefront sensor. Furthermore, because a target typi-

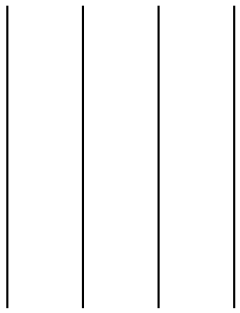

(a) Planar wavefront resulting from perfectly collimated light.

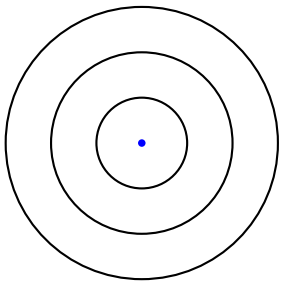

(b) Circular wavefront resulting from a point source (center).

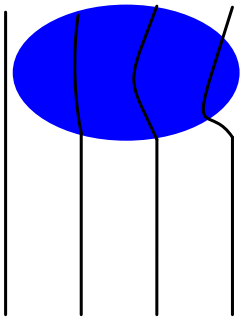

(c) Complex wavefront resulting from a planar wave partially interacting with a non-homogenous refractive index 'bubble'. Out going wavefront is partially planar, whereas the the part cannot be described in terms of planar or spherical components.

Figure 1.2: Illustration of a wavefronts. Lines indicate constant phase (wavefront by definition). For simple cases, such as perfectly collimated light or a point source, the wavefront shape is easy to describe in terms of elementary shapes. In practical applications, an alternative form of wavefront representation is necessary to model complex wavefronts. Note that in 3D wavefronts are surfaces rather than lines.

cally moves (hence the reference star would have to change), HEL engagements are typically horizontal (whereas stars are best observed straight up) and the guide star is affected by the entire atmosphere, one cannot rely on stars to provide a point source. For reliable operation, a TIL system must make a point source available on the target, and use that in the closed loop. This is typically done with a beacon laser, which focuses light to a very small spot (ideally diffraction limited) on the target. The light scattered off the target is then used in a closed loop with a HEL to efficiently deliver energy. The schematic drawing is shown in fig. 1.3a. However, this is a chicken-and-egg problem: how can one focus the beacon to a diffraction limited spot on the target through a turbulent atmosphere, if the atmosphere's turbulence is yet to be known? In practice, the beacon's projection on the target is much larger than the diffraction limited spot size, and is sometimes even left uncompensated (i.e. it does NOT pass through the deformable mirror), as is shown in fig. 1.3b. This leads to the final challenge in TIL discussed in this work's introduction: since the spot size is fairly large, and the beacon laser is fairly coherent, a speckle pattern emerges.

1.4.3 Speckle Pattern

A speckle pattern is a distinctive interference pattern that occurs when coherent light, such as that from a laser, is scattered by an optically rough surface. The roughness of the surface causes the incident light to scatter in multiple directions, with each scattered wave possessing unique phase and amplitude. The scattered waves interfere with one another, producing a granular intensity distribution known as a speckle pattern.

The phenomenon arises because the surface features are not uniform at the scale of the light's wavelength, meaning that different resolution cells (of an imaging system) contain scatterers that vary in size, shape, and orientation. This leads to phase variations of the scatted light,

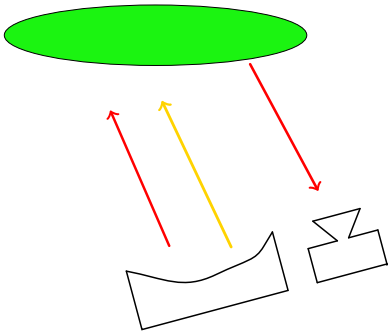

(a) Compensated beacon approach, where both the beacon beam and HEL are pre-compensated with the deformable mirror, and a wavefront sensor records the scatted beacon light off the target for closed loop control.

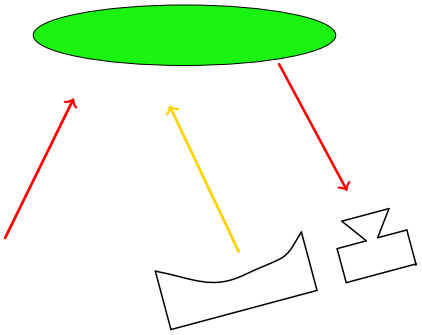

(b) Uncompensated beacon approach, where an outgoing beam is pointed towards a target. The scatted light off the target is recorded by a wavefront sensor, which through the use of a controller and deformable mirror controls the HEL.

Figure 1.3: Compensated versus uncompensated beacon approach. Green is the target, red arrow pointing towards target is beacon, red arrow pointing away from target is the scattered beacon light, and yellow is the HEL. The left black shape is the deformable mirror and the right black shape is the wavefront sensor.

making the resulting pattern highly sensitive to factors such as phase of the incident light, the wavelength, and the angle of incidence. Small changes in these parameters can cause the speckle pattern to change completely. A schematic drawing is shown in fig. 1.4a. Two points are drawn for clarity, however in reality the coherent light interacts with all points on the surface, and hence the interference pattern of all points on the surface forms the characteristic speckle pattern, as shown in fig. 1.4b.

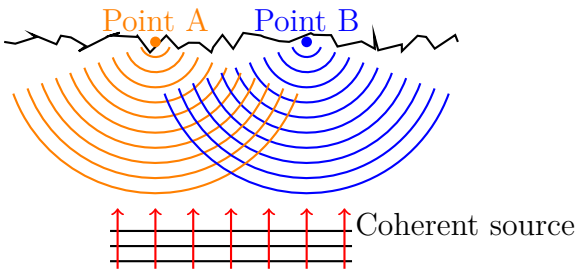

(a) Example of speckle pattern generation. Initially, a coherent source illuminates an optically rough surface. Two infinitesimally small points, points A and B, are drawn on the surface. As the coherent source intersects the surface at different spatial locations, the reflected light from both sources have different phase. The interference of all the points on the surface causes a speckle pattern to arise.

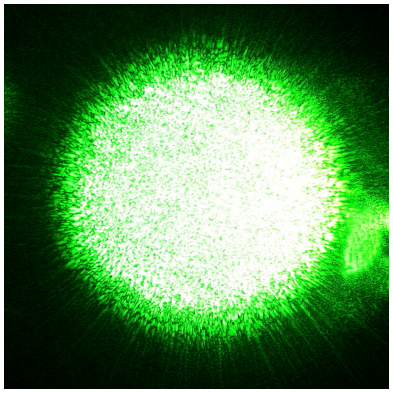

(b) Speckle pattern resulting from laser being shone onto optically rough surface and then imaged. The characteristic dark spots can be observed, especially near the fringes. Image taken from [11].

Figure 1.4: Speckle pattern example generation and imaged speckle pattern.

Speckle can be imaged in two ways, leading to two types of speckle: subjective and objective. Subjective speckle occurs when an optically rough surface is illuminated by a coherent light and it is imaged, hence the speckle pattern is observed in the image plane. Objective speckle occurs when the speckle pattern is incident on another surface, and then is imaged.

A rough surface adds a phase contribution to the incoming light. As such, a wavefront sensor cannot be used in the presence of speckle, as it could not tell the difference between atmospheric induced phase and surface induced phase, complicating TIL even further. More complex algorithms are needed to robustly correct the wavefront in TIL, which will be investigated in this thesis.

1.5 Research Question

As such, the main research question of this work is *How Can Beam Projection Be Optimized using Adaptive Optics on a Target in Target-In-The-Loop Applications in a Turbulent Atmosphere?* This research question has three components to it. (1) Beam projection optimization refers to the optimization of the shape and intensity of the beam on the target, (2) TIL consolidates a non-cooperative target at a distance, and (3) the atmospheric conditions vary according to turbulence. To answer this question, this thesis will investigate the design of a practical setup capable of yielding performance metrics relevant to TIL applications and design a TIL simulation tool that can investigate the expected performance of TIL systems.

Within the main research question, two sub-questions are of particular interest to system integrators, specifically:

- How does performance degrade with increasing turbulence strength?
- How does performance degrade with increasing target distance?

Answering both these questions will help provide a more comprehensive answer to the main research question, and it is the objective of this work to answer all these questions.

1.6 Report Outline

The thesis report begins with an in-depth insight into the literature relevant to TIL systems in chapter 2, including the physical processes associated with laser beam propagation, wavefront sensor working principle, deformable mirror working principle and robust controller design. With the knowledge of chapter 2, the experimental setup will be designed in **??**, including variations that help answer the research question(s). The exact method used to answer the research question(s) is explained in chapter 3. Further, the results of the experimental will be shared in chapter 4 and discussed in chapter 5. After discussion, the conclusions derived from the results and discussion will be drawn in chapter 6. The final chapter, chapter 7, will provide recommendations for future work.

Chapter 2

Theoretical Background

In the previous chapter, the need for robust control of a wavefront was established. In this chapter, the relevant theory and literature is investigated. The literature review begins with an introduction into laser beam theory in section 2.1. Then the controller, which links the wavefront sensor, actuator and beam together, is explained in section 2.2. The findings of the entire chapter are summarized in section 2.3.

2.1 Factors Affecting the Laser

For robust control to take place, first the physical phenomena the laser beam[1] will encounter must be understood. Hence, first, a suitable model for a laser beam is chosen in section 2.1.1. After, the theory outlining the deviation from the ideal case are explored in section 2.1.2, with an explanation of how these imperfections come to exist in a turbulent atmosphere in section 2.1.3. Finally, the theory behind atmospheric modeling is covered in section 2.1.4.

2.1.1 Laser Beam Model

In this work, two laser beam models will be used to model the beam: the Gaussian Beam model and the Wavefront model. They are explained in more detail below.

Gaussian Beam

In laser physics, laser beams often occur in the form of Gaussian beams, where laser light is almost collimated and diffraction effects need to be modeled [12]. In a vacuum, the optical intensity in the transverse plane is given by eq. (2.1), where r is the radial distance from the beam center and z is the distance from the beam focus. Furthermore, $w(z)$ is the beam radius at position z, and P is the optical power of the beam. Equation (2.1) is plotted in fig. 2.1a. The Full Width at Half Maximum (FWHM) is shown, which is a key parameter when discussing laser beam size.

$$I(r,z) = \frac{P}{\pi w(z)^2} \exp\left(-2\frac{r^2}{w(z)^2}\right) \tag{2.1}$$

[1] laser beam in this context refers to the beacon laser, however much of the same theory applies to HELs

The beam radius is given by eq. (2.2), which is of the form of a hyperbola. This form is natural, as by the diffraction limit the beam necessarily diverges with distance from the smallest beam waist radius w_0. In eq. (2.2), z_R, the Rayleigh length, can be computed according to eq. (2.3) (n is the refractive index of the medium).

$$w(z) = \pm w_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2} \tag{2.2}$$

$$z_R = \frac{\pi w_0^2 n}{\lambda} \tag{2.3}$$

The beam radius $w(z)$ is plotted against z in fig. 2.1b. fig. 2.1b shows that any Gaussian beam necessarily diverges, and that in the far field (for large z) a Gaussian beam can even be approximated by a cone with apex angle Θ. Θ is defined in eq. (2.4), where θ is half the apex angle.

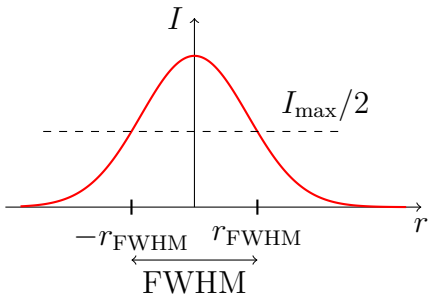

(a) Intensity pattern of Gaussian beam, with spatial coordinate x and intensity I. The full width at half maximum (FWHM) is drawn.

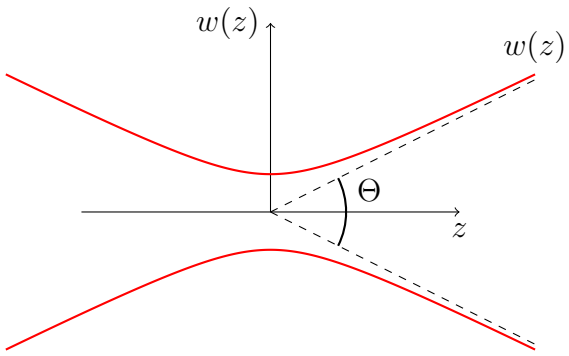

(b) Gaussian beam width $w(z)$ as a function of propagation distance z. The beam waist is modeled using the hyperbolic relation shown in eq. (2.2).

Figure 2.1: Gaussian beam profiles.

$$\frac{\Theta}{2} = \theta = \frac{\lambda}{\pi w_0} \tag{2.4}$$

Of particular interest in eqs. (2.3) and (2.4) is that these values are only true for a perfectly focused beam, and hence are only applicable as a "perfect reference" on which to compare true values.

In reality, the divergence of the beam will be larger than Θ due to external factors, such as imperfect optics and atmospheric turbulence. Atmospheric turbulence will be studied in this work, which further degrades the quality of the beam by spreading out the beam further. The extent to which the beam is changed by atmospheric turbulence will be discussed in subsequent sections.

Wavefront Model

The second model used for beam representation is the wavefront model, which is needed to understand the issues affecting the beam. This method is used as it provides a large flexibility to model many phenomena that may affect a laser beam, such as atmospheric effects, propagation effects (diffraction), effect of a wavefront actuator and wavefront sensor. The theory is derived in appendix B, with two important results moving forward: (1) a wave can be represented in the plane perpendicular to its propagation as a complex field of both phase and amplitude,

and (2) this field can be propagated a distance with a Fourier transform, multiplication and an inverse Fourier transform. This analysis sets the stage for studying wavefront aberrations, discussed in the following subsection.

2.1.2 Wavefront Aberrations

While the Gaussian beam representation was useful to model concepts such as the diffraction limit, in reality beams can take on many different shapes. Three examples are shown in fig. 2.2.

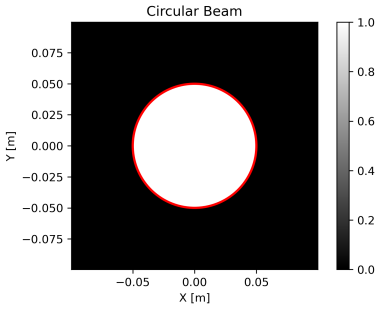

(a) Circular beam cross section. The radius is defined by the radius of the circle.

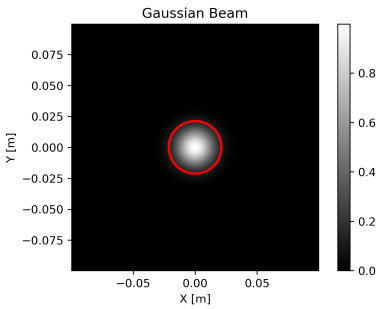

(b) Gaussian beam cross section. The radius is defined by the circle where the intensity is $\frac{1}{e^2} \approx 0.135$ times the peak intensity.

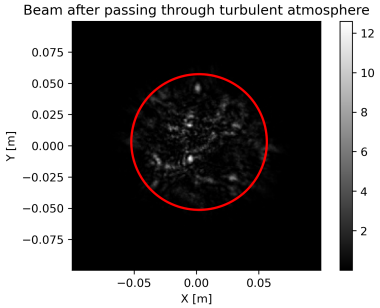

(c) Beam which has gone through atmospheric turbulence. Here, the beam waist is defined by four times the standard deviation.

Figure 2.2: Defining different radii to different beam shapes. There is not one common definition for all conditions, hence once must be chosen for the specific application.

As shown above, beams need not circular, hence a value for r must be chosen to faithfully model the desired effect. For this work, r is the FWHM as shown in fig. 2.1a, however in a different situation a different defintion for r could be chosen.

When discussing aberrations in an optical sense, one is interested in the relative phase difference between different points (x, y) within a beam slice such as the ones shown in fig. 2.2. Phase differences can be represented in many ways, such as in physical distances (e.g. Δz, nm), angular distances (e.g. $\Delta\phi = \frac{2\pi}{\lambda}\Delta z$, rad) or as a ratio of the given wavelength (e.g. $\Delta\lambda = \frac{\Delta z}{\lambda}$, unitless). Importantly, one is not interested in the mean phase difference within an optical slice, hence the mean is often defined as zero when unavailable. Additionally, any deviations from zero phase represent wanted behavior of the optical beam: with zero phase throughout the beam, one can reliably operate on the beam (for example by focusing onto a sensor). Any deviations cause degraded performance.

To simplify the analysis of wavefront aberrations, it is common to represent a wavefront as a linear sum of Zernike polynomials, shown in appendix C.

With a understanding of wavefront representation, the effect of the atmosphere on the wavefront is now explored.

2.1.3 Effect of the Atmosphere on a Wavefront

There are several phenomena that act on laser beam propagation through a real atmosphere. The beam loses its energy quickly due to molecular scattering, molecular absorption and particulate scattering [13]. In the case of turbulence, beam quality is degraded by distorting its

wavefront and by randomly varying the signal power at the receiver [13]. There are several effects that make up turbulence, two of which will briefly be discussed below. More information can be found in [13], however the source writes about atmospheric turbulence in relation to laser communications.

Scintillation The turbulence in the atmosphere slightly spatial shifts the center of the power of the beam, changing with time [13]. This effect is known as scintillation, and is formally defined as the normalized variation of intensity fluctuations. The equation is shown in eq. (2.5).

$$\sigma_I^2 = \frac{< I^2 > - < I >^2}{< I >^2} \tag{2.5}$$

Beam Broadening As was discussed previously, a Gaussian beam necessarily diverges with increasing distance. However, in the presence of a turbulent medium, the divergence is further increased, called beam broadening. This happens as turbulence disrupts the (perhaps uniform) wavefront, comparable to splitting it up into smaller section with marginally different propagation vectors. In a turbulent atmosphere, the beam width is given by eq. (2.6) [13]. C_n^2 is the Refractive Index Structure Parameter, which for very weak turbulence is $1 \times 10^{-18}\,\mathrm{m}^{-2/3}$ and for strong turbulence $1 \times 10^{-13}\,\mathrm{m}^{-2/3}$ [13]. l_0 is the inner scale, and will be discussed below. z is the distance from the beam source.

$$w_t(z) = w(z) + 4.38 C_n^2 l_0^{-1/3} z^3 \tag{2.6}$$

In this work, molecular scattering, molecular absorption and particulate scattering are not considered. Rather, minimizing the energy redistribution due to atmospheric turbulence is addressed. Turbulence originates from chaotic mixing of air, which occurs at high Reynolds numbers [13]. The Reynolds number can be computed with $Re = \nu L/\mu$, where ν is the characteristic flow velocity, L is the characteristic length of the system being considered, and μ is the kinetic viscosity of the fluid. Within the atmosphere, kinetic energy is usually introduced by wind shear and convection from solar energy at lengths greater than L_0, called the outer scale [13]. The kinetic energy is subsequently transferred to smaller and smaller scales through turbulence. At small enough scales, energy dissipation starts becoming viscous [13]. The largest of these small scales is called the inner scale of turbulent flow [13]. Typical values for l_0 are 1 to 10mm, and for $L_0 \sim$ 10 to 100m.

With this information, and eq. (2.6), it is possible to perform a simple calculation to estimate the power loss (in a bucket) due to atmospheric turbulence on an uncorrected (no adaptive optics) beam. The analysis begins with an estimate for the power still contained within a radius R at a distance z, known as $\eta(R, z, w(z))$ and shown in eq. (2.7) (derived in appendix A).

$$\eta(R, z, w(z)) = 1 - \exp\left(-2\frac{R^2}{w(z)^2}\right) \tag{2.7}$$

Next, the ratio of the remaining power within the radius for the atmospheric case versus the

$$\begin{aligned}\eta_{remaining}(R, z, C_n^2, l_0) &= \frac{\eta(R, z, w_t(z))}{\eta(R, z, w(z))} \\ &= \frac{1 - \exp\left(-2\frac{R^2}{w_t(z)^2}\right)}{1 - \exp\left(-2\frac{R^2}{w(z)^2}\right)}\end{aligned} \tag{2.8}$$

Equation (2.8) can be plotted against distance for a better understanding of power loss due to propagation through a turbulent atmosphere. This is shown in fig. 2.3 for various atmospheric turbulences and inner scales.

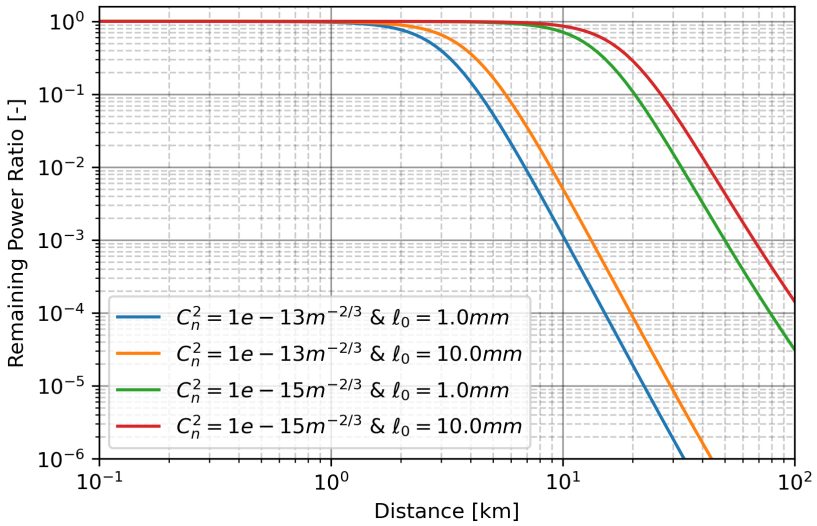

Figure 2.3: Signal attenuation for weak and strong turbulence (C_n^2 is 10^{-15} $\mathrm{m}^{-2/3}$ and 10^{-13} $\mathrm{m}^{-2/3}$, respectively, and different inner scales. Past 3 km the power is reduced by a factor of 10, whereas for the weaker turbulence this happens at 20 km.

Atmospheric Statistics

The small changes in the refractive index can be modeled with eq. (2.9) [13]. Here, $E[\cdot]$ represents ensemble averaging, n_1 is the variation of the refractive index from the mean, which varies with spatial location r, and obeys $E[n_1(r)] = 0$. $E[n(r)] \approx 1$ for Earth's atmosphere.

$$n(r) = E[n(r)] + n_1(r) \tag{2.9}$$

As the atmosphere is stochastic in nature, one would prefer to obtain a statistical function to perform further analysis. When modeling atmospheric turbulence, the statistical function typically used is the structure function, for example eq. (2.10)[14], which describes how a field varies over spatial separations.

$$D_n(r) = E[|n_1(r + r_1) - n_1(r_1)|^2] = \begin{cases} C_n^2 r^{2/3} & l_0 < r < L_0 \\ C_n^2 l_0^{-2/3} r^2 & r < l_0 \end{cases} \tag{2.10}$$

The structure function can be converted to a power spectrum, as shown in eq. (2.11), with κ being the frequency coordinate corresponding to r [14].

$$\Phi(\kappa) = 0.033 C_n^2 \kappa^{-11/3} \ ,1/L_0 < \kappa < 1/l_0 \tag{2.11}$$

With the power spectrum, one can readily generate a field of phase differences with the correct statistics, which will be shown explicitly in section 2.1.5, where the Wiener-Khinchin theorem is used [15]. The refractive indices can be converted to phase shift using eq. (2.12), where λ is the wavelength of light and d is the distance the light travels through the medium.

$$\Delta\phi = \frac{2\pi}{\lambda} \cdot \Delta n \cdot d. \tag{2.12}$$

The culmination of all the theory presented previously results in an $N \times N$ field of phase differences, which can be used to simulate turbulence for a small section of the atmosphere. Additionally, it can be advanced in time by shifting the screen and generating new elements with the same statistical properties using [16]. The phase screens add to the phase of the wavefront as defined in section 2.1.1. In the following section, multiple phase screens will be used to simulate larger extents of the atmosphere.

2.1.4 Atmospheric Modeling

With a single phase screen, ϕ, one can simulate a small section of the atmosphere. With multiple phase screens, ϕ_i where $i = 0, 1, \ldots, n$, one can simulate the entire propagation path through an atmosphere. When disregarding scintillation, the total phase of a light beam is the sum of all ϕ_i eq. (2.13) [16].

$$\phi = \sum_{i=0}^{n} \phi_i \tag{2.13}$$

However, scintillation is expected to cause significant issues in TIL applications [17]. Hence, rather than combine the phase screens into one, the field is propagated (as shown in section 2.1.1) to the next phase screen, and then the phase is added. This process is repeated for each phase screen used. Figure 2.4 shows an example of 3 phase screen separated by a distance d. In practice, the number of phase screens is determined through a combination of turbulence strength and propagation distance, see [17].

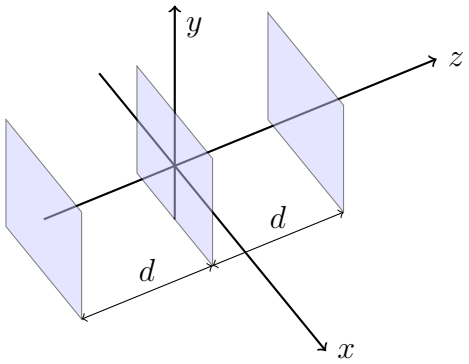

Figure 2.4: 3 phase screens separated by a distance d. The propagation direction is z and x and y are the transverse axes.

The presented wavefront theory regarding atmospheric modeling is only part of the total wavefront needed to continue this work. In the following section, the effect of a surface will be investigated.

2.1.5 Surface Scattering

The surface in TIL applications represents a significant challenge. Assuming a non-cooperative target, no favorable assumptions can be made about the target surface. The most favorable target surface is that of a mirror: the phase of the reflected beam just has a π shift, which does not adversely affect the performance of the TIL system. The most unfavorable target surface is one that reflects very little, and the reflected wave has a random spatially depending phase which rapidly fluctuates. The random phase causes a speckle pattern, While the little reflected light forms part of a link budget, which is outside of this work, the varying phase must be mitigated. The strategies behind the mitigation of speckle begin with an analysis into speckle generation.

To begin the analysis, some assumptions will be made on the statistics of the target surface. Specifically, the target surface will be modeled by the autocorrelation function given by eq. (2.14) and the roughness variance $\sigma_s^2 =< \xi^2(\vec{r}) >$, where $\xi(\vec{r})$ is the surface height due to the rough surface at position $\vec{r}$ and l_s is the transverse roughness correlation length [10]. Note that this assumption does not limit the applicability of this work, as when $l_s \to 0$ any two points on the field become uncorrelated, i.e. white noise.

$$< \xi(\vec{r_1})\xi(\vec{r_2}) >= \exp\Big(\frac{-r^2}{2l_s^2}\Big) \tag{2.14}$$

To convert the autocorrelation $< \xi(\vec{r_1})\xi(\vec{r_2}) >$ into a generated field ξ of realized surface roughness, the Wiener–Khinchin theorem is used, which states that the autocorrelation function and the Power Spectral Density (PSD) are related via the Fourier transform. Further, because a PSD is the square amplitude of the Fourier transform of a signal, one can apply the inverse transformation with fictitious generated phase terms to obtain a realized field ξ, the surface roughness.

Mathematically one has the following. First, the Fourier transform of the autocorrelation function is taken to yield the PSD $S(f)$. Next, the definition of PSD means eq. (2.16) is true. Finally, the two expressions can be combined, as shown in eq. (2.17), to obtain a value for ξ.

$$\mathcal{F}\Big(< \xi(\vec{r_1})\xi(\vec{r_2}) >\Big) = S(f) \tag{2.15}$$

$$S(f) = |\tilde{\xi}(f)|^2 \tag{2.16}$$

$$\mathcal{F}^{-1}\Big(\sqrt{|\tilde{\xi}(f)|^2}e^{2\pi i\phi}\Big) = \xi \tag{2.17}$$

To convert the generated surface profile ξ (in units of length) to a phase shift, eq. (2.18) is used, originating from eq. (2.12).

$$\Delta\phi = \frac{2\pi}{\lambda}\xi \tag{2.18}$$

This method could also be used in the generation of phase screens for atmospheric turbulence, however these computations may also be offloaded to more extended algorithms that provide features like wind simulation, continuous phase screen generation, and more optimized implementations. For the generation of surface roughness profiles, the Fourier method is employed.

From the perspective of simulation, almost all the theory of optical components has been covered: all that is left are the wavefront sensor and actuator, which are discussed in more depth in appendices D and E. In the following section, the controller will be designed and finally selected.

2.2 Controller Design

With the physical theory established, one can now move on to discussing the method by which speckle is mitigated. It is necessary to compensate the atmospheric aberrations to properly focus a beam onto a distant target. Hence, it is necessary to pre-correct the wavefront using a deformable mirror. To do so, the aberrations would need to be known. The linking from measured aberrations to deformable mirrors actuation commands is performed by the controller.

Ideally, as in astronomical adaptive optics, one would have a guide star, from which the aberrated wavefront directly corresponds to the atmospheric induced phase aberrations. In this case, one can use a wavefront sensor to measure the wavefront $W(r, \phi)$, and apply the corrections based off the measurements through a vector $\vec{u}$ to the deformable mirror.

This method works first by calibrating the setup with a planar wavefront $W(r, \phi) = 0$, and then 'poking' the wavefront actuator signals u_i, noting the changes in the wavefront $\partial W_i(r_j, \phi_j)$ at different points (r_j, ϕ_j). Having poked all the actuators, one can now predict what small changes in $\vec{u}$ will have on $W(r, \phi)$. The vector equation is shown in eq. (2.19), where there are n points used for calibration and m actuators. The poke matrix is known as Φ for the rest of the work.

$$\partial W(r,\phi) = \begin{array}{c} \\ j=0 \\ 1 \\ \vdots \\ n \end{array} \begin{array}{c} \begin{array}{cccc} i=0 & 1 & \dots & m \end{array} \\ \begin{pmatrix} \partial W_0(r_0,\phi_0) & \partial W_1(r_0,\phi_0) & \dots & \partial W_m(r_0,\phi_0) \\ \partial W_0(r_1,\phi_1) & \partial W_1(r_1,\phi_1) & \ddots & \vdots \\ \vdots & \ddots & \ddots & \vdots \\ \partial W_0(r_n,\phi_n) & \dots & \dots & \partial W_m(r_n,\phi_n) \end{pmatrix} \end{array} \cdot \partial \vec{u} \tag{2.19}$$

The matrix that relates $\partial W(r, \phi)$ to $\partial \vec{u}$ can be inverted, using the Moore–Penrose inverse, or SVD for more sophisticated implementations [18]. Then, when a wavefront aberration is measured, one can find the required $\partial \vec{u}$ such that the aberration is minimized. More complex controllers would also make use of a state estimator, that can use sensor fusion or atmospheric models to improve performance [18].

The issue in TIL applications is that there is no guide star, and hence the target must be illuminated by some other means. Typically, the beam size is much larger than a diffraction limited spot, a hence some speckle is present. When using a wavefront sensor, the speckle presents itself as another phase term. The goal of a TIL controller is to overcome speckle and apply the well known adaptive optics algorithms to maximize power delivery. The following subsections will contain an overview of the investigated controllers, including the ones without a wavefront sensor, and their working principle, eventually leading to the selection of the most suitable one for this work.

2.2.1 Speckle Average

The first method presented in this work that can mitigate target-induced effects caused by target surface is called *Speckle Average*. The idea behind *Speckle Average* is simple: realize many speckle patterns within one frame of the wavefront sensor. With enough speckle patterns realized, the average measured wavefront will encode just atmospheric-induced phase aberrations, and speckle based phase aberrations will be "averaged out" [10]. One way to realizing many speckle patterns is by using a fast jittering method, where the beam projection is quickly

moved from side to side. The frequency of this movement would have to be much larger than the frequency of the closed loop, which is limited by the wavefront sensor to around 100 Hz. This could mean a jitter rate of around 1 kHz to 10 kHz [10].

While technically achievable with mechanical components, perhaps even a deformable mirror, this is a typical use case of a SLM. Alternatively, [10] proposes rotating the target to realize many speckle patterns. This is not unrealistic to some real-world examples. For example, many missiles and rockets spin along their travel direction to maintain stability, meaning enough speckle patterns would be generated. Within the context of a lab, it is also possible to make use of a hand drill to rotate the target and realize enough speckle patterns. For the purposes of this work, it should be possible to rotate the target, with the understanding that a system based on this work would use a SLM for rapid beam jitter.

With enough speckle patterns realized, the wavefront sensor can measure an adequate wavefront so as to find the correct control vector.

2.2.2 Polychromatic Speckle

A second approach, which was also mentioned in [10], is to use a source with lower coherence, also known as a polychromatic light source [17]. The reason this works is because speckle is sensitive to the wavelength of light used, and hence using a source with many wavelengths (in a continuous fashion the equivalent would be a broadband source) would wash out the speckle induced phase, and leave behind the atmospheric induced phase. There are two key metrics in the polychromatic approach, Object-subaperture fresnel number, N_F, and coherence lengths in a resolution cell, N_c. The equations for both are shown in eqs. (2.20) and (2.21), where D is the subaperture diameter, T is the width of the object or width of the extended beacon (the smallest), λ is the wavelength, Z is the distance between the object and subaperture, l_c is the coherence length and θ is the angle between the objects surface normal and the line of sight [17]. It is important to note that l_c is a measure of the coherence of the light source, with smaller number indicating a more incoherent source.

$$N_F = \frac{DT}{\lambda Z} \tag{2.20}$$

$$N_c = 3.5\frac{\lambda Z}{D l_c}\tan\theta \tag{2.21}$$

In [17], the authors studied the effect of varying N_F and N_c on power-in-the-bucket in weak, strong and up-looking scenarios. For a wavelength of 1.0 µm, distance of 3 km, and subaperture diameter of 3 cm. They studied varying N_F from 0.25 to 2 and N_c from 0.03 to 32. Their results show that a lower N_F resulted in better performance and higher N_c drastically increased performance.

To use polychromatism in this work, N_F would ideally be as small as possible, and N_c would be as high as possible. With the conditions outlined in [17] fulfilled it should be possible to directly measure the wavefront and apply the voltages to the deformable mirror.

2.2.3 Temporal Speckle

The third investigated speckle mitigation strategy is *temporal speckle*. Unlike the previous two methods, this does not make use of a wavefront sensor, but instead a photodiode. Two cases are considered: (1) the beam is not on the target, (2) the beam is on the target.

In the case the beam is not yet on the target, one would like to optimize the power in the bucket metric of the received signal. This is because as soon as some light is scatted off the target and enters the receiver aperture, the system knows that the beam is on target. However, once the beam is on target, the power in a bucket metric is no longer useful as it just depends on the intensity of the outgoing wave. This is where the second metric comes in.

The second metric is the standard deviation of the power the bucket metric (or moving standard deviation for prolonged exposure). [4] proved that maximizing the standard deviation resulted in maximizing the sharpness metric on the target.

Linearly combining these two with a weighting factor results in a compound metric that can be used to optimize the shape on the deformable mirror. The structure of the minimization problem is given eq. (2.22), where $\vec{x}$ is control actuator vector given to the deformable mirror, and $f(\cdot)$ is the compound metric.

$$\min_{\vec{x}} f(\vec{x}) \tag{2.22}$$

Alternatively, fig. 2.5 shows a visual plot of the minimization (which Voronstov framed as a maximization) problem.

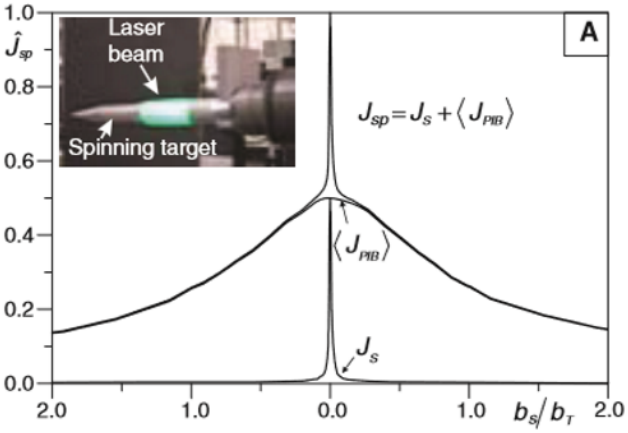

Figure 2.5: Composite cost function as formulated by Voronstov. b_s/b_T is the ratio between beam and target sizes, and $\hat{J}_{sp}$ is the normalized cost function. J_s the standard deviation of recently sampled PIB, and $< J_{PIB} >$ is the average of the recently sampled PIB. Image and figure from [10].

The strategy for temporal speckle is to use a nonlinear optimization algorithm iteratively to find the best $\vec{x}$ such that $f(\vec{x})$ is minimized. Hence, a wavefront sensor is not used in this method.

2.2.4 Spatial Speckle

Similar to the temporal speckle approach, the spatial speckle method does not make use of a wavefront sensor. Unlike the temporal speckle, it uses a imaging sensor instead of a photodetector. It seeks to maximize the speckle size by computing an image metric, that can be used for optimization as in eq. (2.22) [4].

The image optimization metric used is a two step process: firstly, $I_{sp}^{bin} = sign\left[I_{sp}(\vec{r}) - \bar{I}_{sp}\right]$ is computed, where $I_{sp}(\vec{r})$ is the image captured by the imaging sensor and $\bar{I}_{sp}$ is the mean value.

In essense, it is a binary image that is used to enhance the speckle pattern contrast. Secondly, the metric is computed $J_{sp} = \int |\nabla I_{sp}^{bin}|^2 d\vec{r}$, where ∇ is the digital gradient Sobel operator and the integration is performed over the entire image [4].

Similar to temporal speckle, the metric can be used in a nonlinear optimizer to find the best $\vec{x}$ iteratively.

2.2.5 Compensated Beacon

The last approach considered in this work is that of using a compensated beacon or not. [19] studied the effect of compensating the beacon, versus an uncompensated TIL setup. According to the study, using a compensated beacon results in a performance boost of 17 % over uncompensated. This result indicates that the final setup should use a compensated beacon for maximal performance.

Within the paper the authors directly applied the AO control loop, without mitigating speckle [19]. This means that for this work compensated beacon must exist with another speckle mitigation strategy.

2.2.6 Controller Selection

With an overview of the available control schemes complete, it is now possible to judge each with a trade off table. The relevant criteria on which each option will be judged are as follows: follows: follows:

- *Implementation Feasibility* refers to the anticipated computational demands of each controller. Approaches requiring significant computing resources are considered less favorable due to potential limitations in simulation speed and scalability.
- *Integration with Existing Adaptive Optics (AO)* captures how well a method aligns with the current AO simulation framework. Techniques that require substantial modifications or new simulated hardware increase development time and debugging complexity, detracting from core research goals.
- *Available Literature* assesses the presence of prior work supporting the method. Well-documented approaches are preferred, as they provide a foundation for comparison, validation, and further development.
- *Assembly Complexity* estimates the practical difficulty of implementing the method in a hardware system. Simpler physical setups are favored, as they are easier to prototype, test, and iterate upon in future experimental stages.

Finally, with the criteria set, the trade off table is shown in fig. 2.6.

Ideally, an implementation makes use of the extensive research available with adaptive optics, hence a solution that uses a wavefront sensor is preferable over one that does not (speckle average and polychromatic speckle). As it is known that HELs have a wide spectrum around their center wavelength, and it is desirable to choose a method that will converge, speckle average and polychromatic speckle are chosen for this work. Further studies could investigate the effect of compensated beacon (compensating the outgoing wave) to potentially increase performance.

Method	Implementation feasibility	Integration with existing AO	Available literature	Assembly Complexity
Speckle average	Averaging of return wave	Increase exposure time	AO literature	Medium; dithering mechanism required
Compensated beacon	Possible instability	No major changes	Little available	Easy
Spatial speckle	Image processing; 1D optimization on single speckle realization	New control scheme required	Little available	Easy
Temporal speckle	1D optimization over many speckle realizations	New control scheme required	Little available	Medium; dithering mechanism required
Polychromatic speckle	Averaging of return wave	Use polychromatic sensor	Little available	Easy

Figure 2.6: Controller Trade-off with the 5 different methods. Green means greater than desired performance, yellow acceptable performance and red unacceptable performance.

2.3 Literature Summary

To begin the analysis of TIL applications, the literature study began by studying an appropriate method to model the beam in section 2.1. Section 2.2 looked at the various methods to mitigate speckle in TIL applications, and selected a primary speckle mitigation method to test further. With the theory behind TIL covered, the practical considerations of this work will be sorted out in the next section, including a test and simulation plan.

Chapter 3

Methods

Previously, the theory behind target in the loop simulations was presented. In this chapter, the theory will be put into practice through a series of simulations. The first set of simulations, S.I: Individual Simulations (section 3.1), are intended to evaluate at the individual components of the larger simulation. This step is crucial to ensure the components function as intended, and simplifies any problems that may appear in the full simulation later on. Once the individual components are tested, they are put together in S.II: Assembly Simulations (section 3.2). The purpose of S.II is to simulate the expected performance of an actual TIL system. In chapter 4, the methods presented in this chapter will be implemented and the key results will be shared.

Originally, this thesis was intended to include a hardware testing. Due to unforeseen circumstances, it was changed to a purely simulation-based study. However, the planification of the hardware setup is provided in appendix H, and can help explain some design decisions in the simulation.

3.1 S-I: Individual Simulation

The individual simulations of TIL components serve as a starting point to ensure their correction functioning. The simulations focus on the surface roughness and its effect on a wavefront, as this is the main focus of this work. As a standard library will be used to simulate the other components, which has already passed testing, these will not be tested as thoroughly [16].

3.1.1 S-I.1: Surface Roughness Statistics

The first simulation that will be carried out is one used to confirm the statistics of the generated surface roughness profiles. This is done by generating multiple surface profiles, and generating two key statistics from the field: (1) the autocorrelation function as a function of distance, and (2) the two dimensional autocorrelation field. The reason the two methods are used is because the first method can directly be used to compare the quality of the generated statistics, however it scales with the fourth power of the input (each pixel in the two dimensional space requires a computation for each other pixel), and hence is expected to take an unreasonable amount of time for medium to large resolutions. The second method generates the autocorrelation field using the PSD and Wiener–Khinchin theorem, applied in the reverse direction as was used to generate the fields. A success on both these tests would indicate that the generated field has the correct statistics, and hence can be used for TIL simulation.

3.1.2 S-I.2: Speckle Mitigation: Speckle Average

The second simulation will investigate the feasibility of speckle mitigation in favorable conditions: no atmospheric effects and short propagation distances. After generating multiple surface roughness profiles, the incoming wavefront will be propagated a distance, and at distance the incoming wavefront will be averaged. A successful test would be one where the effect of speckle can be observed to decrease with increasing number of surface roughness profiles.

3.1.3 S-I.3: Speckle Mitigation: Polychromatic Speckle

The third experiment in the simulation evaluates the effectiveness of using a polychromatic light source rather than a monochromatic one. This simulation will first generate one surface roughness profile, and simulate the interaction of this surface with multiple wavelengths of light. Similar to simulation S-I.2, the scattered wave off the target will be propagated a short distance away, after which it will be averaged. A successful test would be one where the effect of speckle can be observed to decrease with increasing number of wavelengths.

Having concluded the S-I series of simulations, the next round of simulations will integrate all the components into one simulation: Target-in-the-loop simulation.

3.2 S-II: Assembly Simulations

Having studied the individual aspects of a full TIL system, they can be now brought together to test and evaluate the performance of a TIL system. The work begins by first integrating all the components except for the speckle pattern, simulating a mirror-like (simulating a retro-reflection) reflection off the target in section 3.2.1. Next, in section 3.2.2, the mirror like reflection is replaced with the wave scattering discussed and simulated in section 2.1.5, forming the complete TIL system. The full simulation overview is presented in section 3.3, and a chapter summary is provided in section 3.4.

3.2.1 S-II.1: Speckle-Free TIL

The first simulation of TIL feasibility will be in a speckle-free scenario, i.e. zero phase contribution from the reflection off the target. The schematic of the simulation setup is shown in fig. 3.1, which was inspired by [17]. The Tx/Rx optics consist of just the known aperture, and the wave passing through it. This field is propagated using the Fourier transform to the first phase screen, which possess known statistical properties. 5 other phase screens a distance of $\frac{d}{6}$ are simulated, eventually reaching the target surface. Because this simulation is speckle free, the target surface contains zero phase contributions, and the wave is just reflected, in the opposite order, through the phase screens to reach the Rx optics, where a wavefront sensor is present.

The parameters of the closed loop control will be determined by a standard calibration technique, where the deformable mirror actuators are each 'poked', and the response is recorded by the wavefront sensor. Physically, the calibration procedure is analogous to having a point source close to the Tx/Rx optics, and changing the shape of the deformable mirror to register the changes in phases at the wavefront sensor. The schematic for the setup is shown in fig. 3.2, where the same Tx/Rx optics as in fig. 3.1 are shown, however now a reflective surface is used in close proximity to calibrate the optics.

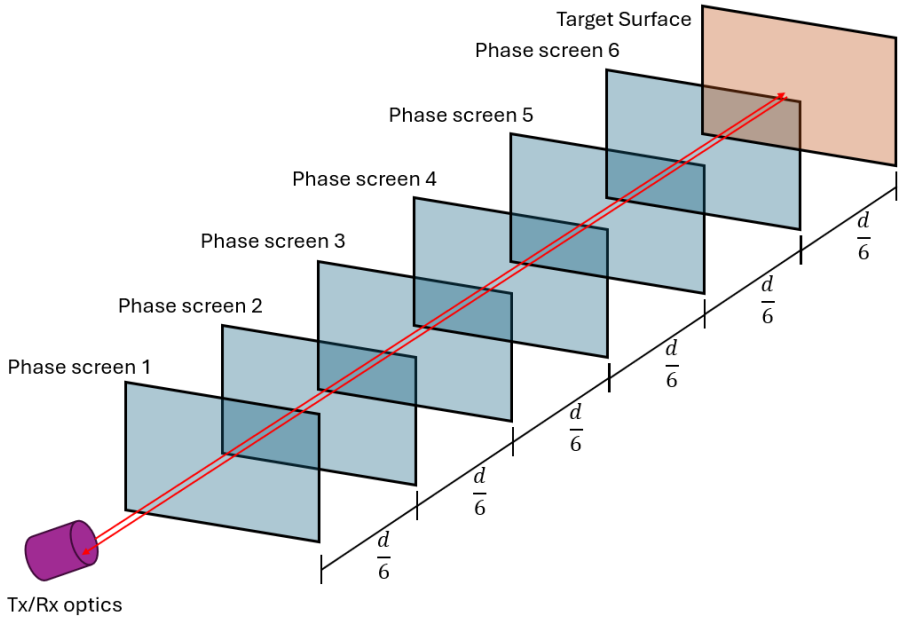

Figure 3.1: Simulation setup, showing the Tx/Rx optics, after which the wavefront is propagated through 6 phase screens intended to simulate the atmosphere. After the phase screens, the wavefront encounters the target surface and is scattered back to the Tx/Rx optics, again going through the phase screens.

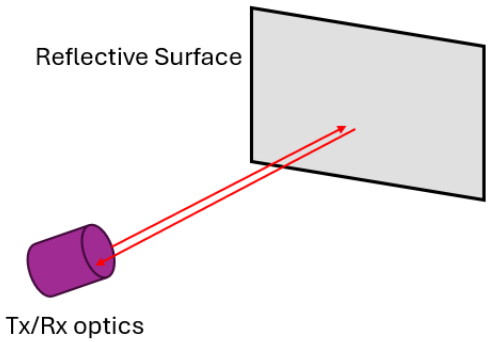

Figure 3.2: Calibration setup, where Tx/Rx optics are calibrated against a perfectly reflective target surface that is undisturbed by distortions.

Having a solid fundamental of the TIL system working without speckle, the next simulation will introduce speckle and investigate its effect on performance.

3.2.2 S-II.2: Full TIL

The final simulation will be that of the full TIL system in action: the reflected waves will be affected by speckle, and using the speckle mitigation techniques investigated earlier the system will attempt to mitigate the speckle effects to achieve performance similar to the speckle-free TIL simulation. This is a simple modification of fig. 3.1, whereby instead of a reflective target surface, a rough scattering surface is used.

3.3 Simulation Setup

With a representation of the speckle free and full TIL simulation in mind, it is now possible to place the simulation components in such a way that they can be used to generate results. This is done in section 3.3.1, where the setup is shown. With an overview of the setup, it is possible to delve into the preparation for the main simulation (section 3.3.2), and then the actual simulation execution in section 3.3.3. However, the direct results of the simulation need to be processed to get a better picture of what is happening, which is done in section 3.3.4.

3.3.1 Simulation Components

The path the beam takes throughout the simulation is shown in fig. 3.3. The first step, as shown in fig. 3.3a, in the control loop is to propagate a planar wavefront forward through a deformable mirror, applying a phase shift that's intended to compensate for the following atmospheric turbulence. Then it propagates through the atmosphere, accumulating wavefront aberrations along the way, in addition to the diffractive effect of passing through the transmitting optics. At the target surface, zero phase shift is applied, conforming to the mirror-like reflection of the speckle-free simulation. Then, following the return path fig. 3.3b, it propagates in reverse order the atmosphere, eventually reaching the transmitting optics, passing backwards through the deformable mirror and then into the wavefront sensor.

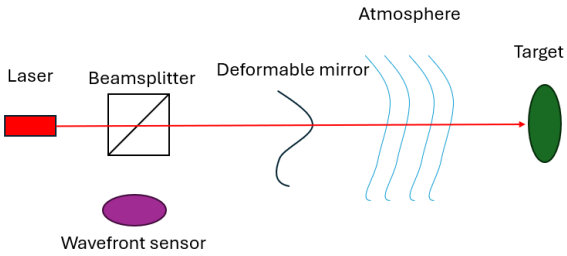

(a) Outgoing wave path, from laser, deformable mirror, beam-splitter (not simulated), atmosphere and then the target.

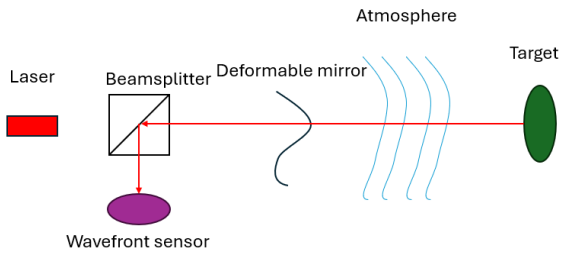

(b) Return wave path, from target, through atmosphere, deformable mirror, beam-splitter (not simulated) and into the wavefront sensor.

Figure 3.3: Simulation order for the outgoing and return wave. The beamsplitter is not modeled in the simulation, however the differences in the outgoing and return paths are simulated.

With a good overview of the main components of the simulation, the initialization of each component will be covered in the next subsection.

3.3.2 Simulation Precursor and Values

A critical aspect of a successful simulation is the correct initialization of all components. Specifically, while the individual components can be initialized with the theory presented previously, the most complex interaction is between the wavefront sensor and the deformable mirror. Specifically, before the main simulation can take place, the matrix that relates the wavefront sensor measured slopes to the deformable mirror actuators must be created. It must also be created in a realistic way: one cannot create such a matrix based off a target a certain distance away assuming very little atmospheric turbulence. For this reason, in this work the matrix is created using the standard approach in typical adaptive optics setups.

The calibration procedure is assumed to take place long before an actual engagement with a target, in a controlled environment. A planar wavefront first passes through the deformable mirror and is then registered on the wavefront sensor. Using the theory from section 2.2, a matrix can be generated, which relates the two components through a single matrix. The matrix is called the poke matrix.

There is one issue, however: the poke matrix relates a known actuator movement to an unknown wavefront sensor slope change. In a true setup, one measures the wavefront slopes to then obtain the required actuator input vector. Furthermore, the poke matrix cannot be assumed to be square, as different setups may use different equipment. Even so, methods exist to obtain a generalized inverse of a matrix, which can then be used to reverse the transformation from a known slope measurement to a unknown actuator command. One such example is the The Moore-Penrose Inverse [20]. However, based off the authors previous experience, this inverse can be malformed: small changes in the slope due to an actuator, for example an actuator transversely far away from the lenslet causes a very small change in slope. The very small change in slope, when the inverse is taken, produces a very large value for which the system may become unstable [18]. For this reason, the Tikhonov regularization method is used to obtain the inverse, which suppresses small changes in slopes [21].

The simulation parameters used are shown in table 3.1. These are the default values used unless otherwise noted.

Table 3.1: Simulation parameters used unless otherwise noted.

Parameter Name	**Parameter Symbol**	**Value**	**Justification**
(Center) Wavelength	λ	500 nm	Wavelength used for practical testing
Deformable Mirror Actuators	-	30	Comparable to the 40 actuators present on the DMH40-P01 [22]
Tx/Rx Diameter	-	10 cm	
Simulation Grid Size	-	20 cm	Double the Tx/Rx diameter
Number of Lenslets	-	100	Comparable to the 121 spots on the WFS20-5C/M [22]
Correlation Length	l_s	2 mm	Larger value comparable to simulation grid step size
Surface Roughness Std. Dev.	σ_s	1 µm	[10] states greater than wavelength
Number of pixels	-	256×256	

3.3.3 Simulation Execution

Assuming the required parameters are setup as explain in the previous section, it is now possible to focus on the execution of the simulation. The simulation algorithm is shown in section 3.3.3. It is run until the the total number of iterations are reached such that the

controller can converge to a steady state solution. Initially, the wavefront from the laser source is wf, which has zero phase aberrations and remains constant throughout each iterations. First, the forward path is computed by first modifying the phase of the wavefront using the deformable mirror, and then propagating the wave through the atmosphere. Importantly, $atms\_forward$ is used, which is an object with the atmospheric layers ordered in an increasing manner (i.e. $1, 2, ...$). At this point, the wavefront has theoretically reached the target plane, where a phase modifying element would be added if the simulation were not speckle free. Instead, the return path is directly simulated, first by propagating the wave back through the same atmosphere (using $atms\_backward$, which has its layers ordered in decreasing fashion). The phase is again modified by the deformable mirror, after which the slopes are directly obtained from a Shack-Hartmann wavefront sensor. The final part of the algorithm uses a Kalman filter. The state that is estimated is the ideal actuator input, $\vec{x}$. Because $\vec{x}$ cannot directly be observed, it is estimated using a state transition matrix that assumes very little change across iterations. Finally, the actuators of the deformable mirror are updated as same as the estimated state of the system.

Algorithm 1 Wavefront Correction using Kalman Filter

for $i = 0$ to $num\_iterations$ **do**
    # Forward path
    $wf\_dm \leftarrow deformable\_mirror.forward(wf)$
    $wf\_tar \leftarrow atms\_forward.forward(wf\_dm)$

    # Return path
    $wf\_atm \leftarrow atms\_backward.forward(wf\_tar)$
    $wf\_after\_dm \leftarrow deformable\_mirror.forward(wf\_atm)$
    $slopes \leftarrow shwfs.get\_slopes(wf\_after\_dm)$
    $slope\_err \leftarrow slopes - slopes\_ref$

    # Kalman filter prediction step
    $x\_pred \leftarrow A \cdot x\_est + B \cdot slope\_err$
    $P\_pred \leftarrow A \cdot P\_est \cdot A^T + Q$

    # Kalman gain calculation
    $K \leftarrow P\_pred \cdot C^T \cdot (C \cdot P\_pred \cdot C^T + R)^{-1}$

    # Update step
    $x\_est \leftarrow x\_pred + K \cdot (slope\_err - C \cdot x\_pred)$
    $P\_est \leftarrow (I - K \cdot C) \cdot P\_pred$

    # Apply correction to the deformable mirror
    $deformable\_mirror.actuators \leftarrow x\_est$
end for

3.3.4 Simulation Results Processing

After the simulation is run, there are many different variables which could be looked at, including across the iterations. For this reason, once a simulation is run, the results must be correctly processed to allow the efficient understanding. One of the main issues with this is

that the wavefront is composed of a two dimensional complex matrix, with both a phase and intensity that may be interesting.

For this work, the most valuable information is contained in the wf_{tar} matrix, which contains the wavefront at the target. Ultimately, the most important measure of success relates to the energy delivered to the target, hence the intensity is of main interest at the target. Further, one may also be interested in the phase term of $wf_{afterdm}$, as in theory the controller should work to minimize this term. Hence, the two main results are the intensity at the target plane and the phase right before the wavefront sensor.

Additionally, to aid in the efficient evaluation of a simulation, an algorithm must automatically compute a single numerical value can be extracted from the intensity profile at the target plane. This is the performance metric discussed in appendix H.5.1. The peak intensity is easy to determine: the maximum pixel value of the intensity image is taken. The power in the bucket is slightly more involved, taking 3 steps: (1) A mask of the 'bucket' needs to be created, with the correct dimensions, then (2) the intensity pattern is convolved with the mask to produce the convolved intensity pattern, and finally (3) the maximum value of the convolved intensity pattern corresponds to the PIB metric at that radius, centered at the pixel containing the maximum value. This algorithm ensures that the maximal power in the bucket is extracted from the intensity profile. Additionally, the PIB is normalized to the total amount of energy in the image, such that PIB is always expressed as a percentage of the total energy of the beam.

A visual representation of the process is shown in fig. 3.4 for 1 cm radius. The mask shown in fig. 3.4a was created by setting all spatial positions within 1 cm of the center to 1, otherwise 0. The mask was then convolved with a sample intensity pattern shown in fig. 3.4b to obtain fig. 3.4c. Figure 3.4c contains all the information about the PIB metric for 1 cm: the location of the brightest pixel corresponds to the center of the PIB bucket (circle), with the value of the pixel at that location equal to the PIB metric.

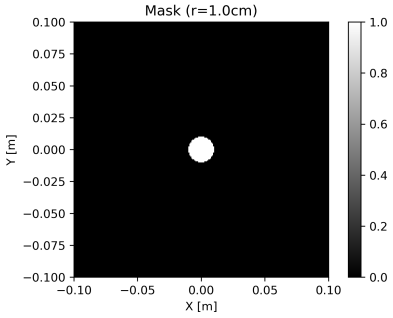

(a) 1 cm sample mask generated. Yellow indicates a value of 1, whereas purple 0.

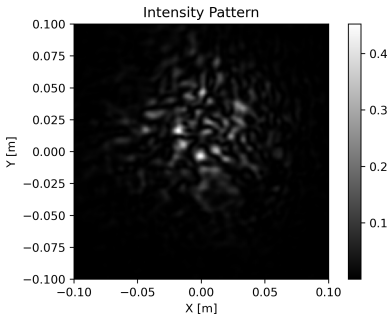

(b) Intensity pattern as generated by the simulation.

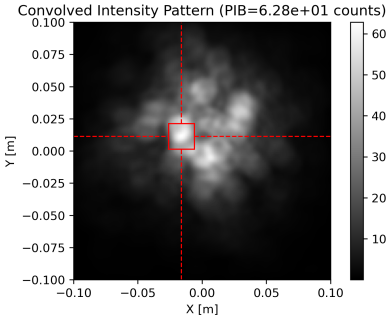

(c) Convolution of fig. 3.4a with fig. 3.4b.

Figure 3.4: Sample mask (fig. 3.4a) and intensity pattern (fig. 3.4b), which when convolved result in fig. 3.4c. This can be used to calculate the PIB given an intensity pattern and a PIB radius.

The reason for using a convolution-based approach rather than a fixed (x, y) for PIB is to ensure that the most favorable PIB is chosen in each case. To demonstrate, the same algorithm that computed fig. 3.4 was used on fig. 3.5a. Figure 3.5a is a generated intensity field that shows one pixel with a brightness of 100, and a 1 cm circle with a brightness of 10 for each pixel within that circle. This image is meant to exaggerate a very common scenario where a very bright pixel is present, however it may not be near other areas with high brightness. In this exaggerated case, the PIB metric for 1 cm shows that even though it does not contain the pixel with peak intensity, the circle with values 10 contains more energy than the single pixel. Hence,

the PIB algorithm is robust to intensity fluctuations as are to be expected with the turbulent atmosphere and speckle pattern.

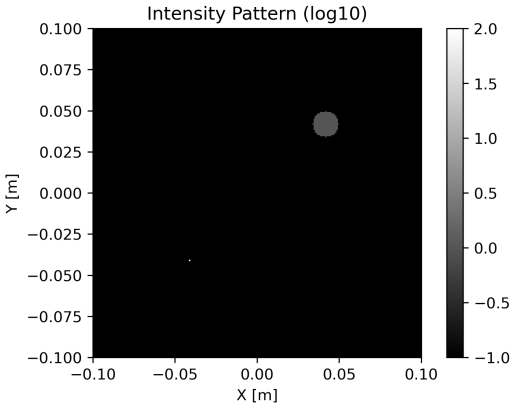

(a) $1\,\mathrm{cm}$ sample mask generated. Two regions are present: ont he bottom left a single pixel with value 100 and in the top right a $1\,\mathrm{cm}$ circle with brightness of 10. All other pixels are 0. The colorbar on the right shows the $\log_1 0$ value for brightness.

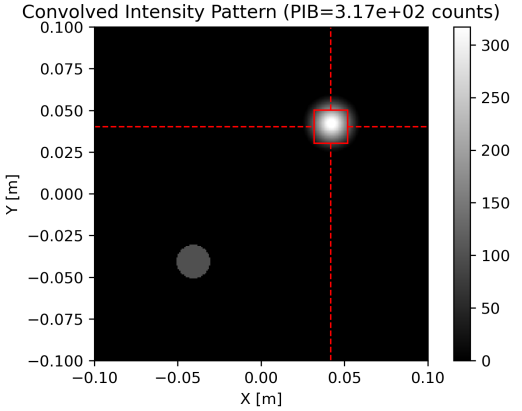

(b) Convolved intensity pattern, showing the top right $1\,\mathrm{cm}$ circle with higher energy than the bottom left.

Figure 3.5: Simulated intensity pattern and the PIB algorithm used to evaluate the performance. Despite containing the peak intensity, the single pixel does not have the maximum PIB for $1\,\mathrm{cm}$.

3.4 Summary

In this chapter, the overview of the simulation steps was provided, starting with the individual simulations in section 3.1. The individual simulations pave the way for the assembly simulations (section 3.2), which puts together the individual components which were tested before. Finally, the full simulation procedure is presented in section 3.3. In the following chapters, the results, discussions, and conclusions of the simulated environment are presented in chapters 4 to 6, which make use of the methods of this chapter to come to meaningful results.

Chapter 4

Results

Having presented TIL theory in previous chapters, it is now possible to begin generating some results. The results chapter is split into the two major parts identified in the methods section: the first section covers the single simulation results in section 4.1, where the individual test of the simulated components are performed. This is done to ensure an in-depth understanding of the results. In order to make more substantiated claims, these results are combined into statistical models from which conclusions can be drawn in the second major part, section 4.2. Finally, the results are summarized in section 4.3.

4.1 Single Simulation Results

The first step of this work was to ensure that the individual simulations could be run, and output the expected the results. The objective of this section is to present some key results, explain how they were made and the implications. The section begins with the S-I simulations in section 4.1.1, and continues to the TIL simulation in section 4.1.2. While not part of the core of this work, the verification of the atmospheric turbulence propagation software is performed in appendix I.

4.1.1 S-I Simulations

The first simulations to be carried out are the individual and subassembly simulations. This subsection begins with the surface roughness statistics simulations in section 4.1.1. Next, the speckle mitigation strategies are tested in section 4.1.1, respectively.

S-I.1: Surface Roughness Statistics

The first simulation carried out was generating a surface that conformed to the statistics outlined in section 2.1.5. As discussed, there are two ways of confirming the correct generation of the field: 2D autocorrelation, which uses the generation technique in reverse to generate the autocorrelation function, and the 1D autocorrelation, which uses a completely different method to calculate the autocorrelation, which is less efficient but does not rely on the technique used for generation. Use of both methods ensures the results are correct.

To increase the efficiency of subsequent simulations, 100 speckle fields were generated and saved, with statistical parameters $l_s = 2\,\text{mm}$ and $\sigma_s = 10\,\mu\text{m}$. These two values represent somewhat of

a worst case scenario: l_s is rather small compared to the evaluation grid of the optics (aperture in the order of 10 cm and grid spanning 20 cm across. This means that l_s is about 2 % of the total aperture. Hence, decreasing l_s would not significantly worsen performance. Increasing l_s would mean more gradual phase changes due to the surface and hence better performance of the overall system. A surface roughness of $\sigma_s = 10\,\mu m$ is also chosen to generate the 100 fields. This value is also somewhat of a worst case scenario, where the standard deviation of the surface roughness is significantly larger than the wavelength of light. This means that significant speckle is expected, and increasing σ_s above 10 µm would mean a significantly rough surface. Decreasing σ_s would be beneficial as scattered light is more coherent. Hence, both l_s and σ_s are chosen to represent a worst-case scenario.

Figure 4.1a shows the first of 100 surface roughness profiles generated, ξ_0. Alongside it, fig. 4.1b shows the 2D autocorrelation function recovered from just the field. A couple observations can be made, specifically about (1) the size of the bright spot of fig. 4.1b, (2) the size of the speckles of fig. 4.1a, and (3) the scale of the autocorrelation fig. 4.1b. These observations are qualitative, with quantitative reasoning presented after.

The first noteworthy observation of the two figures is the size of the bright spot of fig. 4.1b. From theory, it is expected that the bright spot is around 1 % of the length of the field. Strictly, at $r = l_s$, around $e^{-1/2} \approx 60\,\%$ of the maximal value occurs. With a grid size of 256 pixels would correspond to a bright spot diameter (from 60 % to 60 % across) of around 5 pixels. Although not immediately verifiable from the figure, this value seems plausible.

The second observation relates the spot size of fig. 4.1b to the speckle size of fig. 4.1a. It can by qualitatively assessed that the two sizes roughly correspond, as the most dark and bright spots of fig. 4.1a are about the size of the bright spot of fig. 4.1b. This would make sense, as values within the the l_s radius are auto-correlated with each other. Hence, again one could qualitatively assess that the statistics match to the expected value.

Thirdly, and finally, the scale of the autocorrelation fig. 4.1b is assessed. From the definition of the autocorrelation, the brightest spot in fig. 4.1b is expected at the center, with a value of $\exp(-0/l_s) = 1$. Indeed, the bright spot of fig. 4.1b is just below 1 presumably owing to the fact that this is drawn from a random sample. Moving outside $r = 0$, one can also observe the steep drop off from 1 into 0, however well outside $r = l_s$ the field seems to fluctuate around 0. Again, this is likely due to the fact that being a finite generated random sample there is some autocorrelation at these scales, albeit very small. This finding further confirms the consistencies between the rough surface statistics and the generated random surface.

However, to be certain that the statistics of the generated random surface correspond to the desired statistical parameters, a 1D analysis is carried out by plotting autocorrelation versus distance r. This is done in two ways: (1) flattening the 2D field into a 1D distance, (2) computing the product $\xi(\vec{r_1})\xi(\vec{r_2})$ for various $(\vec{r_1}, \vec{r_2})$ pairs. Figure 4.2a plots the reconstructed 1D autocorrelation function versus the the theoretical autocorrelation shown in eq. (2.14). Next, the directly computed products are added onto the figure, shown in fig. 4.2b. A more detailed explanation of the generation techniques and analysis follows.

Figure 4.2a shows the autocorrelation as shown in fig. 4.1b, however the distance from the center was first calculated for each pixel and then plotted in a scatter plot in Figure 4.2a. The red line shows a direct computation from eq. (2.14). Figure 4.2a shows that the method of generating ξ corresponds very well to the theory, even near $r = 0$ where there is less pixel data available. Indeed, the same phenomena of fig. 4.1b, namely the fluctuation near 0 when $r >> l_s$ can be observed. Figure 4.2a is yet another piece of evidence supporting the validity of the generated roughness profile.

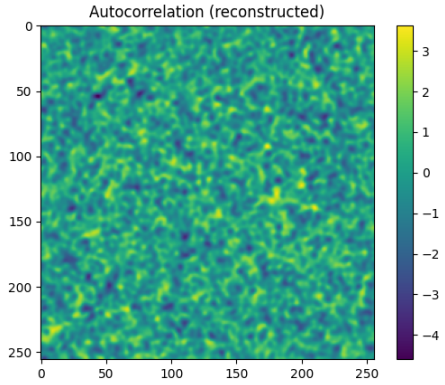

(a) Generated field ξ.

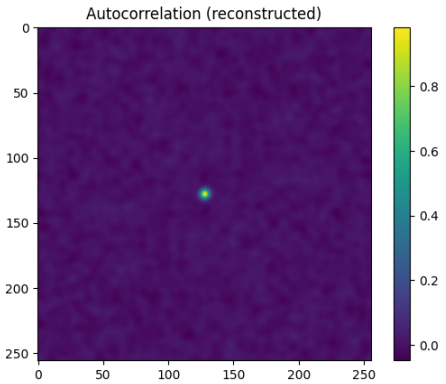

(b) Two dimensional autocorrelation of first field.

Figure 4.1: ξ and autocorrelation $< \xi(\vec{r_1})\xi(\vec{r_2}) >$ plotted in two dimensional plots.

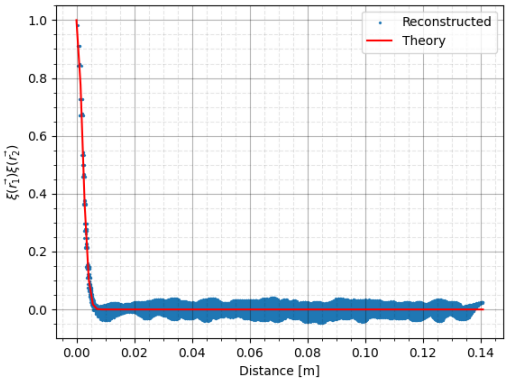

(a) One dimensional autocorrelation, with results from theory and reconstructed using the Wiener–Khinchin theorem. The field used is shown in fig. 4.1a.

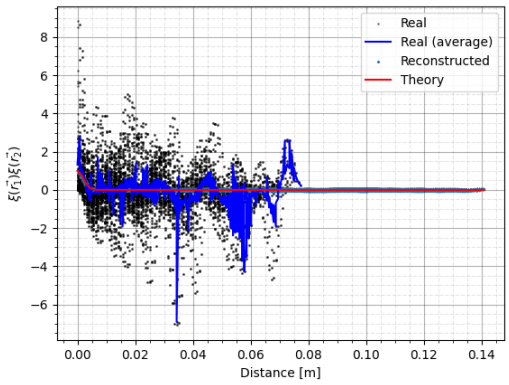

(b) One dimensional autocorrelation, with results from theory, from reconstruction using the Wiener–Khinchin theorem, and a direct computation (black dots). An average for each distance is shown in blue.

Figure 4.2: One dimensional autocorrelation plots extracted from fig. 4.1a. These figures are used to confirm that the generated field corresponds to the intended statistics.

To extend the analysis further, one final step is made: for the first 100×100 pixels the autocorrelation was directly computed. This resulted in 1×10^8 data points for $\xi(\vec{r_1})\xi(\vec{r_2})$ and $r = |\vec{r_2} - \vec{r_1}|$, which were plotted on top of fig. 4.2a to produce fig. 4.2b. Additionally, for each r $\xi(\vec{r_1})\xi(\vec{r_2})$ was averaged to obtain an estimate for $< \xi(\vec{r_1})\xi(\vec{r_2}) >$, shown in blue. The individual data points are shown in black. As can be observed from the figure, the directly computed data (real data) carries significant amplitude, and even when averaged there are significant spikes in the data. Nevertheless, given that the data was computed on a relatively small region of the field ξ (with around 40 times less data than physically available), and that may result in additional correlations, the variance is excused. Importantly, the averaged direct computation shows a clear relationship with the red theoretical line near $r = 0$, providing further evidence that the surface roughness profile was generated correctly.

The culmination of all results in S-I.1: Surface Roughness Statistics were intended to demonstrate that the generation of the surface roughness profiles indeed correspond to the statistical

parameters that were used to generate them. From observing the 2D autocorrelation function and comparing it to the field in fig. 4.1, to flattening the data into 1D in fig. 4.2, it was shown that expected statistics were re-obtained from the generated field, hence further providing evidence that roughness profiles were generated correctly. These results demonstrate that the rough surface profile generated does conform to the given statistics, and can be used in further work.

S-I.2: Speckle Mitigation: Speckle Average

Having successfully simulated roughness profile, it should be now possible to simulate a speckle pattern. As discussed previously, a speckle pattern occurs when a highly coherent source is incident on a rough surface, introducing large phase variations throughout the field of view. Speckle is disadvantageous for wavefront sensors, as it is impossible to distinguish between atmosphere-induced phase aberrations and the target's rough surface-induced phase aberrations. One way of mitigating the speckle pattern is to observe and average many speckle patterns, so as to 'average out' the surface-induced phase and re-obtain just the atmospheric phase.

In S-I.2: Speckle Mitigation: Speckle Average, the goal is to test and demonstrate the capability of averaging multiple surface roughness realizations. This is done by simulating a highly coherent source coincident with the target surface, and register the speckle pattern as observed some distance away, as shown in fig. 4.3. In this simulation, no atmosphere was simulated; hence the expected phase is almost planar near the center and some slight curving at the outer edges, as diffraction does take place. Hence, the phase distribution is expected to be spread out over the aperture.

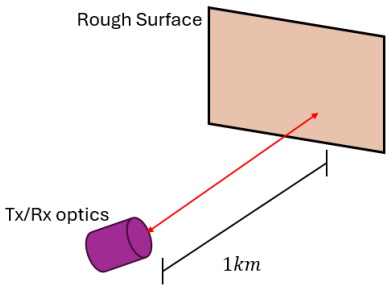

Figure 4.3: Setup used for testing the speckle averaging technique. The transmitting optics passes through a turbulence free medium 1 km towards the target rough surface. The rough surface scatters the incoming light, which is propagated back to the Rx optics.

The first result is the baseline: what would an ideal, mirror-like surface reflect? The result of the simulation is shown in fig. 4.4. When zero phase is induced by the target surface, one would expect a circular intensity profile at the receiving optics, with uniform phase difference near the center, with a slight curvature due to diffraction. Figure 4.4 shows the returned intensity pattern, a speckle-free intensity pattern. The histogram to the right shows the distribution of the different phases in the aperture. Importantly, there is a peak near -1.8 rad, with a little spread. This result is in line with what is expected: an almost uniform distribution small variation of phase values due to diffraction. The value -1.8 rad corresponds to the phase unwrapping algorithms, and the value itself is insignificant (as it corresponds to the piston mode of phase) .

In the ideal case, with perfect speckle mitigation one would expect to obtain a figure as shown in fig. 4.4. However, in practice it is impossible to average and obtain such perfect results.

The results of the Speckle Average simulation are shown in fig. 4.5, with the same structure as fig. 4.4.

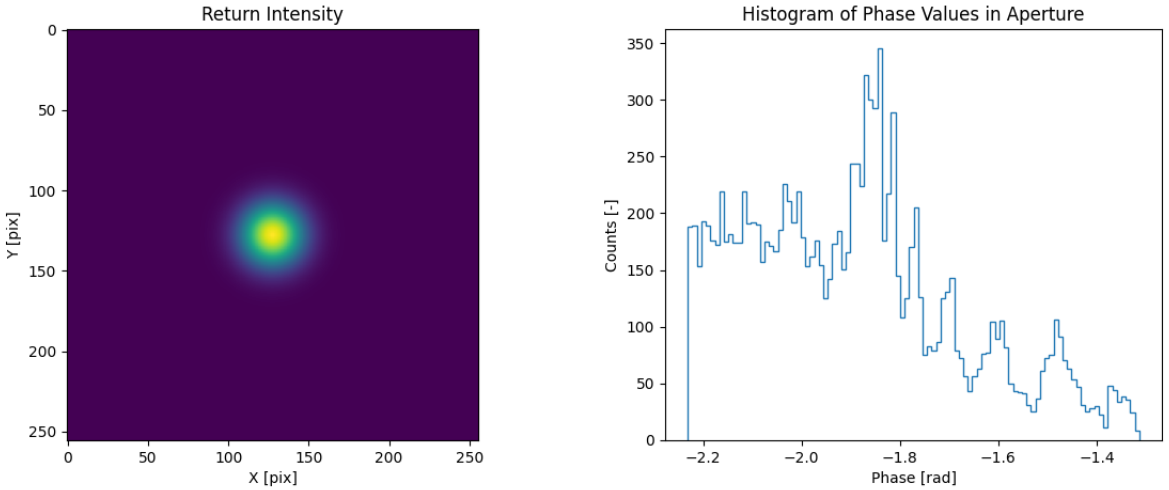

Figure 4.4: Mirror surface simulation. A coherent light source close to a mirror illuminates it, and the image is registered $1000\,\mathrm{m}$ away as the intensity profile on the left. The image on which the histogram is based on cannot be directly registered by a camera, however it can be retrieved from the simulation.

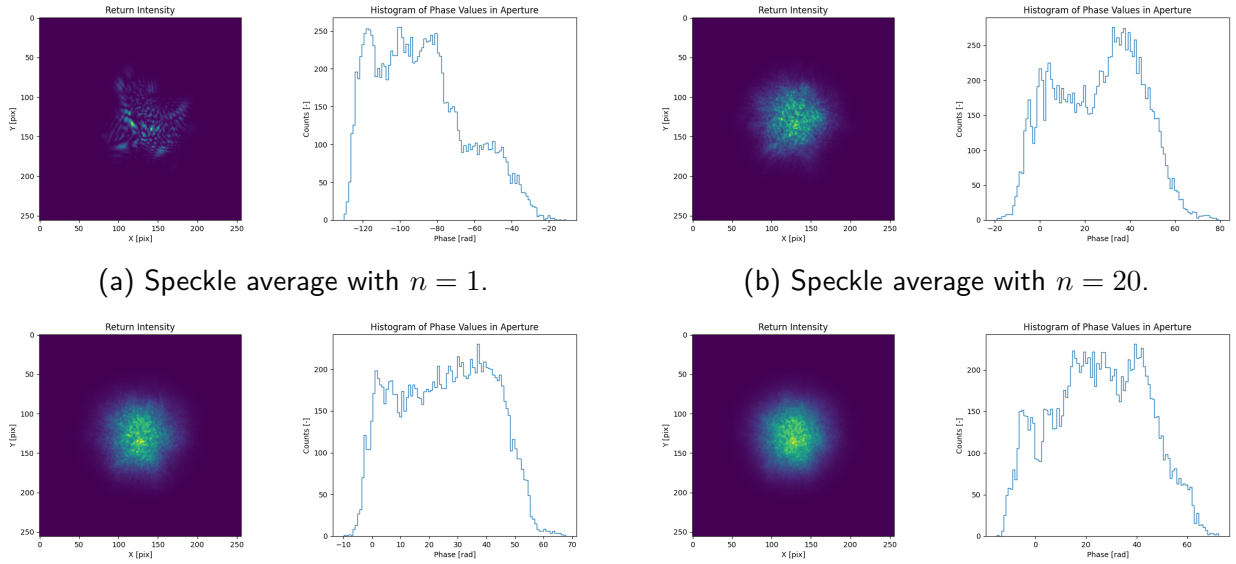

(a) Speckle average with $n = 1$.

(b) Speckle average with $n = 20$.

(c) Speckle average with $n = 50$.

(d) Speckle average with $n = 100$.

Figure 4.5: Speckle average for multiple surface roughness realizations. On the left the intensity plot shows the intensity at the receiving optics (counts), and on the right the distribution of unwrapped phase (counts). In all cases, a coherent light source is shone on the target surface and the intensity profile is registered by receiving optics $1\,\mathrm{km}$ away. The phase histograms are also shown on the right.

To begin the analysis, fig. 4.5a shows what the receiving optics $1000\,\mathrm{m}$ away would register if it just captured one sample of surface roughness in the absence of atmospheric turbulence. The intensity profile shows a clear speckle pattern, with bright spots scattered throughout the aperture. This result makes sense: the surface roughness was generated to be significant, and hence where one would desire a circular reflected beam there is just scatted light throughout.

The second Speckle Average simulation averages the propagation of 20 surface profiles in the receiving optics. The result is shown in fig. 4.5b. Immediately, it is possible to infer that the illumination's source is one noisy surface, rather than the multiple small glints that could be interpreted from fig. 4.5a. The effect of averaging is clear over just the one sample, however the question remains whether or not speckle averaging can produce better results.

The next two plots, fig. 4.5c and fig. 4.5d, show the effect of averaging using 50 and 100 surface roughness realizations. Through the intensity profiles one can observe a decreasing speckle influence, especially in the $n = 100$ figure.

The intention of simulation S-I.2: Speckle Mitigation: Speckle Average was to show that speckle averaging is a valid approach to mitigating speckle. The ideal, mirror-like reflection was simulated and the results are shown in fig. 4.4. When simulating speckle, the results were shown in fig. 4.5. The results point to speckle averaging as a viable technique in speckle mitigation, even with realtivly few

S-I.3: Speckle Mitigation: Polychromatic Speckle

The second mitigation technique investigated in this work is polychromatic speckle. In this simulation, rather than use multiple realizations of surface roughness, the idea is to use a wideband source that contains many different wavelengths, hence it is a polychromatic source. The source is defined by a center wavelength λ_c, equal to 634 nm for the simulation. The spectral width of the source is defined by w, which varies to mimic a polychromatic source, and could take values such as 1 nm, to mimic a highly coherent laser, or 10 nm, a more polychromatic laser. A single sample consists of 100 equidistant samples between $\lambda_c - w/2$ and $\lambda_c + w/2$, and then performs the averaging at the receiver optics.

The simulated environment in which the tests take place is the same for S-I.2: the source and receiver are placed 1000 m away from the rough surface of the target. Hence, the ideal case is still that shown in fig. 4.4 and the worst case as shown in fig. 4.5a. Figure 4.6 shows the results of polychromatic speckle mitigation.

To begin unpacking the results from fig. 4.6, first attention is paid to fig. 4.6a. Unsurprisingly, this first result paints a similar picture to that shown in fig. 4.5a, i.e. the worst case. The worst case is that of a highly coherent laser with the receiving optics just registering the first speckle realization. This confirms that the simulation is working as expected.

Figures 4.6b to 4.6d paint a better image for polychromatic speckle, however the results are still below that of taking 20 surface roughness realizations and averaging them, as shown in fig. 4.5b. This is likely due to the low surface roughness ($\sigma_s = 1\,\mu m$) chosen, which means that the polychromatic approach is insufficient to cancel out the phase induced by the rough surface.

To investigate this, an extreme example was chosen where $l_s = 5\,mm$, $\sigma_s = 100\,\mu m$ and w of 1 nm and 20 nm. The results are shown in fig. 4.7.

Simulation S-I.3: Speckle Mitigation: Polychromatic Speckle has shown the advantages, and limitations, of using a polychromatic source to mitigate speckle. While in practice the use of a polychromatic source is easy, requiring minimal hardware changes, it is not as effective as speckle average in mitigating speckle. Only in an exaggerated scenario, with significant surface roughness and shallow target angle, the speckle pattern begins to disapear. This result is somewhat expected, as [10] states four conditions for polychromatic speckle: (1) relatively small values of coherence length, (2) coherence length times wavenumber times field of view must be much less than 1, (3) absolute value of the gradient of the target shape must be less

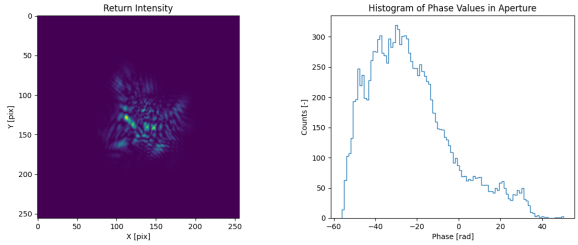

(a) Polychromatic speckle with w =1 nm

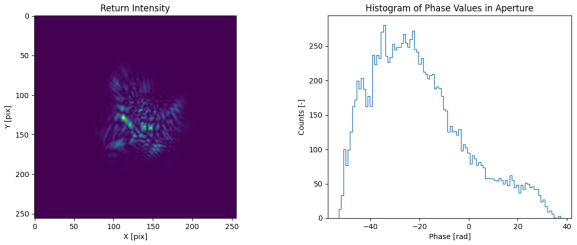

(b) Polychromatic speckle with w =10 nm

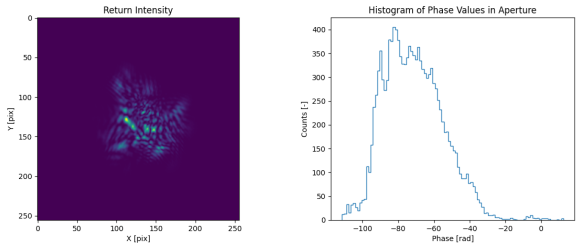

(c) Polychromatic speckle with w =20 nm

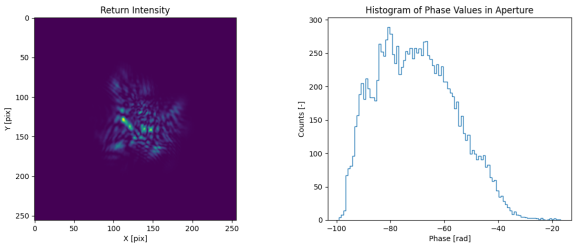

(d) Polychromatic speckle with w =50 nm

Figure 4.6: Results of polychromatic speckle mitigation for various FWHM w. In all cases, a light source is shone on the target surface and the intensity profile is registered by receiving optics 300 m away.

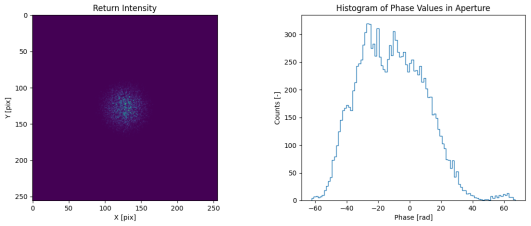

(a) Polychromatic speckle with w =1 nm.

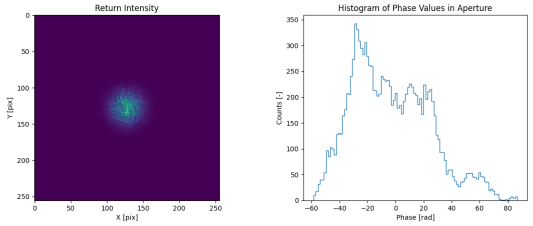

(b) Polychromatic speckle with w =20 nm.

Figure 4.7: Polychromatic speckle with modified surface roughness parameters. In all cases, a light source is shone on the target surface and the intensity profile is registered by receiving optics 300 m away.

than the field of view, and (4) $\theta_s \gg \Theta_0$, with $\theta_s \approx \sigma_s / l_s$ [10]. These conditions are investigated in the following:

(1) Coherence Length To estimate coherence length ℓ_Γ, Equation (12) of [23] is used (eq. (4.1)).

$$\begin{aligned} \ell_\Gamma &= \frac{2\ln 2}{\pi} \frac{\lambda_0^2}{\Delta\lambda} \\ &= \frac{2\ln 2}{\pi} \frac{634\,\mathrm{nm}^2}{20\,\mathrm{nm}} \\ &\approx 9.00\,\mu\mathrm{m} \end{aligned} \tag{4.1}$$

$\ell_\Gamma = 9.00\,\mu$m is rather small, especially when comparing the FWHM given by [17] (largest value being 0.66 nm). Hence, this condition is satisfied.

(2) Coherence Length × Wavenumber × Field of View In [10], this is mathematically described as shown in eq. (4.2), with k being the wavenumber and Θ_0 the field of view.

$$\ell_\Gamma k \Theta_0 \ll 1 \tag{4.2}$$

The remaining values can be estimated as shown below.

$$k = \frac{2\pi}{\lambda} = \frac{2\pi}{634\,\text{nm}} \approx 9.91\,\mu\text{m}^{-1} \tag{4.3}$$

$$\Theta_0 \approx \frac{W}{D} \approx \frac{20\,\text{cm}}{3\,\text{km}} \approx 66.7\,\mu\text{rad} \tag{4.4}$$

Substituting the derived values into the original equation:

$$\ell_\Gamma k \Theta_0 \approx 9\,\mu\text{m} \times 9.91\,\mu\text{m}^{-1} \times 66.7\,\mu\text{rad} \approx 5.95\,\text{mrad} \ll 1 \tag{4.5}$$

Hence this condition is also satisfied.

(3) Absolute value of Surface Gradient In [10], this is mathematically described as shown in eq. (4.6).

$$\left|\nabla_R S(R,t)\right| < \Theta_0 \tag{4.6}$$

Since $S(R,t) = 0$ for this work, it is clear that this condition is satisfied.

(4) Characteristic Angle for Target Surface Roughness The characteristic angle for target surface roughness, θ_s, can be estimated through eq. (4.7) and must be much greater than Θ_0.

$$\theta_s \approx \frac{\sigma_s}{l_s} = \frac{100\,\mu\text{m}}{5\,\text{mm}} \approx 20\,\text{mrad} \tag{4.7}$$

As $20\,\text{mrad} \gg 66.7\,\mu\text{rad}$, the condition for speckle mitigation condition is satisfied. Decreasing σ_s by a factor of 100 would mean $\theta_s \approx 200\,\mu\text{rad}$, which results in polychromatic speckle not working as it is no longer much greater than the field of view.

The first three conditions are all satisfied, even in the unmodified case. In the modified case, the fourth condition is satisfied, hence the speckle pattern is mitigated as shown in fig. 4.7. For a less favorable σ_s, polychromatic speckle would not work, as shown in fig. 4.6.

4.1.2 S-II Simulations

Having successfully tested the individual components, it is now possible to integrate the entire TIL system into one closed loop control simulation. First, a TIL system without scattering is simulated, which is done to have a baseline off which to compare the results when speckle effects are introduced. This is done in **S-II.1: Speckle-Free TIL** below. Upon success, the speckle effect is added in **S-II.2: Full TIL** further down.

S-II.1: Speckle-Free TIL

To properly assess the performance of the TIL system, first a speckle free idealization is simulated. Before the results are presented, first a few details of the implementation are shared. These include the deformable mirror configuration, the wavefront sensor configuration and the atmosphere. All these elements are already tested individually by [16], however they have parameter within themselves that are worth explaining.

Deformable Mirror The technology behind the deformable mirror simulated is not modeled, but rather, 20 influence functions are evenly spread out throughout the specified aperture. Applying a control input to the influence function results in a simulated displacement of the reflective surface, causing a phase shift. This input/output structure is analogous to the input/output structure of bimorph piezoelectric mirrors, consistent with the trade off. With knowledge of the wavefront aberrations present, one can theoretically translate them into control inputs to the deformable mirror to counteract the aberrations and obtain the desired wavefront.

Wavefront Sensor The wavefront sensor is the second element in the simulation. Its goal is to estimate the wavefront aberrations at each point on the aperture. As the chosen wavefront sensor was a Shack-Hartmann Wavefront Sensor, that is the wavefront that was modeled and simulated for the experiments. 10 lenslets are simulated along each axis, resulting in 100 lenslets overall. Note that some of these lenslets do not contribute any information as they are covered by the circular aperture of the sensor.

Atmosphere The final element in the speckle-free simulation is the atmosphere. The atmosphere was modeled using 6 layers, each with a C_N^2 of $1 \times 10^{-13}\,\mathrm{m}^{-2/3}$, and along a distance of 3 km, in accordance with the strong turbulence case presented in [17].

Calibration The calibration of the TIL system is an essential aspect of the entire process, as it ensures that the system knows what aberrations can be compensated for by the correct deformable mirror actuators. As the system must be calibrated before an engagement, and irrespective of target distance, the calibration procedure used was that of typical adaptive optics: a guide star (point source) is placed in front of the receiving optics in an aberration-free environment. Then, by actuating each actuator on the deformable mirror and recording the changes in the sensor, one can invert the relationship and use that for robust control.

Simulation Overview The path the beam takes throughout the simulation was shown in fig. 3.3.

No Controller Simulation The first TIL test was one without any adaptive optics. This represents the case of using a fixed, flat mirror (e.g. non-active DM) to in the place of the one used to compensate for the atmosphere. The results are displayed in fig. 4.8.

Figure 4.8 displays what a relatively simple HEL system could achieve: without any closed loop control, one could focus a large amount of energy onto a target far away. The key metrics of fig. 4.8 are the peak intensity and PIB figures above the right figure. The PIB is primarily considered in this work. The PIB of the no controller simulation is 34.4 % for 1 cm radius. This value will be later compared to the controller PIB.

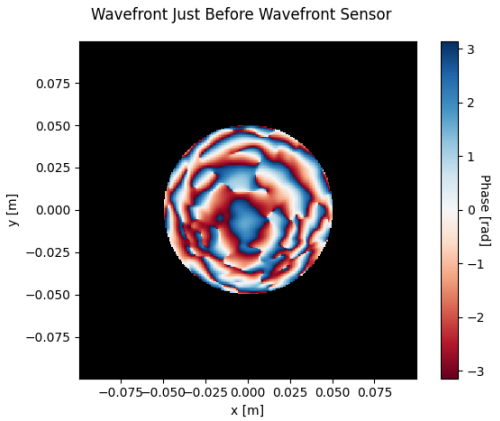

(a) Phase of wavefront just before wavefront sensor.

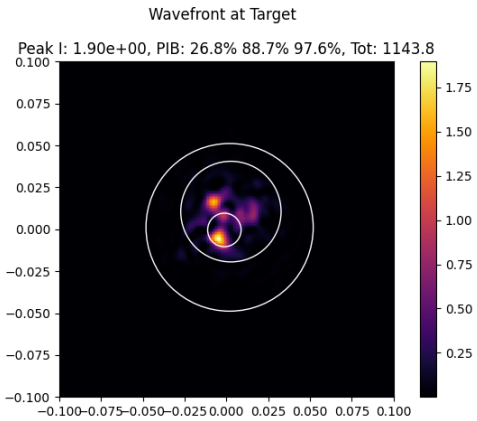

(b) Intensity at target plane.

Figure 4.8: Simulation result at $3\,\mathrm{km}$ with strong turbulence ($1 \times 10^{-13}\,\mathrm{m}^{-2/3}$), no speckle, and no controller. The left figure shows the phase just before the wavefront sensor in radians. The right image shows peak intensity and can be read after Peak I in the title. The white circles indicate the PIB radii $1\,\mathrm{cm}$, $3\,\mathrm{cm}$, $5\,\mathrm{cm}$, also corresponding to the percentage values in the title after PIB. The aperture size is $10\,\mathrm{cm}$. The beam diameter at aperture is around $6\,\mathrm{cm}$.

Controller Simulation The result of the simulation carried out at a distance of $3\,\mathrm{km}$ is shown in fig. 4.9. The controller was run for 100 iterations in a steady atmosphere, emulating the most challenging environment: the beginning of an engagement with a target.

Immediately, the advantages of using a closed loop control are apparent: from a PIB figure of $34.4\,\%$ to $45.0\,\%$ represents an increase in PIB of $11.6\,\%$. The intensity difference between the two is also formidable, from 0.88 to 1.50, a $71\,\%$ increase in peak intensity. These results on a single realization of the atmosphere are promising, however now it is possible to investigate the effect of speckle and speckle mitigation of the simulation.

S-II.2: Full TIL

Having individually tested and confirmed the viability of the optical setup as well as the speckle mitigation strategy, it is now possible to simulate the two and draw some conclusions on the expected performance of the simulated TIL system. This section covers the simulations done with speckle, using both speckle average and polychromatic speckle.

Speckle Average The speckle average speckle mitigation strategy seemed to be the most promising of the two strategies, based off previous results. In this simulation, the optical setup is kept the same, as well as the target distance, and rather than mirror-like reflection off the target, the surface roughness of the target introduces phase terms which present themselves as speckle on the receiving optics. As in section 4.1.1, several values of n are tested to understand their effect on the performance of the system. Figure 4.10 shows the results for n 1 and 100.

The first observation from fig. 4.10 is the characteristic speckle pattern shown in the top left for both figures. It qualitatively corresponds to the results when a coherent source was shone onto a target surface and the scattered wave was observed some distance away. Despite the $n = 1$ case achieving a slightly higher PIB than the no speckle case, using 100 roughness realizations results in even more energy onto the target. With fig. 4.10, it appears that speckle averaging is a very useful technique to mitigate speckle, and can be used in a practical scenario.

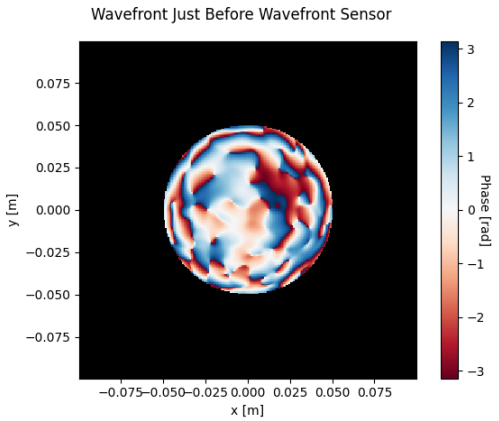

(a) Phase of wavefront just before wavefront sensor.

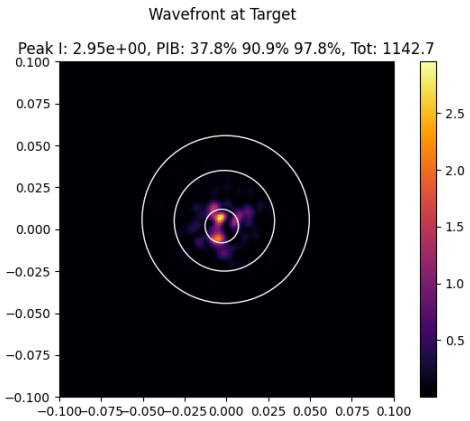

(b) Intensity at target plane.

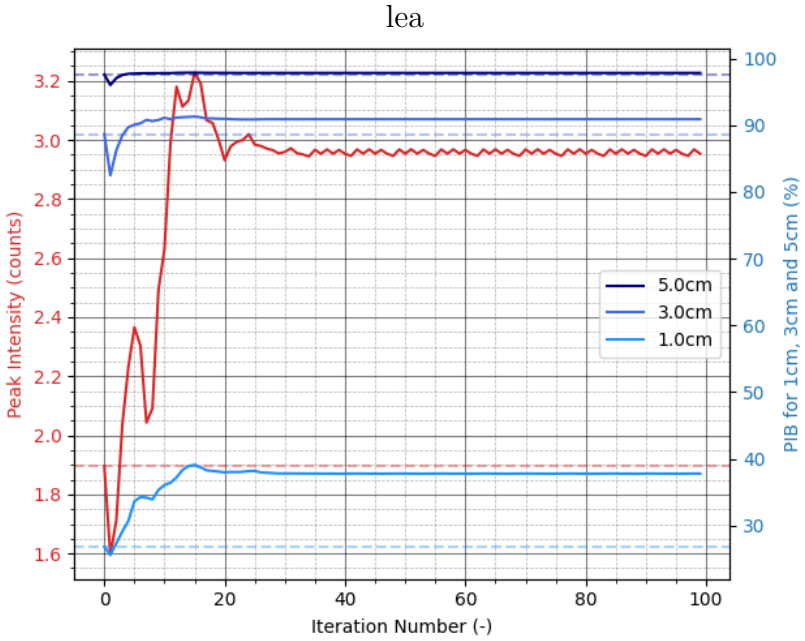

(c) Key metrics plotted against iteration number. After about 20 iterations the metrics stabalize and reach their approximate final values. The dashed lines of the same color represent the initial (no AO) values.

Figure 4.9: Simulation result at 3 km with strong turbulence and no speckle. The plots display the same information as in fig. 4.8 in the case a controller is applied. Additionally, fig. 4.9c shows the key metrics as they evolved with iteration number.

Polychromatic Speckle Section 4.1.1, specifically speckle mitigation using the polychromatic approach showed the limitations of using a polychromatic approach to mitigate speckle. Still, however, the polychromatic mitigation technique was simulated in the hopes that perhaps it can provide some degree of correction. The results are shown in fig. 4.11 for $w = 1\,\text{nm}$ and $50\,\text{nm}$.

Figures 4.11a and 4.11b highlight the weakness of using a polychromatic source to mitigate speckle. Between fig. 4.11a and fig. 4.11b, there is very little difference, both in the received speckle after the atmosphere and in the phase after the deformable mirror. This lack of difference translates to virtually no differences in target energy delivery, with peak intensity and power in a bucket at nearly the same values. Despite using a highly incoherent source, speckle

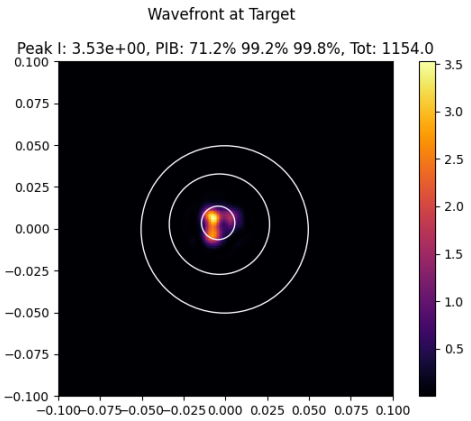

(a) Results for $n = 1$.

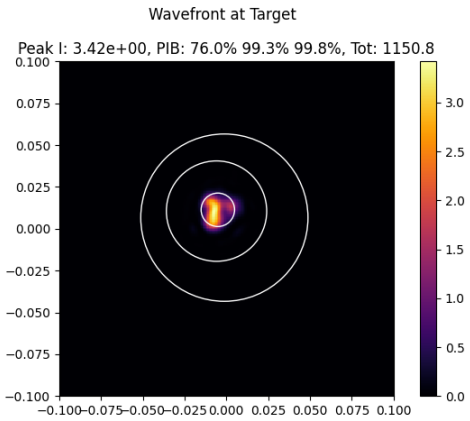

(b) Results for $n = 100$.

Figure 4.10: Speckle average TIL full simulation at a distance $3\,\mathrm{km}$ away for different n. 6 phase screens were used with a turbulence strength of $C_N^2 = 1 \times 10^{-13}\,\mathrm{m}^{-2/3}$. The surface roughness characteristics are l_s =2 mm and σ_s =10 µm.

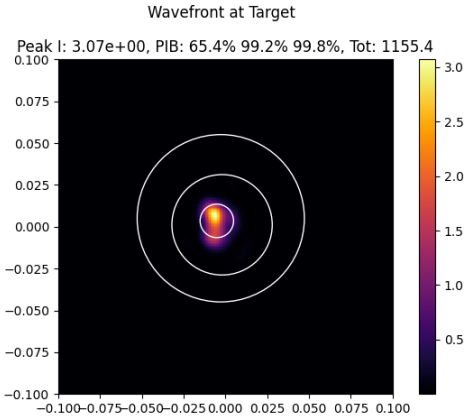

(a) Results for polychromatic speckle with w =1 nm.

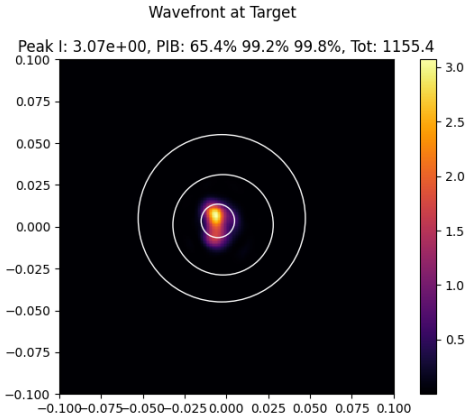

(b) Results for polychromatic speckle with w =50 nm.

Figure 4.11: Polychromatic speckle TIL full simulation at a distance $3\,\mathrm{km}$ away for different w. 6 phase screens were used with a turbulence strength of $C_N^2 = 1 \times 10^{-13}\,\mathrm{m}^{-2/3}$. The surface roughness characteristics are l_s =2 mm and σ_s =10 µm.

is still unable to be mitigated, and hence the optimization is suboptimal. As was the conclusion of S-I.3, the polychromatic approach could possibly be used with another approach to be acceptable, and on its own does not mitigate speckle.

4.2 Results Culminated

Up until this point the results have been from individual tests. While the individual tests provide an in-depth understanding of the functioning behind the technique, they cannot provide enough evidence on their own to make a claim that would be able to address the research question of this work.

To address this, many individual results are culminated in this section, which will be used later on to substantiate claims related to the research questions. The analysis begins by assessing the effect turbulence strength has on key performance indicators in section 4.2.1. Next, a similar

analysis is performed on target distance in section 4.2.2.

4.2.1 Turbulence Strength

To being the analysis, many simulations were run with varying turbulence strengths, C_n^2, and target distances, with and without controller.

First, the culmination of the results for varying C_n^2 is shown in fig. 4.12. The peak intensity and PIB for radii 5 cm, 3 cm and 1 cm are shown with and without control for 100 independent simulations in a box-and-whisker plot.

The comparison between the controlled and uncontrolled cases demonstrate some trends in the behavior of peak intensity and PIB across different turbulence strengths. In general, the controlled case consistently displays similar variability and improved performance compared to the uncontrolled scenario. Additionally, as the turbulence strength decreases, the benefits of the controlled system begin to show, indicating that the effect of control is most effective after a certain atmospheric turbulence strength.

For peak intensity, the uncontrolled case shows a significant spread at high turbulence levels ($C_n^2 = 1.0 \times 10^{-13}\,\mathrm{m}^{-2/3}$), indicating strong fluctuations in intensity in the target plane. The controlled case seems to not improve the average performance nor decrease the variability. As turbulence decreases, both cases converge toward a higher average peak intensity, while the variability remains high. This suggests that control is particularly effective in mitigating the effects of medium to little turbulence.

A similar trend is observed in PIB at 1 cm, where the controlled case consistently achieves higher values than the uncontrolled case in the medium to little turbulence cases. The difference is most pronounced at the medium and low turbulence levels, where the average difference is greatest. The controlled case maintains a higher average PIB, indicating that the system is effectively reducing intensity distortions caused by atmospheric turbulence.

For PIB at 3 cm and 5 cm, the controlled case shows a slight advantage over the uncontrolled case. In each case, the variability is around the same. By the lowest turbulence level ($C_n^2 = 1.0 \times 10^{-15}\,\mathrm{m}^{-2/3}$), the results for both cases are nearly identical, suggesting that in near-ideal atmospheric conditions, the benefits of control are less pronounced.

These observations seem to suggest that the benefits to adaptive optics only apply to the turbulence strengths around $1.0 \times 10^{-14}\,\mathrm{m}^{-2/3}$ and $1.0 \times 10^{-15}\,\mathrm{m}^{-2/3}$, while at $1.0 \times 10^{-13}\,\mathrm{m}^{-2/3}$ the benefits observed are insignificant. To investigate this, fig. 4.13 was created as a comparison of the target plane intensity pattern for two different turbulence cases.

Figure 4.13 provides some insight into the results plotted in fig. 4.12. Firstly, it shows that at $C_n^2 = 1.0 \times 10^{-13}\,\mathrm{m}^{-2/3}$ the returned pattern clearly resembles a turbulence-based speckle pattern as the adaptive optics is incapable of correcting the wavefront. This singular result can be used to argue why all results at this turbulence level in fig. 4.12 are so similar: the adaptive optics system is doing almost nothing. On the other hand, fig. 4.13b shows that no longer is there a speckle pattern, but rather an almost optimized focused point. This result suggest that between $C_n^2 = 1.0 \times 10^{-13}\,\mathrm{m}^{-2/3}$ and $1.0 \times 10^{-14}\,\mathrm{m}^{-2/3}$, there is a point (for these simulation settings) where adaptive optics begins to offer some improvement.

Overall, the results indicate that the implementation of control mechanisms significantly improves the performance of the system, especially at medium to low turbulence levels. At high turbulence levels, such as $C_n^2 = 1.0 \times 10^{-13}\,\mathrm{m}^{-2/3}$, this configuration of adaptive optics is incapable of correcting the target plane intensity. The overall trends suggest that control remains

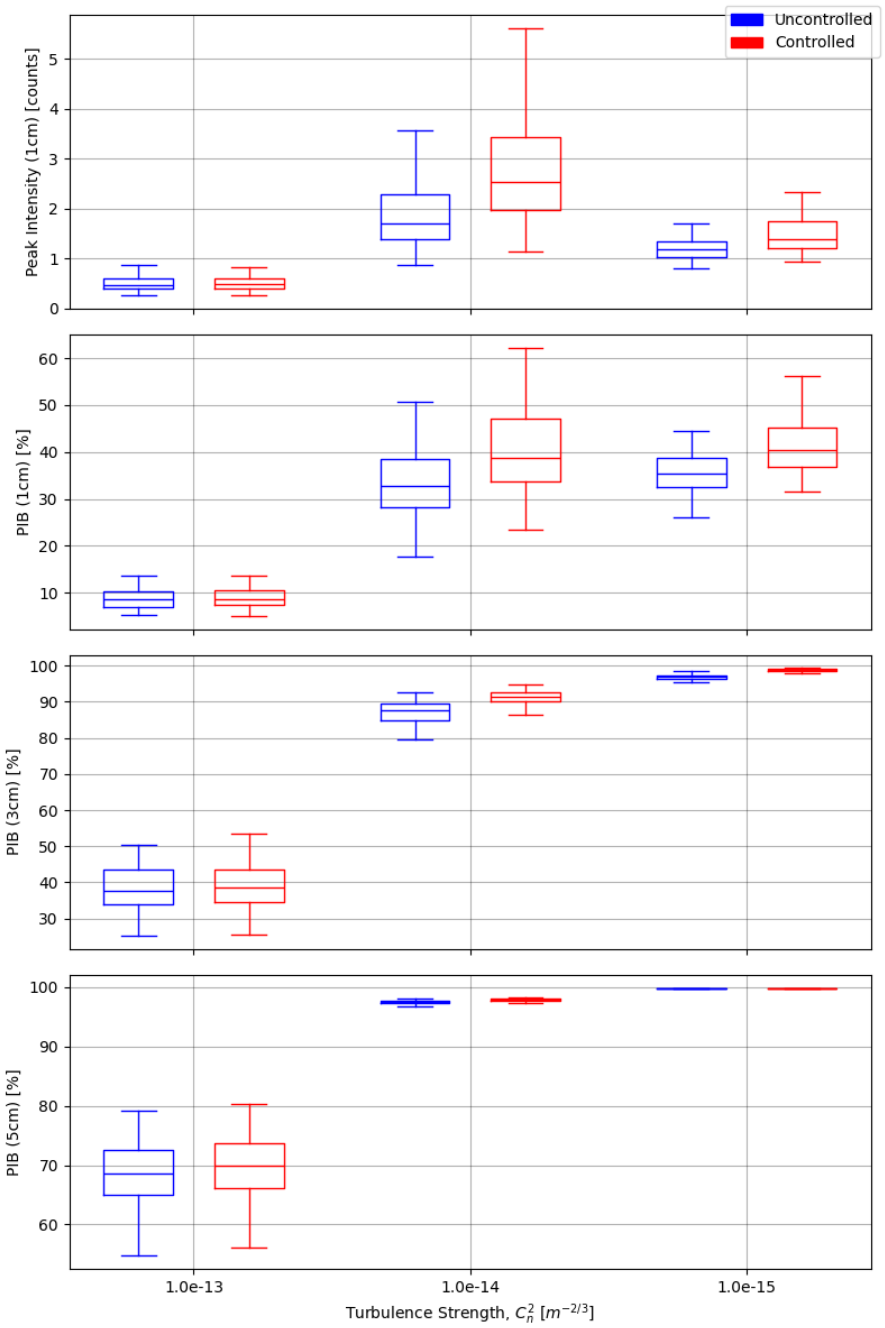

Figure 4.12: Power in the bucket versus atmospheric turbulence strength when using a controller for various radii. Data was generated by simulating a target at a distance of $3\,\mathrm{km}$ away at varying turbulence conditions and recording the power in the bucket for $r = 1\,\mathrm{cm}$, $3\,\mathrm{cm}$ and $5\,\mathrm{cm}$. This process was repeated 20 times to generate a dataset for both the controlled and non-controlled scenarios. Note that both the controlled and non-controlled were evaluated at the same atmospheric turbulence strength and are not overlapping for visual purposes.

beneficial in improving optical performance under realistic atmospheric conditions, and never decreases the performance of the system.

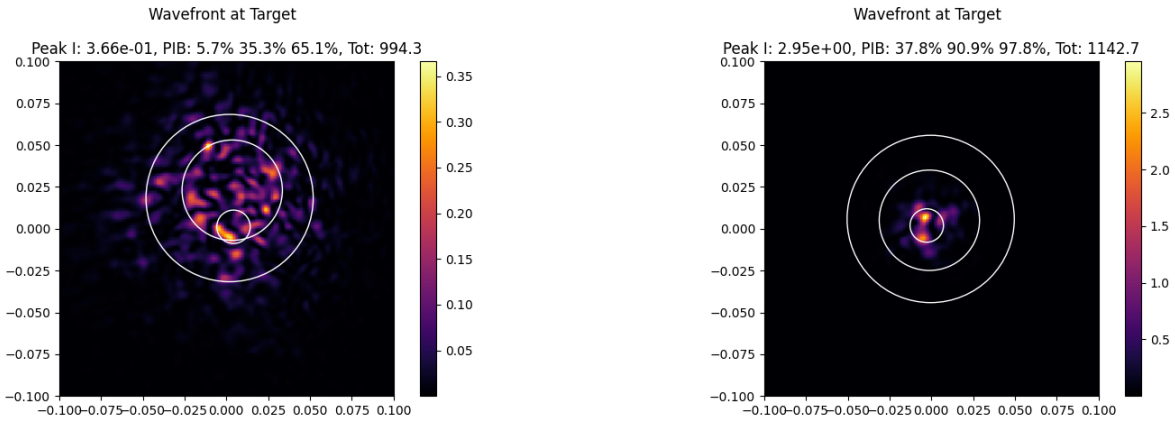

(a) Result for $C_n^2 = 1.0 \times 10^{-13}\,\mathrm{m}^{-2/3}$.

(b) Result for $C_n^2 = 1.0 \times 10^{-14}\,\mathrm{m}^{-2/3}$.

Figure 4.13: Target plane intensity profiles for two runs at two different turbulence strengths levels. Both images are after 100 iterations at a target 3 km away.

4.2.2 Target Distance

The second result culmination performed relates to target distance, and is shown in fig. 4.14. Similar to section 4.2.1, the PIB for radii 5 cm, 3 cm, and 1 cm are shown with and without control, for 100 independent samples. The results demonstrate that control generally improves performance by reducing variability and enhancing stability, particularly at longer distances where turbulence effects are stronger.

At shorter distances (≤ 1 km), the benefits of the controlled case over the uncontrolled case are very visible, especially in peak intensity and PIB (1 cm). As the distance increases to 3 km and beyond, the uncontrolled and controlled cases begin to behave in a similar way, with similar means and variability.

For PIB at 3 cm and 5 cm at the shorter target distances (≤ 1 km), the mean values are almost the same, and the variability is greatly reduced. This could be due to the fact that at such short distances the effect of turbulence is greatly decreased. This is doubly true as the return path involves propagating twice the distance.

The final observation is that at 5 cm PIB, the ≤ 1 km cases show that approximately 100 % of the energy is still contained. The 3 km case is already known to be unstable (with the 5 km case even more turublent), and indeed the difference between it and the ≤ 1 km cases is very clear.

4.2.3 Aggregate Results

To investigate the relationship between the strength of the turbulence and the distance from the target, fig. 4.15 was generated.

Figure 4.15 shows results for various different atmospheric turbulence strengths and target distances. The figure shows the percentage increase over the non-adaptive optics case for peak intensity, and PIB for 1 cm, 3 cm and 5 cm. Four observations are made of this plot.

The figures indicate that the percentage difference is always greater than 0 %, meaning that the AO system will always be at least as good as the non-adaptive optics case. However, the changes are most noticeable at a certain combination of distance and atmospheric turbulence strength.

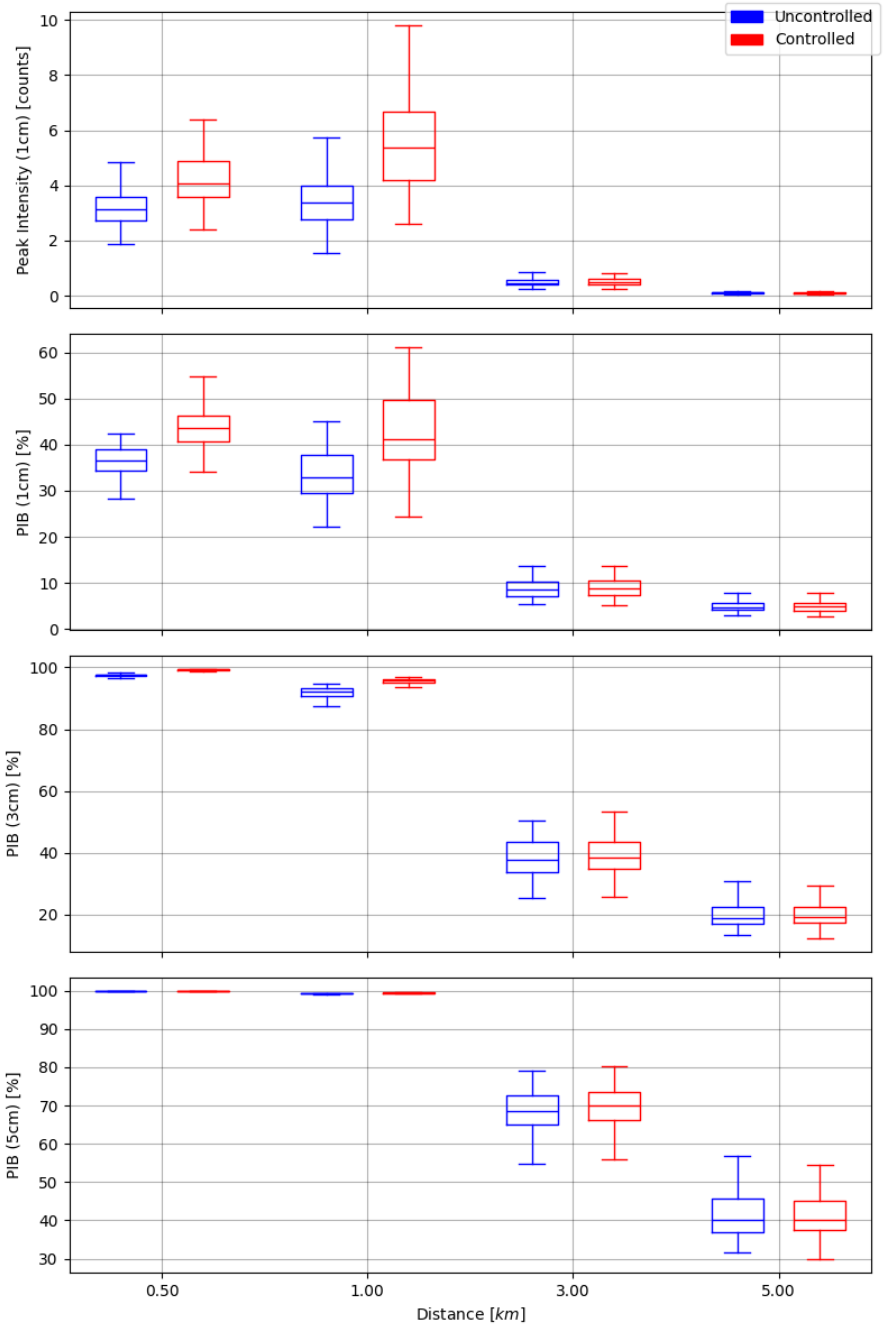

Figure 4.14: Power in the bucket versus distance when using a controller for various radii. Data was generated by simulating a turbulence strength of $1 \times 10^{-13}\,\mathrm{m}^{-2/3}$ away at varying distances and recording the power in the bucket for $r = 1\,\mathrm{cm}$, $3\,\mathrm{cm}$ and $5\,\mathrm{cm}$. This process was repeated 20 times to generate a dataset for both the controlled and non-controlled scenarios. Note that both the controlled and non-controlled were evaluated at the same distances and are not overlapping for visual purposes.

At shorter distances, the percentage difference shows only a marginal increase, indicating that the AO system can only slightly improve the performance of the system over non-adaptive optics systems at these short ranges. This suggests that over shorter ranges, the laser beam retains much of its original characteristics, with only slight distortions. The relatively low standard deviation in these regions supports the idea that fluctuations remain controlled until turbulence has had sufficient distance to influence the beam significantly.

Beyond a certain distance, the percentage difference approaches zero, meaning that both adap-

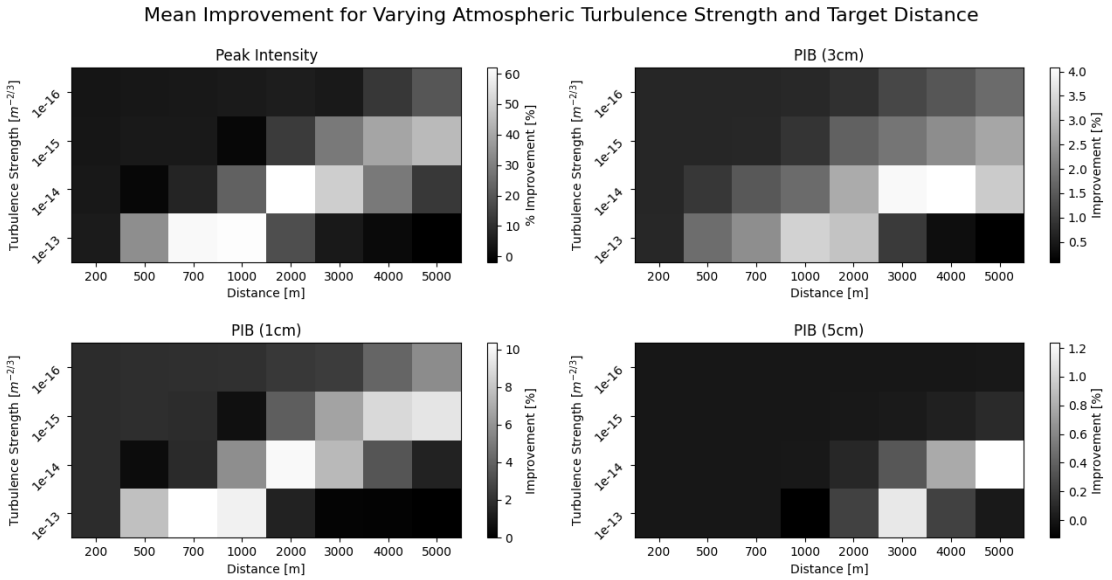

Figure 4.15: Aggregate Results for 4 unique turbulence strengths and 8 unique distances. For each subfigure, the horizontal axis is distance in meters, and the vertical axis is the turbulence strength in $\mathrm{m}^{-2/3}$. The results are based off 100 simulations for each data point. The left figures show the mean improvement when using no adaptive optics versus using a deformable mirror. The right 4 figures show the standard deviations.

tive optics and non-adaptive optics systems would suffer from too much turbulence. This implies that after a critical range, the beam can no longer be correctly steered. This is particularly evident in the PIB at 1 cm versus 5 cm, where the effects of AO are "pushed back" in the case of 5 cm.

The figure for PIB at 1 cm and peak intensity are nearly identical, suggesting that at small radii, PIB behaves similarly to peak intensity. This would seem to make sense as the limit of PIB as the radius decreases is equivalent to peak intensity.

Another key observation, related to the similarities between the PIB at 1 cm and peak intensity, are the two dark spots at $(500\,\mathrm{m},\ 1\times10^{-14}\,\mathrm{m}^{-2/3})$ and $(1000\,\mathrm{m},\ 1\times10^{-15}\,\mathrm{m}^{-2/3})$. These values are pitch black, suggesting a 0 % mean improvement. Due to the 100 samples used, it is likely that this result is not due to noise, but some other phenomena. It was suspected that the gains used for the Kalman controller had a great influence on these results, hence to investigate a different set of Kalman gains were used, which were tweaked to perform better for PIB at 1 cm and peak intensity. The results are shown in fig. 4.16 on a limited domain due to computational reasons.

Figure 4.16 suggests a very interesting observation: the gain of the system can be chosen to optimize for a certain metric: peak intensity or PIB for 1 cm, 3 cm, or 5 cm. In this case, peak intensity was optimized, with little regard for PIB 1 cm and no regard for the other two metrics. As an example, one simulation result is shown in fig. 4.17 for $(1\,\mathrm{km},\ 1\times10^{-14}\,\mathrm{m}^{-2/3})$.

It is shown that for the two extraneous cases shown in fig. 4.15, they no longer lack performance, demonstrating that the effect is due to the chosen Kalman gains.

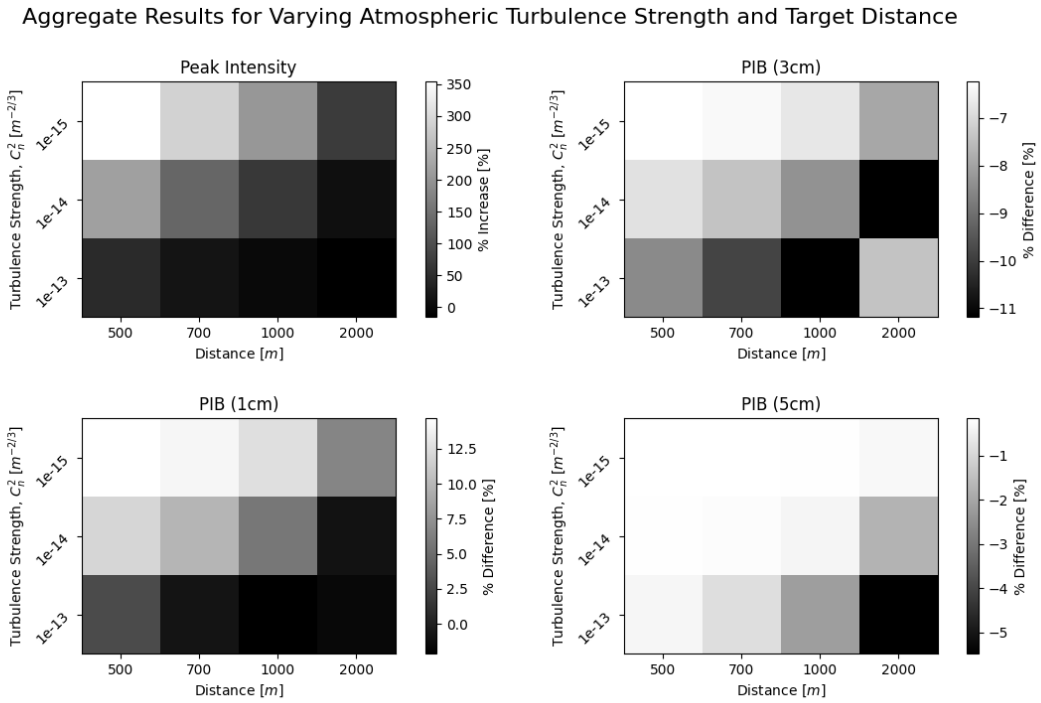

Figure 4.16: Results with modified Kalman gains. The gains were optimized on a single simulation ($1\,\mathrm{km}$, $1 \times 10^{-14}\,\mathrm{m}^{-2/3}$) to perform much better in the metrics of PIB at $1\,\mathrm{cm}$ and peak intensity.

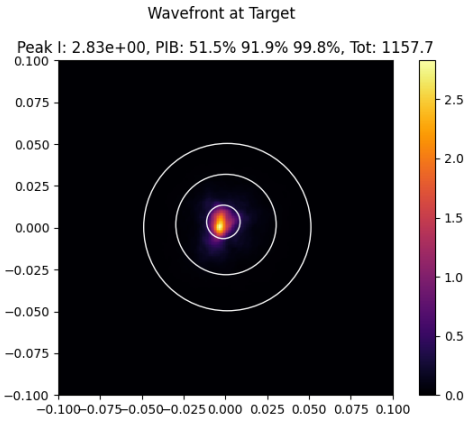

(a) Wavefront intensity image at the target after 100 iterations.

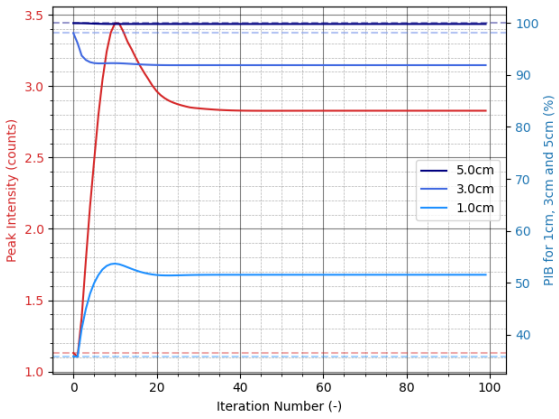

(b) Key metrics against iteration number with the modified gains.

Figure 4.17: Results for ($1\,\mathrm{km}$, $1 \times 10^{-14}\,\mathrm{m}^{-2/3}$) with the modified Kalman gains, optimized using a single simulation with parameters ($1\,\mathrm{km}$, $1 \times 10^{-14}\,\mathrm{m}^{-2/3}$). The mean improvement for peak intensity is $150\,\%$, and for PIB from smallest to largest is $15.7\,\%$, $-6.2\,\%$ and $-0.15\,\%$. This result complements fig. 4.16 as it shows how optimizing for the Kalman gains is imperative to the key performance metric chosen.

4.3 Results Overview

In the results chapter, the experimental work outlining the simulation work was carried out and presented. Section 4.1 saw the results of the individual simulations, and a detailed account into the conditions under which the results were created in. Section 4.2 combined many of the simulations carried out previously into a statistical results, which can be used to substantiate claims about the performance of the system. Having presented the main results of this work, it

is now possible to move onto the discussion of results, where an in-depth analysis on the result will be performed.

Chapter 5

Discussion

The results presented in the previous section provide valuable insights into the effectiveness of TIL systems. This discussion interprets these findings in the context of the previous findings shown in chapter 4 (section 5.1), the interpretation of results (section 5.2), then the comparison with existing literature (section 5.3). Finally, the broader implications are explored in section 5.4, as well as the limitations in section 5.5.

5.1 General Overview of Findings

From chapter 4, some key findings were identified. The goal of this section is to summarize the findings such that they can be used in the further sections of chapter 5.

Firstly, chapter 4 saw the successful realization of surface roughness profiles according to the statistical parameters that were used to generate them. Further, it was shown that these surface roughness profiles, when a shone with highly coherent light and imaged some distance away, result in a characteristic speckle pattern.

The speckle pattern was shown to significantly degrade the integrity of the intensity pattern as observed from receiving optics. While for a mirror-like reflection the receiving optics may register a circular pattern, when scattered off a rough surface the return wave had many bright and dark spots. These bright and dark spots were a result of the phase induced by the rough surface, which were theorized to significantly degrade the performance of an adaptive optics system which relied on the phase measurements of the return wave. To mitigate the effect of speckle, two approaches were studied: speckle average and polychromatic speckle.

Speckle average relies on realizing many surface roughnesses (for example by jittering a beam very fast across a target surface or having a fast spinning target). Through an example plot, it was shown that as one increases the number of speckle realizations used for the average from $n = 1$, 20, 50 and 100, the phase of the incoming light was approaching that of the mirror-like reflection. For more favorable surface roughness conditions the mirror-like reflection was almost completely restored.

The second method investigated to mitigate speckle was polychromatic speckle, which relies on using a source that emits various wavelengths of light. The qualitative principle behind this method is that many wavelengths of light interact with the surface differently, resulting in an averaging similar to speckle average. However, this method proved to be ineffective at flattening the incoming wavefront, even in extremely favorable conditions.

With the individual tests successfully completed, full TIL simulations could now be carried

out. The first such simulation eliminated the scattered wave off the target and replaced it with a mirror-like reflection. This was done to ensure the optical setup would work in speckle-free conditions. It was shown that the setup could indeed optimize the projected pattern on the target, with very favorable results for PIB with radius 1 cm.

From here, the effect of speckle was added, along with the mitigation strategies investigated before. As with the individual tests, speckle averaged performed very well, whereas polychromatic speckle could not be used to optimize the projection.

Next, similar simulations were run for varying atmospheric strengths and target distances, resulting in a culmination of results that could be used to substantiate claims about the effectiveness of the system. It was found that while a certain set of gains could be used throughout the operational envelope, these could be optimized for the specific conditions to achieve even greater performance.

5.2 Results Interpretation

The first key results are those that are culminated across 100 independent simulations for varying atmospheric strengths and target distances. These provide some idea behind the variances between results, as well as the type of performance that can be expected, outside of individual results which may misrepresent the performance. Figures 4.12 and 4.14 show the culminated results, with the overall interpretation coming from fig. 4.15.

Figure 4.15 shows that the main benefit of using AO is only at longer target distances or more turbulent atmospheres, up to a point. All the figures seem to be separated into two: the turbulence induced speckle, the turbulence induced speckle-free regions, and the intermediate region between them. Starting with the turbulence induced speckle-free region, found in the upper left area of the figures, uncontrolled and controlled systems perform about the same. The qualitative explanation is that turbulence is too weak to greatly affect the beam over the relatively short distance, and hence the projected pattern on the target surface is about the same as if it were propagated through free space. In this situation, the only benefit an AO system can provide is focusing the light, however given the aperture to target distance ratio the effect is minimal. This can be observed with the PIB 1 cm and 3 cm, where the speckle free region is not quite black but rather a shade of gray, indicating an above 0 % performance increase.

The second of three regions is the speckle region, found in the bottom right of each figure. Qualitatively, the turbulence over the target distance is too strong (too string to preperly measure the phase) even for the adaptive optics system, and an optimization of the beam projection on the target is impossible to realize. This presents itself as a 0 % performance increase between a controlled and uncontrolled system, the phase measurement is insufficient to properly apply corrections.

The third and final region is the transition between speckle-free and speckle. In the previous two regions, the benefits of AO were marginal, however it is in this region that AO benefits the most. Qualitatively, the turbulence degrades the beam to an extent where phase measurements can be performed, and with an inverse operation from the deformable mirror some additional performance can be extracted. This result points to one major outcome of this work: AO is mostly useful to extend the operation range and/or atmospheric strength of a system, with marginal benefits inside the existing operating conditions and no benefits well beyond existing operating conditions.

Additionally, the PIB for 3 cm and 5 cm see the main improvements at longer distances the the PIB for 1 cm. This points to the final interpretation: the AO control algorithm developed in this work serves to further contain the beam spot size on the target.

Overall, these results suggest that the control mechanism is particularly effective for concentrating power in smaller buckets and at longer distances with medium turbulence. The improvement in PIB indicates that controlled beams maintain their power concentration more effectively, leading to potential benefits in applications in TIL systems. For larger radii, the effect of control is less pronounced, and delayed, as the beam waist radius is already constrained by the atmospheric turbulence strength and its natural diffraction characteristics.

The second key finding relates to the effectiveness of the speckle mitigation strategies investigated. For speckle average, it was shown to be very effective, achieving a relatively unperturbed wavefront after only 20 samples, with increasing effectiveness at 50 and 100 samples. Practically, this means that in the time it takes the Shack-Hartmann sensor to integrate the incoming light, the beam must move across the surface a distance at most 20 times the width of the beam. Alternatively, a rotating target would have to rotate a distance equal to 20 times the beam-width. Polychromatic speckle, on the other hand, was shown to be very limited in its applications, as it could only be used to effectively mitigate speckle under very forgiving circumstances.

With these two main results, one can now compare them with the existing literature.

5.3 Comparison with Existing Literature

The equations used to model the rough surface, and by extension the speckle average and polychromatic approach, stem from [10] and [17]. In his paper [10], Voronstov assumes that enough roughness realizations have taken place such that the registered wavefront has no phase contributions from the rough surface. In [17], a similar assumption is made: the source is polychromatic enough such that speckle is mitigated. In this work, their work was extended in two major ways: (1) the two methods were seamlessly combined, such that a combination of the two is possible, and comparisons between the two can also be made, and (2) the question of exactly how much speckle mitigation is necessary was answered by simulating speckle and the speckle mitigation strategy. These two extensions made it possible to simulate the performance of a real-world system in many different applications, and obtain statistical parameters as to its expected performance.

In [10], it was derived that for enough speckle mitigation, the surface roughness profile is washed out and then the typical adaptive optics algorithms can be used to optimize the outgoing beam. This conclusion can also be observed in the results of this work, where for more speckle averages the return pattern resembles more and more the one with perfect reflection.

It was shown that speckle mitigation can only be achieved with speckle average, and that polychromatic speckle can only be used when a sufficiently large surface roughness is present.

With respect to the aggreate results, [17] suggested that in the case of strong turbulence, ”degradation is mostly due to scintillation”. Indeed, some scintillation can be observed in fig. 4.13a.

5.4 Broader Implications

The broader implications of this work stem mainly from the simulation carried out, which imply that speckle mitigation should mainly be done through speckle averaging rather than polychromatic speckle.

Additionally, the planning of the practical work that can be used to validate the simulation work was presented, and can be used in future works.

5.5 Limitations

The limitation of this work relates to the amount of power incident on the target. While in this work it was assumed that the atmosphere did not absorb any energy from the laser beam, in reality this is not the case. To make the situation even worse, the absorbed energy by the atmosphere creates a thermal blooming effect, further spreading out the energy of the beam. This causes a positive feedback look, as the AO system seeks to minimize the spot size, further heating the local atmosphere. The result is that the real performance of the TIL system is not simulated, but rather just the atmospheric effects. In reality, a power budget would need to be used to estimate the power incident on the target, including the effect of atmospheric absorption and degraded performance due to thermal blooming. Another limitation is tip/tilt modeling as the software used to propagate the wavefront centers the results automatically.

Chapter 6

Conclusions

This study investigated TIL systems, aiming to solve the double path and speckle problems associated with it. Through simulations and planned practical testing, this work derived several important conclusions, shared in section 6.1. Further, the main achievements and limitations of this work are highlighted in section 6.2. Finally, concluding remarks are shared in section 6.3.

6.1 Summary of Key Findings

Target in the loop systems suffer from engineering problems that need to be solved before deploying such a system in the field. In this work, two main problems, namely the double path problem and the speckle problem were addressed. The double path problem was explained by arranging the optical equipment in such a configuration so as to provide a source of light which can be used as the signal for the control loop. However, due to practical limitations this light is projected over a large area, introducing a speckle pattern at the receiving optics. This work investigated the speckle pattern, as well as ways to mitigate its effect, through the use of two method: speckle average and polychromatic speckle.

It was found that speckle average consistently outperformed polychromatic speckle in every configuration, leading to the finding that speckle averaging should be prioritized over polychromatic speckle.

With this finding, the double path problem was address by simulating a system where speckle average is used. This system was simulated under a variety of turbulence strengths and distances, resulting in consistent increases in power in the bucket metrics over the system with no active control. The turbulence strengths simulated included extreme turbulence ($C_n^2 = 1 \times 10^{-13}\,\mathrm{m}^{-2/3}$) to light turbulence ($C_n^2 = 1 \times 10^{-15}\,\mathrm{m}^{-2/3}$), and the distances simulated ranged from 500 m to 5 km. In all cases, compensating for the atmospheric turbulence resulted in increases in energy delivery.

With these findings, it was also found that in optimizing for a slightly focused return wave, the metrics PIB 1 cm, 3 cm, 5 cm and peak intensity improved after being controlled. Further improvements could be achieved at the cost of other metrics by tweaking the gains.

The key findings of this work are:

- The two speckle mitigation strategies investigated function, with speckle averaging working for all the tested surface roughness realizations, but taking comparatively longer and requiring specialized hardware in practice. Polychromatic speckle was shown to work

only with more favorable surface roughness parameters, however is relatively much easier to implement.

- The intensity pattern on the target could be optimized, with general purpose gains that work throughout the operational envelope. The gains could be further optimized for a specific metric (in this work peak intensity), to achieve much higher performance at the cost of the other key metrics.
- All of the above could be simulated, facilitating the study and simulation of a concept project.

6.2 Achievements and Limitations

The key achievements of this work lie in working towards addressing two big issues of TIL systems. Typically a TIL application would suffer from speckle patterns and difficulties optimizing the energy delivery due to the double path problem, however in this work steps were taken to mitigate these effects. The result was a simulation tool capable of simulating TIL systems and estimating their performance for a variety of turbulence conditions and target distances. Additionally, this simulation tool could be used in further design to inform key decisions such as the need for adaptive optics, High-Energy Laser power output and optical configuration.

Along with these achievements, several limitations were also identified. Due to simulating exclusively atmospheric effects, no attempt was made to simulate other power losses throughout the laser beams trajectory. Specifically, as was shown when preparing for the practical work, major losses are incurred as a result of scattering off the target, which are not modeled in this work. Additionally, thermal blooming, which results by the local absorption of the laser by the atmosphere, causing a diverging lens effect along the laser beam path. These are all limitations which should be addressed in future iterations of this work.

6.3 Concluding Remarks

To conclude, in this work a TIL simulation tool was developed, along with the plan for practical validation step. The performance of a specific TIL system was shown, along with its performance across various atmospheric turbulence and target distances.

Chapter 7

Recommendations

Based on the findings of this study, several recommendations can be made to improve TIL systems. These suggestions aim to address key challenges and guide future work in this field. These recommendations are provided in the order in which the author of this work recommends they should be achieved. The recommendations for future work are as follows: (1) investigation of fusion of speckle average and polychromatic speckle, (2) general controller optimization, (3) carry out practical work to validate the simulations of this work, and (4) add an SLM to simulate jitter to the practical setup.

Speckle Mitigation Fusion The first recommendation to continue this work is to investigate the desirableness of fusing different speckle mitigation strategies. As was seen in section 4.1.1, the polychromatic approach on its own is rather weak, but some desirable effects are visible. It it were to be combined with the much more capable speckle averaging technique, more impressive results could be attained. For now, however, the speckle averaging technique is much preferred over polychromatic speckle.

Validation of Results Significant effort was put in this work to verify the results obtained. However, for a proper analysis this would should be accompanied by some practical work that validates the results of this work. For example, it is surprising to the author that the polychromatic approach performed so poorly, and practical work could confirm or disprove the efficacy of this method.

SLM for Beam Jitter Lastly, as an extension to the validation of the results an SLM to induce beam jitter would be very useful. Rather than realize different roughness profiles by simulating uncorrelated fields, or spinning a target about one of its axes, using an SLM would fully decouple this work from assumptions made of the target, as enough speckle realizations would be obtained so as to not need a spinning target.

Bibliography

[1] Mariano Zafra et al. *How drone combat in Ukraine is Changing Warfare*. Ed. by Mike Collett-White, Simon Scarr, and Pravin Char. Mar. 2024. URL: https://www.reuters.com/graphics/UKRAINE-CRISIS/DRONES/dwpkeyjwkpm/.

[2] Iain Boyd. *Israel's Iron Dome Air Defense System Works well – here's how Hamas got around it*. Sept. 2023. URL: https://www.colorado.edu/asmagazine/2023/10/13/israels-iron-dome-air-defense-system-works-well-heres-how-hamas-got-around-it.

[3] Defence Science and Technology Laboratory and Ministry of Defence. *Advanced future military laser achieves UK first — gov.uk*. https://www.gov.uk/government/news/advanced-future-military-laser-achieves-uk-first. [Accessed 20-11-2024].

[4] Mikhail A. Vorontsov. In: *Advanced Optical Technologies* 2.5-6 (2013), pp. 369–395. DOI: doi:10.1515/aot-2013-0054. URL: https://doi.org/10.1515/aot-2013-0054.

[5] Patrick S. Salter and Martin J. Booth. "Adaptive optics in laser processing". In: *Light: Science & Applications* 8.1 (Nov. 2019), p. 110. ISSN: 2047-7538. DOI: 10.1038/s41377-019-0215-1. URL: https://doi.org/10.1038/s41377-019-0215-1.

[6] R. L. Carman and P. L. Kelley. "TIME DEPENDENCE IN THE THERMAL BLOOMING OF LASER BEAMS". In: *Applied Physics Letters* 12.8 (Apr. 1968), pp. 241–243. ISSN: 0003-6951. DOI: 10.1063/1.1651973. eprint: https://pubs.aip.org/aip/apl/article-pdf/12/8/241/18421018/241_1_online.pdf. URL: https://doi.org/10.1063/1.1651973.

[7] R. Paschotta. *Thermal Blooming*. RP Photonics Encyclopedia. DOI: 10.61835/tka. URL: https://www.rp-photonics.com/thermal_blooming.html (visited on 11/21/2024).

[8] Roderik A. Overzier et al. "High energy lasers as an anti-satellite capability". In: *High-Power Lasers and Technologies for Optical Countermeasures II*. Ed. by Willy L. Bohn, Marc Eichhorn, and Gareth D. Lewis. Vol. 13201. International Society for Optics and Photonics. SPIE, 2024, 132010G. DOI: 10.1117/12.3031520. URL: https://doi.org/10.1117/12.3031520.

[9] Karen M. Hampson et al. "Adaptive optics for high-resolution imaging". In: *Nature Reviews Methods Primers* 1.1 (Oct. 2021), p. 68. ISSN: 2662-8449. DOI: 10.1038/s43586-021-00066-7. URL: https://doi.org/10.1038/s43586-021-00066-7.

[10] Mikhail A. Vorontsov and Valeriy Kolosov. "Target-in-the-loop beam control: basic considerations for analysis and wave-front sensing". In: *J. Opt. Soc. Am. A* 22.1 (Jan. 2005), pp. 126–141. DOI: 10.1364/JOSAA.22.000126. URL: https://opg.optica.org/josaa/abstract.cfm?URI=josaa-22-1-126.

[11] Wikimedia Commons. *File:Laser speckle.jpg — Wikimedia Commons, the free media repository*. [Online; accessed 22-November-2024]. 2023. URL: %5Curl%7Bhttps://commons.wikimedia.org/w/index.php?title=File:Laser_speckle.jpg&oldid=831871312%7D.

[12] R. Paschotta. *Gaussian Beams*. RP Photonics Encyclopedia. Available online at https://www.rp-photonics.com/gaussian_beams.html. DOI: 10.61835/mla"title="Thislinkwillreloadthecurrentpage.. URL: https://www.rp-photonics.com/gaussian_beams.html (visited on 11/22/2024).

[13] K. S. Shaik. *Atmospheric Propagation Effects Relevant to Optical Communications*. Tech. rep. 20040191353. Washington, D.C.: Jet Propulsion Laboratory, National Aeronautics and Space Administration (NASA), Apr. 1988, pp. 1–21. URL: https://ntrs.nasa.gov/api/citations/20040191353/downloads/20040191353.pdf.

[14] Larry C. Andrews. *Field Guide to Atmospheric Optics, Second Edition*. SPIE, Jan. 2019. ISBN: 9781510619388. DOI: 10.1117/3.2318080. URL: http://dx.doi.org/10.1117/3.2318080.

[15] Norbert Wiener. "Generalized harmonic analysis". In: *Acta Mathematica* 55.none (1930), pp. 117–258. DOI: 10.1007/BF02546511. URL: https://doi.org/10.1007/BF02546511.

[16] E. H. Por et al. "High Contrast Imaging for Python (HCIPy): an open-source adaptive optics and coronagraph simulator". In: *Adaptive Optics Systems VI*. Vol. 10703. Proc. SPIE. 2018. DOI: 10.1117/12.2314407. URL: https://doi.org/10.1117/12.2314407.

[17] Noah R. Van Zandt and Mark F. Spencer. "Improved adaptive-optics performance using polychromatic speckle mitigation". In: *Appl. Opt.* 59.4 (Feb. 2020), pp. 1071–1081. DOI: 10.1364/AO.379972. URL: https://opg.optica.org/ao/abstract.cfm?URI=ao-59-4-1071.

[18] Jakub Kraciuk et al. *SC42065 Adaptive Optics Design Project*. Tech. rep. Author's own work. Mekelweg 5, 2628 CD Delft: Delft University of Technology, June 2023.

[19] Matthias T. Banet and Mark F. Spencer. "Compensated-beacon adaptive optics using least-squares phase reconstruction". In: *Opt. Express* 28.24 (Nov. 2020), pp. 36902–36914. DOI: 10.1364/OE.409134. URL: https://opg.optica.org/oe/abstract.cfm?URI=oe-28-24-36902.

[20] David A. Harville. "The Moore-Penrose Inverse". In: *Matrix Algebra From a Statistician's Perspective*. New York, NY: Springer New York, 1997, pp. 497–519. ISBN: 978-0-387-22677-4. DOI: 10.1007/0-387-22677-X_20. URL: https://doi.org/10.1007/0-387-22677-X_20.

[21] Anurag Singh Tomar et al. "Chapter 9 - Comprehensive survey of face super-resolution techniques". In: *Digital Image Enhancement and Reconstruction*. Ed. by Shyam Singh Rajput et al. Hybrid Computational Intelligence for Pattern Analysis. Academic Press, 2023, pp. 213–233. ISBN: 978-0-323-98370-9. DOI: https://doi.org/10.1016/B978-0-32-398370-9.00016-0. URL: https://www.sciencedirect.com/science/article/pii/B9780323983709000160.

[22] *AOK5(/M), AOK8(/M), AOKWT1(/M) Adaptive Optics Kits User Guide*. Rev. D. Thorlabs. Aug. 2024.

[23] Ceyhun Akcay, Pascale Parrein, and Jannick P. Rolland. "Estimation of longitudinal resolution in optical coherence imaging". In: *Applied Optics* 41.25 (Sept. 2002), p. 5256. ISSN: 1539-4522. DOI: 10.1364/ao.41.005256. URL: http://dx.doi.org/10.1364/AO.41.005256.

[24] Jeroen Kalkman. *Wavefield propagation*. 2022. URL: https://qiweb.tudelft.nl/aoi/wavefieldpropagation/wavefieldpropagation.html.

[25] Frits Zernike. "Beugungstheorie des schneidenver-fahrens und seiner verbesserten form, der phasenkontrastmethode". In: *Physica* 1.7 (1934), pp. 689–704. ISSN: 0031-8914. DOI: https://doi.org/10.1016/S0031-8914(34)80259-5. URL: https://www.sciencedirect.com/science/article/pii/S0031891434802595.

[26] Frits Beukers. *Gauss' hypergeometric function*. Oct. 2009. URL: https://webspace.science.uu.nl/~beuke106/GaussHF.pdf.

[27] *ANSI Z80.28-2017 Ophthalmics - Methods of Reporting Optical Aberrations of Eyes*. Standard. Alexandria, Virginia, United States of America: American National Standard for Ophthalmics, Aug. 2017.

[28] Jacopo Antonello, Théo Lequesne, and adriaanph. *zernike*. https://github.com/jacopoantonello/zernike. 2024.

[29] François Roddier. "Curvature sensing and compensation: a new concept in adaptive optics". In: *Appl. Opt.* 27.7 (Apr. 1988), pp. 1223–1225. DOI: 10.1364/AO.27.001223. URL: https://opg.optica.org/ao/abstract.cfm?URI=ao-27-7-1223.

[30] Roberto Ragazzoni. "Pupil plane wavefront sensing with an oscillating prism". In: *Journal of Modern Optics* 43.2 (Feb. 1996), pp. 289–293. ISSN: 1362-3044. DOI: 10.1080/09500349608232742. URL: http://dx.doi.org/10.1080/09500349608232742.

[31] Phasics. *QWLSI: a powerful wavefront technology overcoming limitations of Shack-Hartmann wavefront sensors and Fizeau interferometers*. URL: https://www.phasics.com/en/company/unique-wavefront-sensing-technology/.

[32] N. T. Pelekanos et al. "All-optical spatial light modulator with megahertz modulation rates". In: *Optics Letters* 20.20 (Oct. 1995), p. 2099. ISSN: 1539-4794. DOI: 10.1364/ol.20.002099. URL: http://dx.doi.org/10.1364/OL.20.002099.

[33] Thorlabs. *Adaptive Optics kits*. 2024. URL: https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=3208.

[34] H. Ottevaere and H. Thienpont. "OPTICAL MICROLENSES". In: *Encyclopedia of Modern Optics*. Ed. by Robert D. Guenther. Oxford: Elsevier, 2005, pp. 21–43. ISBN: 978-0-12-369395-2. DOI: https://doi.org/10.1016/B0-12-369395-0/00923-4. URL: https://www.sciencedirect.com/science/article/pii/B0123693950009234.

[35] Tony Lacey. "Tutorial: the Kalman Filter 11.1 Introduction 11.2 Mean Squared Error". In: 1998. URL: https://api.semanticscholar.org/CorpusID:48574040.

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my advisors, Dr.ir. Rudolf Saathof and Sven van Binsbergen for their invaluable guidance, support, and patience throughout the course of this work. Their insight and mentorship have shaped the direction of this thesis and my growth as an aspiring engineer.

On a personal note, I want to thank my friends and family for their unwavering encouragement and understanding.

To everyone who contributed to this journey in ways big or small — thank you.

Appendix A

Integration of Intensity for Power

The derivation begins with eq. (2.1).

$$I(r,z) = \frac{P}{\pi(w(z)^2/2)} \exp\left(-2\frac{r^2}{w(z)^2}\right) \tag{A.1}$$

To find the total power within a circular area of radius R, integrate over the area:

$$P(R,z) = \int_0^{2\pi} \int_0^R I(r,z)\, r\, dr\, d\theta \tag{A.2}$$

Substitute the expression for $I(r,z)$:

$$P(R,z) = \int_0^{2\pi} \int_0^R \frac{P}{\pi(w(z)^2/2)} \exp\left(-2\frac{r^2}{w(z)^2}\right) r\, dr\, d\theta \tag{A.3}$$

Separate the integrals:

$$P(R,z) = \frac{P}{\pi(w(z)^2/2)} \cdot \int_0^{2\pi} d\theta \cdot \int_0^R r \exp\left(-2\frac{r^2}{w(z)^2}\right) dr \tag{A.4}$$

Evaluate the angular part:

$$P(R,z) = \frac{P}{\pi(w(z)^2/2)} \cdot 2\pi \cdot \int_0^R r \exp\left(-2\frac{r^2}{w(z)^2}\right) dr \tag{A.5}$$

Simplify the constants:

$$P(R,z) = \frac{2P}{w(z)^2/2} \cdot \int_0^R r \exp\left(-2\frac{r^2}{w(z)^2}\right) dr \tag{A.6}$$

Now evaluate the integral. Let

$$u = \frac{2r^2}{w(z)^2} \quad \Rightarrow \quad du = \frac{4r}{w(z)^2} dr \quad \Rightarrow \quad dr = \frac{w(z)^2}{4r} du$$

Then the integral becomes:

$$\int_0^R r \exp\left(-2\frac{r^2}{w(z)^2}\right) dr = \frac{w(z)^2}{4}\left[1 - \exp\left(-2\frac{R^2}{w(z)^2}\right)\right] \tag{A.7}$$

Substitute back into the expression for $P(R, z)$:

$$P(R, z) = \frac{2P}{w(z)^2/2} \cdot \frac{w(z)^2}{4}\left[1 - \exp\left(-2\frac{R^2}{w(z)^2}\right)\right] \tag{A.8}$$

Simplify:

$$P(R, z) = P\left[1 - \exp\left(-2\frac{R^2}{w(z)^2}\right)\right] \tag{A.9}$$

Thus, the fraction of total power contained within a radius R is:

$$\boxed{\eta(R, z) = \frac{P(R, z)}{P} = 1 - \exp\left(-2\frac{R^2}{w(z)^2}\right)} \tag{A.10}$$

Appendix B

Wavefront Methods

The other model used for beam representation is the wavefront model, which is used when a more in-depth analysis of the issues affecting the beam is needed. The equation for the electric field $\vec{E}$ (which varies with space and time) is given by eq. (B.1), where c is the speed of light in the medium, and ∇ is the Laplacian operator.

$$(c^2\nabla^2 - \frac{\partial^2}{\partial t^2})\vec{E} = \vec{0} \tag{B.1}$$

For a monochromatic wave, $\vec{E}$ takes the form eq. (B.2), where $k = \frac{\omega}{c} = \frac{2\pi f}{c}$ is the wavenumber. The wavevector $\vec{k}$ is the direction of the propagation of energy, whose magnitude equals the wavenumber: $|\vec{k}| = k$.

$$\vec{E}(\vec{r}, t) = \vec{E}(\vec{r})e^{-i\omega t} \tag{B.2}$$

Substituting eq. (B.2) into eq. (B.1) yields eq. (B.3). Equation (B.3) is very useful, as it has separated the time dependence of the field from the spatial domain.

$$\left[\nabla^2\vec{E}(\vec{r}) + k^2\vec{E}(\vec{r})\right]e^{-i\omega t} = \vec{0} \tag{B.3}$$

For eq. (B.3) to hold for any t, it must be that eq. (B.4) is true.

$$\nabla^2\vec{E}(\vec{r}) + k^2\vec{E}(\vec{r}) = \vec{0} \tag{B.4}$$

To simplify the analysis, the assumption is made that the electric field is tightly concentrated along one axis, called in this work the optical axis (z axis). In a TIL application, this axis would be collinear with the line joining the optical system and the target. As the target is only a few meters long measured perpendicular to the optical axis, and the optical system to target length can be in the order of kilometers, this assumption is natural. The implication for the mathematics is that the wavevector is or is almost collinear with the z axis, hence (where subscript indicates direction in Cartesian space): $k \approx k_z$.

A general solution to eq. (B.4) can then be expressed as eq. (B.5), where $U(x, y, z)$ is the amplitude of the wave at a spatial location.

$$\vec{E}(\vec{r}) = U(x,y,z)e^{i(k_x x + k_y y + k_z z)} \tag{B.5}$$
$$\approx U(x,y,z)e^{ik_z z} \tag{B.6}$$
$$\approx U(x,y,z)e^{ikz} \tag{B.7}$$
$$\tag{B.8}$$

Substituting eq. (B.5) into eq. (B.4) yields eq. (B.9).

$$\nabla^2(Ue^{ikz}) + k^2 U e^{ikz} = 0 \tag{B.9}$$

Using the chain rule, eq. (B.9) can be expanded into the transverse and longitudinal Laplacian ($\nabla_\perp^2 U$ and $\frac{\partial^2 U}{\partial z^2}$, respectively), as shown in eq. (B.10).

$$0 = \Big(\nabla_\perp^2 U + \frac{\partial^2 U}{\partial z^2} + 2ik\frac{\partial U}{\partial z} - k^2 U\Big)e^{ikz} + k^2 U e^{ikz} \tag{B.10}$$
$$= \Big(\nabla_\perp^2 U + \frac{\partial^2 U}{\partial z^2} + 2ik\frac{\partial U}{\partial z}\Big)e^{ikz} \tag{B.11}$$

Since $|\frac{\partial^2 U}{\partial z^2}|$ will tend to be significantly smaller than $|k\frac{\partial U}{\partial z}|$, it is further eliminated from the equation. The resulting equation is shown in eq. (B.12), known as the Helmholtz equation [24].

$$2ik\frac{\partial U}{\partial z} + \nabla_\perp^2 U = 0 \tag{B.12}$$

The final step is to integrate eq. (B.12) along the z dimension to propagate the wave. Although numerical methods could be used, it is typical to use a series of Fourier transforms to simplify implementation and guarantee numerical stability [24]. At a high level, these methods include applying the Fourier transform to the wavefront field, multiplying the wavefront in the frequency domain by eq. (B.13), and then applying the inverse Fourier transform [24].

$$H(f_x, f_y) = \exp\left[i2\pi z\sqrt{\frac{1}{\lambda^2} - f_x^2 - f_y^2}\right] \tag{B.13}$$

Appendix C

Zernike Polynomials

To simplify the analysis of wavefront aberrations, it is common to represent a wavefront as a linear sum of Zernike polynomials, as defined in the polar equation eq. (C.1) [25]. With $n \geq m \geq 0$ (n is known as radial order and m as azimuthal frequency), ϕ is the azimuthal angle, and r is the normalized radial distance $0 \leq r \leq 1$ of an optical system. A_{nm} and B_{nm} are the coefficients that determine the amplitude of the wavefront aberrations.

$$\begin{aligned} W(r,\phi) &= \sum_{n,m} \Big(A_{nm} \sin m\phi + B_{nm} \cos m\phi \Big) R_n^m(r) \\ &= \sum_{n,m} A_{nm} Z_n^{-m}(r,\phi) + B_{nm} Z_n^m(r,\phi) \end{aligned} \tag{C.1}$$

Sometimes it is preferable to look at a specific order m and n, mathematically shown as Z_n^m. The expressions for both are shown in eqs. (C.2) and (C.3).

$$Z_n^{-m}(r,\phi) = R_n^m(r)\cos(m\phi) \tag{C.2}$$

$$Z_n^m(r,\phi) = R_n^m(r)\sin(m\phi) \tag{C.3}$$

In eq. (C.1), r and ϕ are the polar coordinates, related to (x, y) by the following relations $x = r\cos\phi$ and $y = r\sin\phi$. A_{nm} and B_{nm} are coefficients that determine the extent of each aberration. $R_n^m(r)$ is called the radial polynomial, which is defined in eq. (C.4) [25].

$$R_n^m(r) = (-1)^{\frac{n-m}{2}} \binom{\frac{n+m}{2}}{m} r^m F\left(\frac{n+m+2}{2}, -\frac{n-m}{2}, m+1, r^2\right) \tag{C.4}$$

In eq. (C.4), F is the called Gauss's hypergeometric function, defined by eq. (C.5) [26]. It uses the notation $(x)_0 = 1$ and $(x)_n = x \cdot (x+1) \cdots (x+n-1)$.

$$F(a,b,c,z) = \sum \frac{(a)_n (b)_n}{(c)_n n!} z^n \tag{C.5}$$

For example, if one wanted to compute $Z_n^m(r,\phi)$ for $n = 1$ and $m = 1$, one could start to compute compute $R_1^1(r)$ as in eq. (C.6).

$$\begin{aligned} R_1^1(r) &= (-1)^{\frac{1-1}{2}} \binom{\frac{1+1}{2}}{1} r^1 F(\frac{1+1+2}{2}, -\frac{1-1}{2}, 1+1, r^2) \\ &= r \cdot F(2, 0, 2, r^2) \\ &= r \cdot \sum \frac{(2)_n (0)_n}{(2)_n n!} r^{2n} \\ &= r \end{aligned} \tag{C.6}$$

Knowing the value $R_1^1(r)$, finding $Z_1^1(r, \phi)$ is a simple substitution, as shown in eq. (C.7).

$$\begin{aligned} Z_1^1(r, \phi) &= \cos 1\phi R_1^1(r) \\ &= r \cos \phi \end{aligned} \tag{C.7}$$

Within the context of software design, it is preferable to have a single index rather than use the radial order and azimuthal frequency. To convert from m and n to j, where j begins at 0 and extends indefinitely, one can use the relation eq. (C.8) [27].

$$j = \frac{n(n+2)+m}{2} \tag{C.8}$$

The common names until $j = 9$ are shown in table C.1. While more common names exist for higher orders, the higher orders tend to be less prevalent in adaptive optics and hence their ANSI index is used instead. Each Zernike mode up until $j = 27$ is visualized in fig. C.1.

Table C.1: Table showing (n, m), j and the Common Name for various Zernike Modes

(n, m)	j	**Common Name**
$(0, 0)$	0	Piston
$(1, -1)$	1	Tilt (Y)
$(1, 1)$	2	Tilt (X)
$(2, -2)$	3	Astigmatism (Oblique)
$(2, 0)$	4	Defocus
$(2, 2)$	5	Astigmatism (Vertical)
$(3, -3)$	6	Trefoil (Oblique)
$(3, -1)$	7	Coma (Y)
$(3, 1)$	8	Coma (X)
$(3, 3)$	9	Trefoil (Vertical)

As discussed before, linear combinations of Zernike polynomials can be used to represent a complex wavefront. For example, fig. C.2 makes use of 3 Zernike modes to display a slightly curved wavefront in the x direction, and a very curved wavefront in the y direction. The z units are arbitrary.

In practical scenarios, one evaluate each Zernike polynomial up to the desired order on a grid, and then use least squares to efficiently compute the correct Zernike coefficients.

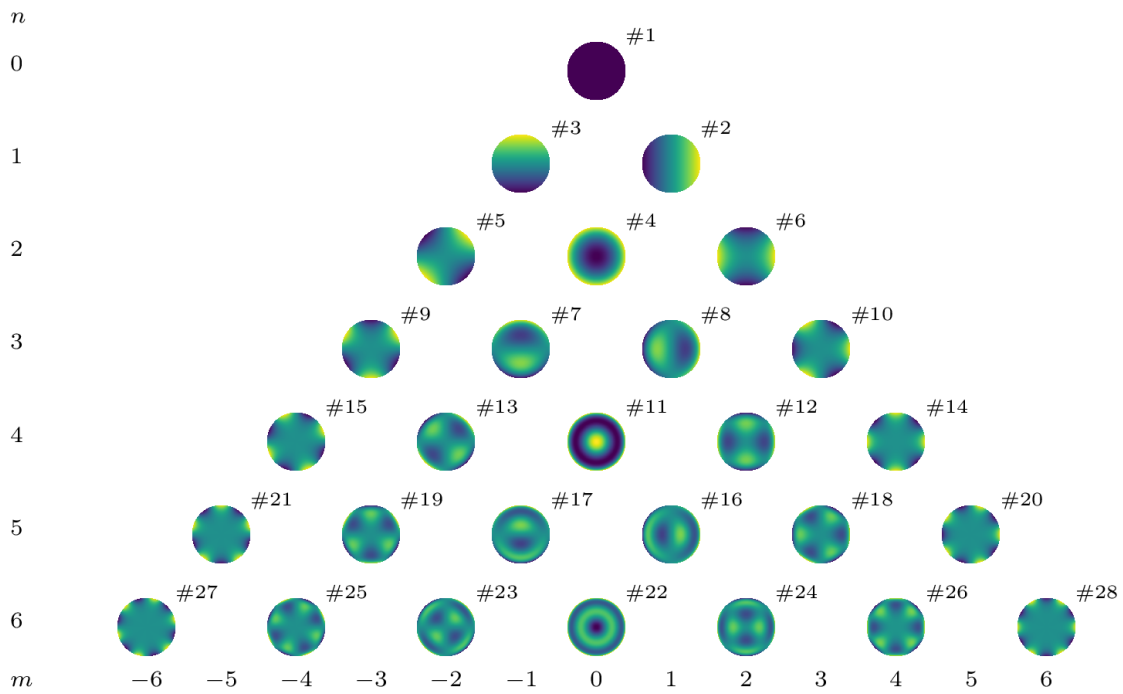

Figure C.1: Zernike table showing the different Zernike modes up to order $j = 27$. Lighter colors indicate larger values. From [28].

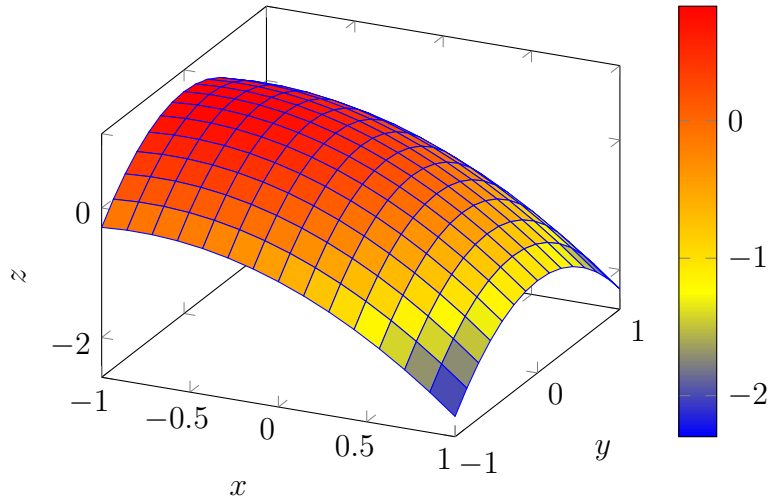

Figure C.2: Example wavefront with Zernike coefficients (using ANSI indexing): $C_2 = -1$, $C_4 = -0.25$ and $C_5 = 0.1$. Arbitrary z units.

Appendix D

Wavefront sensing

Wavefront sensing involves sensing the phase differences between different spatial locations of the beam. The measured the phase differences can then be used with the controller and deformable mirror to make up the control loop in controllers that require a wavefront sensor.

Ideally, one would measure such phase differences directly, with a sensor such as a Charge-Coupled Device (CCD) sensor. However, the frequency of IR light is around $\frac{3 \cdot 10^8}{1 \cdot 10^{-6}} = 3 \cdot 10^{14} =$ 300THz. The Nyquist sampling theorem states that the sample rate should be at least double the frequency to prevent aliasing, meaning a sample rate of 600THz. In any case, both these frequencies are well outside what contemporary sensors are capable of, and hence more creative phase measurement techniques were devised for adaptive optics each with advantages and disadvantages.

The key performance indicators of a technology for wavefront sensing are:

- The **sampling rate** sampling rate is a very important metric. With increasing turbulence strength, the rate of change of the atmosphere increases too. With a high sample rate one can overcome a more turbulent atmosphere.

- **Spatial resolution** is a measure of how many data points of the wavefront can be collected by the sensor. The more points, the better the wavefront estimation. For TIL applications, spatial resolution is not as important as the sampling rate as higher orders of abberations tend not to be as prevalent.

- **Dynamic range** is the ratio between the highest and lowest value of a wavefront measurement. This figure is important as it limits the resolution of the system. This figure just limits the spatial resolution, which for TIL applications is not as critical as the sampling rate.

- The **theoretical maturity** of the measurement device is a very important metric for this thesis, as the objective of this thesis is not to develop a new measurement device, but rather integrate existing systems. If a new system were to be used, much effort would go into developing it rather than assembling and testing a TIL system. Hence, it is critical that the chosen sensor is theoretically mature.

- **Ease of procurement** is a critical metric. As this is a practical thesis, the chosen sensor must be easily procured. This will typically mean a commercial-off-the-shelf (COTS) solution.

What follows is a brief overview of a selection of methods, including a critical analysis of the suitability of the sensor to TIL applications. In the end, one of the methods will be chosen for the experimental work of this thesis.

D.1 Shack-Hartmann Wavefront Sensor

The Shack-Hartmann wavefront sensor makes use of a micro-lens array (MLA) and a CCD sensor. The MLA is a small plastic component typically manufactured with photolithography with many very small lenses. Each lens is in the order of 100 µm in diameter, in a arrays of 10×10 or more. The key assumption of the Shack-Hartmann sensor is that within the 100 µm diameter lens the wavefront just possesses a tip/tilt component, and that higher order Zernike modes are so small they are negligible. Figure D.1a shows a figure of this phenomenon in two dimensions. In three dimensions, two shifts will be registered: Δx and Δy. The equation for the two is given in eqs. (D.1) and (D.2) [18].

$$\frac{\partial W}{\partial x} = \frac{\Delta x}{f} \tag{D.1}$$

$$\frac{\partial W}{\partial y} = \frac{\Delta y}{f} \tag{D.2}$$

Considering just 4 lenses of possibly hundreds, one can obtain a spotfield image such as fig. D.1b. Note that the grid has to be compted from a calibration step, where a known planar wavefront is used to calibrate the sensor. As for a planar wavefront all the spots should be in an ideal grid, the spot centers can be recorded and then used to compare to an unknown wavefront.

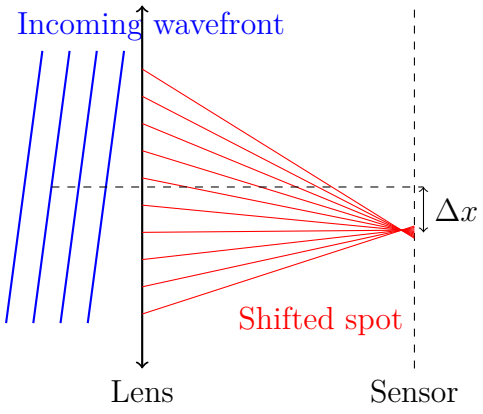

(a) Single lens in a MLA. The incoming wavefront just has a tip/tilt aberration, resulting in a spot shift Δx that is measured on a wavefront sensor. There are some higher order abberations which cause the focal point to be before the sensor plane, however these are too small to influence the spot shift measurement.

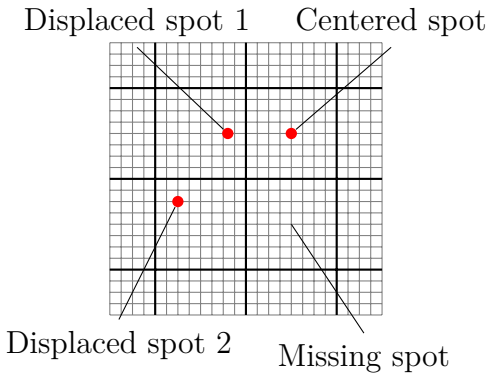

(b) Simulated spotfield image with grid overlayed various cases: *centered spot*, *displaced spot* and *missing spot*. The *centered spot* shows that locally the field is planar, whereas for the two *displaced spots* the field has some tip/tilt to it. The *missing spot* could be a result of dead pixels, faulty MLA or extreme tip/tilt that causes the spot to be outside the expected grid.

Figure D.1: Two diagrams showing the basic operation of a Shack-Hartmann wavefront sensor. Figure D.1a shows a microlens focusing a small portion of light to a spot. Figure D.1b shows the culmination of 4 lenses and the resulting spots shifts.

Table D.1 shows the evaluation of a Shack-Hartmann sensor based on theory. While spatial resolution and dynamic range are poor for a Shack-Hartmann sensor, the sampling rate of

100 Hz is decent. The theoretical maturity and ease of procurement make this sensor especially attractive.

Table D.1: Table showing evaluation of Shack-Hartmann wavefront sensor

Criteria	**Evaluation**
Sampling rate	Limited by the frame rate of the sensor and the sensor's ability to capture and amplify the energy of photons. Practically around 100Hz sampling rate should be expected.
Spatial resolution	Poor spatial resolution as many pixels are used for a single spot.
Dynamic range	Dependent on the sensor dynamic range, which is typically less than photodetectors.
Theoretical maturity	Very theoretically mature, with extensive literature available. The author of this work also has practical experience with Shack-Hartmann wavefront sensors.
Ease of procurement	Shack-Hartmann wavefront sensors are COTS components can can be bought individually [1] or as kits [2].

Note that some modifications could be applied to a Shack-Hartmann wavefront sensor, such as a specialized quad-photodiode to estimate each spot shift, however these solutions would not be theoretically mature and would not be easy to procure.

D.2 Curvature Sensor

A curvature sensor works in a similar fashion as the Shack-Hartmann wavefront sensor, in that it uses a MLA to focus light. However, instead of using a CCD like the Shack-Hartmann sensor, it uses a pair of photodiodes per microlens. A schematic drawing of a single lens is shown in fig. D.2, where two detector planes measure the intensity on either side of the focal plane of the lens. The difference in measured intensities is proportional to the wavefront curvature in the microlens, which can be used to reconstruct the wavefront for use in the controller [29]. Alternatively, as the measured value is the second derivative of the wavefront, it can directly be applied to some deformable mirrors (covered later), meaning a completely analog closed loop control. The curvature sensor is evaluated in table D.2.

While a curvature sensor does present some significant technical advantages over a Shack-Hartmann sensor, the practical disadvantages associated with its use mean that it is much less desirable in this work.

D.3 Pyramid Sensor

A Pyramid wavefront sensor is another type of wavefront measurement device. Figure D.3 shows a schematic of how such a sensor works. Light is focused with a mirror to nearly the apex of a pyramid optic. The image is split into 4, with which the help of an optical relay is projected onto a detector. In operation, the pyramid optic is oscillated, and differences

[1] https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=5287
[2] https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=3208

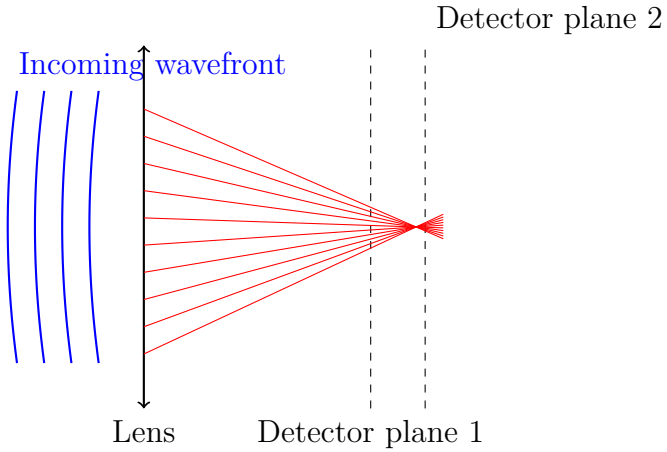

Figure D.2: Wavefront curvature sensor diagram.

Table D.2: Table showing evaluation of Wavefront curvature sensor

Criteria	**Evaluation**
Sampling rate	Dependent on photodiode speed and electronics, however can easily exceed 100s of kHz, possibly in excess of 10 MHz.
Spatial resolution	Much better than a Shack-Hartmann sensor, as two photodetector per microlens array is used
Dynamic range	Limited to dynamic range of photodiodes, which is much greater than a Shack-Hartmann sensor
Theoretical maturity	Not as developed as Shack-Hartmann sensor, due to not being adopted at a large scale
Ease of procurement	No commercial curvature sensors were found

in intensity patterns in each of the 4 pupils correspond to phase differences in the incoming wavefront. A perfect planar wavefront would hence register as a uniform intensity distribution.

The detected intensity $I(x, y)$ of each quadrant is related to the derivatives of the wavefront as shown in eqs. (D.3) and (D.4).

$$I(x, y) \propto \frac{\partial W}{\partial x} \tag{D.3}$$

$$I(x, y) \propto \frac{\partial W}{\partial y} \tag{D.4}$$

A short evaluation of the Pyramid wavefront sensor is shown in table D.3.

A Pyramid wavefront sensor is a good alternative for AO applications that require a high spatial resolution and where sampling rate is not as critical. In TIL, the opposite is true: high sample rate at low spatial resolutions is acceptable. Hence, the Pyramid wavefront sensor is not a suitable alternative.

D.4 Foucalt Knife Edge Test

Another way to measure the wavefront, that is typically used in amateur astronomy is the (Foucalt) knife edge test. This test involves a moving knife edge (a very well defined edge)

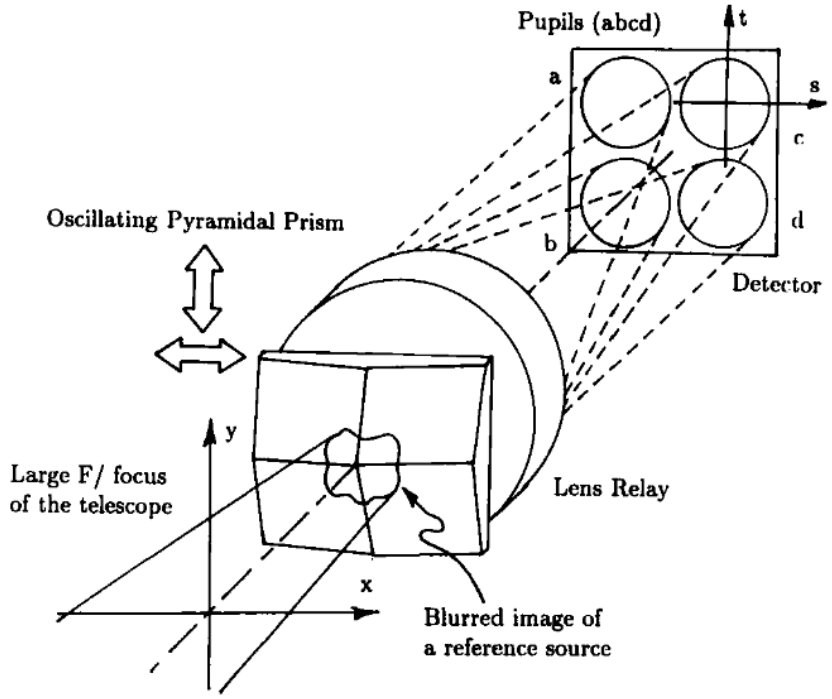

Figure D.3: Pyramid wavefront sensor. Image from [30].

Table D.3: Table showing evaluation of Pyramid wavefront sensor

Criteria	**Evaluation**
Sampling rate	Like the Shack-Hartmann sensor, the sampling rate is limited by the sampling rate of the CCD sensor. Hence, around 100 Hz is expected.
Spatial resolution	As the pyramid amplifies the wavefront slopes, the spatial resolution is better than a Shack-Hartmann sensor
Dynamic range	The dynamic range is dependent on the type of sensor used, which is less than photodetectors
Theoretical maturity	Not adopted at a large scale, hence not as theoretically mature as a Shack-Hartmann sensor
Ease of procurement	No commercial pyramid sensors were found

across the z axis. As the knife edge moves across the wavefront, the wavefront slopes can be observed and integrated, resulting in the wavefront. The evaluation of the Foucalt knife edge test is shown in table D.4.

The Foucault knife edge test is good for simple, inexpensive testing of equipment. While the test could be adapted to act like a wavefront sensor, its fundamental limitation is a move part, which must accelerate and slow down frequently, limiting the sampling rate significantly.

D.5 Lateral Shift Interferometer

The last wavefront sensor considered in this work is the lateral shift interferometer. These devices use the interference of light to measure the wavefront. For this specific analysis, the Quadriwave Lateral Shearing Interferometer (QWLSI) by Phasics will be analyzed [31]. Oper-

Table D.4: Table showing evaluation of Foucault knife edge test wavefront sensor

Criteria	Evaluation
Sampling rate	Slow, as knife edge needs to scan across the aperture.
Spatial resolution	Dependent on positioning accuracy of knife edge, which could be very precise.
Dynamic range	Dependent on CCD sensor, less than a photodetector
Theoretical maturity	Very mature, however no widely used practical implementation
Ease of procurement	No commercial Foucalt knife edge test sensors were found

ationally, QWLSI works with a similar principle as a Shack-Hartmann wavefront sensor, but rather than using a MLA, QWLSI uses a diffractive grating. The diffractive grating splits the incoming beam into 4, after which they interfere and interfere. The interference pattern encodes information about the phase gradients, which can be reconstructed with the detector. The main advantages of this technology is the very high spatial resolution and high dynamic range. Figure D.4 shows the different masks used for Hartmann (not used in this work), Shack-Hartmann and Phasics QWLSI sensors. QWLSI is evaluated in table D.5.

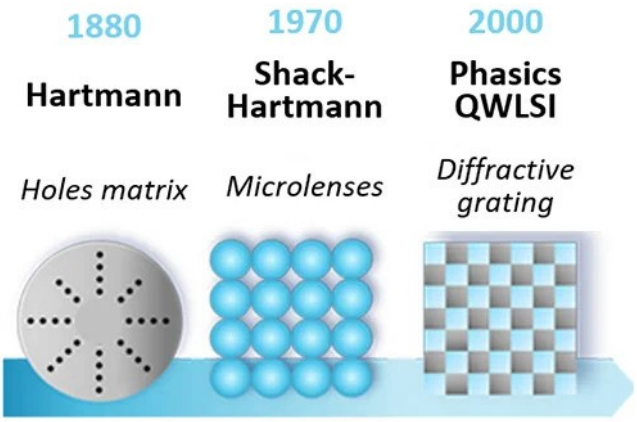

Figure D.4: Quadriwave Lateral Shearing Interferometer mask, compared to Shack-Hartmann sensor and the earlier Hartmann Holes matrix. Image taken from [31].

Table D.5: Table showing evaluation of Lateral shift interferometer wavefront sensor

Criteria	Evaluation
Sampling rate	Similar to Shack-Hartmann sensor, limited by imaging system frame rate. About 100 Hz
Spatial resolution	Much better spatial resolution, with Phasics advertising 852×720 spots [31]
Dynamic range	Much higher than Shack-Hartmann sensor, with advertising of 500 µm PtV [31]
Theoretical maturity	Introduced in 2004, the system has extensive use.
Ease of procurement	Available for purchase as individual unit [31]. Not part of kit

QWLSI is a definitive technological improvement over a traditional Shack-Hartmann wavefront

sensor. Unfortunately, all the improvements are in the parts of the sensor that are not as useful for TIL. Additionally, procurement of a QWLSI would be for an individual wavefront sensor, and cannot be bought as a kit. This makes it slightly less appealing than a Shack-Hartmann sensor.

D.6 Wavefront Sensor Selection

With 5 wavefront sensor technologies investigated, one can now select which sensor is best for TIL. The trade off is shown in table D.6 and is based off the investigation of each technology.

The final sensor must be easy to procure and use, which eliminates all but the Shack-Hartmaan sensor and QWLSI. Both perform identically in the most important aspects (sampling rate, theoretical maturity and ease of procurement), however the Shack-Hartmaan sensor performs worse in terms of spatial resolution and dynamic range, both of which are not as relevant in TIL. As such, both are essentially equally as appealing to use, however the Shack-Hartmaan sensor is chosen as it is available in an AO kit which can be modified into a TIL test setup.

Table D.6: Trade off table of the various wavefront sensors. The width of the column represents the importance. Green means good, yellow intermediate, and red not good. Ratings are based off the other candidates.

Sensor name	**Sampling rate**	**Spatial resolution**	**Dynamic range**	**Theoretical maturity**	**Ease of procurement**
Shack-Hartmann	Sensor-based	Limited pixel array size	Less than photodetectors	Used extensivly in industry	Available to purchase
Curvature sensor	Photodetector-based	Two photodetectors used per microlens	Photodetector-based	Not as developed as Shack-Hartmann	No commercial product
Pyramid sensor	Sensor-based	Wavefront slopes amplified	Less than photodetectors	Not adopted at large scale	No commercial product
Foucalt Knife Edge	Sensor-based	Can be very precise	Less than photodetectors	Not widely used	No commercial product
Lateral Shift Interferometer	Sensor-based	852×720	500 µm	Operational since 2004	Available to purchase

Appendix E

Wavefront Actuator

The second of third component required for closed loop control is an actuator: how the wavefront will be manipulated. Within wavefront manipulation, two main technologies exist: deformable mirrors and spatial light modulators (SLM). While the method of functioning is different for both, their end use is the same: modify an incoming wavefront. Figure E.1 shows two examples of wavefront manipulation: fig. E.1a of a deformable mirror and fig. E.1b of a SLM.

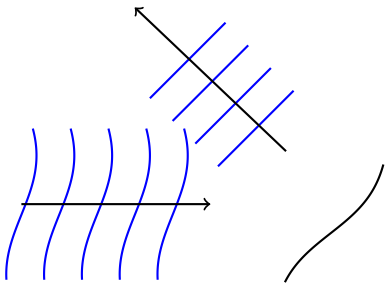

(a) Deformable mirror example.

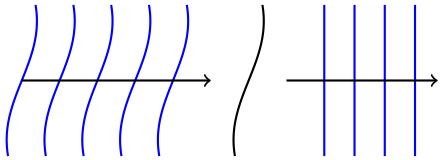

(b) Transmissive spatial light modulator example. Reflective SLMs exist too.

Figure E.1: Two main wavefront manipulation strategies: deformable mirrors and spatial light modulators. Both accept an incoming wavefront, and ideally correct it such that the outgoing wavefront is planar.

In the following subsections, different types of deformable mirrors will be investigated and evaluated, and later spatial light modulators will be investigated. Then the most suitable actuator for this work will be chosen.

E.1 Deformable Mirror

Within the realm of deformable mirrors, there are various types of constructions. One difference between mirrors is segmented versus continuous mirrors. Segmented mirrors are composed of many different segments that move independently of each other. Continuous segmented mirrors are made of one continuous material that is reflective on one side. When this material is acted on by a force, it changes the shape of the mirror.

Two deformable mirrors are needed for this work: one to simulate the atmosphere and another to compensate the beam. Simulating the atmosphere with a deformable mirror allows for precise control over the strength and speed of the turbulence, allowing for more quantitative statements to be made about the performance.

In the following, 3 different deformable mirrors will be investigated. They are chosen due to being easily procured, especially relevant, or due to their novelty. It is not an exhaustive list.

E.1.1 Bimorph Mirror

A bimorph mirror is typically a segmented mirror, and is formed by a plurality of layers of different materials. Active layer(s) are made of a piezoelectric material, and when a voltage is applied to the piezoelectric material, it deforms and hence moves the mirror, up or down, making a single actuator. With multiple segments one can make a deformable mirror capable of correcting high order aberrations. Piezoelectric bimorph segmented mirrors are available for purchase, for example [1].

E.1.2 MEMS Mirror

MEMS mirrors are continuous mirrors with a thin reflective membrane that is controlled by micromachined actuators. These actuators on MEMS mirrors can be extremely small, meaning much higher aberrations can be corrected for than bimorph mirrors. Typically, MEMS devices can become very cheap as they can greatly benefit from economies of scale. MEMS mirrors are available for purchase, for example[2].

E.1.3 Membrane Mirror

Membrane mirrors are another type of continuous mirrors made of a thin conductive and reflective membrane that is stretched over a frame. Voltages can be applied to electrodes on the other side of the mirror, repelling or attracting the membrane. This changes the surface shape, actuating the mirror.

The equation governing a membrane mirror is of the form eq. (E.1) (known as the Poisson equation), where z is the membrane surface, P is the electrostatic pressure applied, and A and B are constants [29].

$$\frac{\partial^2 z}{\partial t^2} = A\nabla^2 z + BP \tag{E.1}$$

This sensor is of particular relevance as it can be used with a curvature sensor in an analog fashion. This is possible as the curvature sensor measures the local curvature, whereas applying an electrostatic pressure results in a change of curvature of the mirror. Hence, the analog signal from the curvature sensor need only be amplified and inputted to the deformable mirror, and a high bandwidth closed loop system is formed. As it is known that a curvature sensor will not be used, the specific advantages of using a membrane mirror are not present. However, if future work considers using a curvature sensor then a membrane mirror may become a much better candidate.

[1] https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_ID=5056
[2] https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=3258

E.2 Spatial Light Modulator

Spatial light modulators are another type of wavefront actuator. Their working principle is a crystal that changes refractive index based on the applied electric field. As deduced by eq. (2.12), a change in refractive index is proportional to a phase shift. Spatial light modulators typically cannot have a high actuator count as other wavefront technologies, however they have no moving parts and hence can operate at frequencies well in the MHz range [32].

Figure E.2 shows an example transmissive SLM with two regions: one with an applied electric field and the other without. The phase delay in the crystal is such that the complex input wavefront is converted to a planar output. Note that reflective SLMs also exist, and work in a similar fashion.

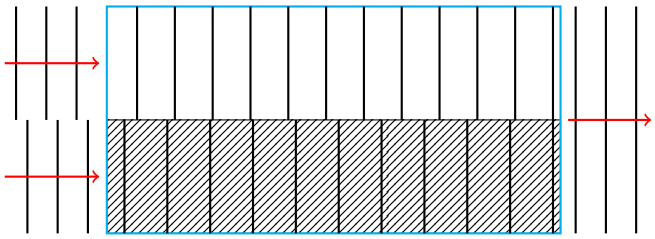

Figure E.2: Spatial light modulator crystal (cyan rectangle) simple example with two modes: one activated (hatched pattern) and the other deactivated (non-filled area inside crystal). An electric field is applied to the hatched area to incur a phase delay relative to the top wave, leading to an planar output wave.

E.3 Wavefront Actuator Selection

With the available technology reviewed, the final selection for wavefront actuator can be selected. Two wavefront actuators must be selected: one for control of the precompensated beam and the other for atmospheric turbulence simulation. Table E.1 shows the trade off for the various actuators. Actuator count represents the amount of actuators that are typical on such a mirror, actuator stroke is the maximum deflection an actuator can perform.

Having performed a trade off in table E.1, it is clear that bimorph and MEMS mirrors are desirable. For the turbulence inducing mirror, as multiple phase screens are to be used one would benefit from a high spatial resolution (large actuator count) to incorporate the effects of multiple phase screens. Hence for simulating the atmospheric turbulence a MEMS based mirror is chosen. For beam compensation, a large actuator stroke is beneficial as it allows for correcting aberrations caused by slight misalignment in passive optical elements. As higher order aberrations need not be corrected in TIL applications, a bimorph mirror can be used.

With the wavefront actuators chosen, the TIL controller can now be designed.

Table E.1: Trade off table of the various wavefront actuators. The width of the column represents the importance. Green means good, yellow intermediate, and red not good. Ratings are based off the other candidates.

Actuator name	**Actuator count**	**Actuator stroke**	**Ease of procurement**
Bimorph mirror	Actuator stroke limited by small actuator footprint	Large stroke	Available individually or as a kit
MEMS mirror	Many actuators	Medium stroke	Available individually or as a kit
Membrane mirror	Actuator stroke limited by small actuator footprint	Medium stroke	No commerical products
SLM	Few modes	Single wavelength stroke common	Outside budget

Appendix F

Power Loss due to Lambertian Reflection

To estimate the power loss due to Lambertian reflection, one can first draw a diagram of the situation, as shown in fig. F.1.

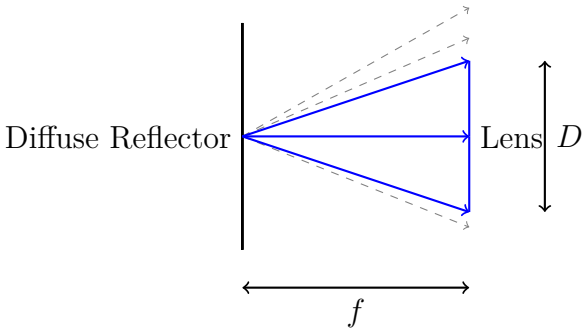

Figure F.1: Diagram of Lambertian reflection off a diffuse reflector.

The radiance is proportional to the cosine of the angle between the surface normal and the direction of emission in Lambertian reflection, as shown in fig. F.1

$$I(\theta) = I_0 \cos\theta \tag{F.1}$$

As the differential surface area element on a unit sphere is $dA = \sin\theta d\theta d\phi$ and intensity is the power per unit area, the total power is given by eq. (F.2), with A_0 being the surface area element, θ is the polar angle and ϕ the azimuthal angle. θ_{max} is defined as equal to $\tan\frac{D}{2f}$.

$$\begin{aligned} P &= \int_{A_0} I dA \\ &= \int_{A_0} I_0 \cos\theta dA \\ &= \int_0^{2\pi}\int_0^{\theta_{max}} I_0 \cos\theta \sin\theta d\theta d\phi \\ &= 2\pi I_0 \cdot \int_0^{\theta_{max}} \cos\theta \sin\theta d\theta \\ &= 2\pi I_0 \frac{\sin^2\theta_{max}}{2} \\ &= \pi I_0 \sin^2\theta_{max} \end{aligned} \tag{F.2}$$

The total power irradiated is given by $P_{tot} = \pi I_0 \sin^2 \pi/2 = \pi I_0$. The power captured by the lens can be found by replacing θ_{max} in terms of D and f, as show in eq. (F.3). To simplify the expression the assumption $\frac{D^2}{4f^2} \ll 1$ is used. This is typically true as the focal length tends to be much larger than the diameter of the optic.

$$\begin{aligned} P_{captured} &= \pi I_0 \left(\frac{\frac{D}{2}}{\sqrt{f^2 + \frac{D^2}{4}}} \right)^2 \\ &= \pi I_0 \frac{\frac{D^2}{4f^2}}{1 + \frac{D^2}{4f^2}} \\ &\approx \pi I_0 \frac{D^2}{4f^2} \end{aligned} \tag{F.3}$$

The ratio of the power captured and power irradiated is:

$$\frac{P_{captured}}{P_{tot}} = \frac{\pi I_0 \frac{D^2}{4f^2}}{\pi I_0} = \frac{D^2}{4f^2} \tag{F.4}$$

The fraction of power lost is:

$$\frac{P_{loss}}{P_{tot}} = 1 - \frac{D^2}{4f^2} \tag{F.5}$$

Appendix G

T-I: Individual and Subassembly Testing

The purposes of T-I tests are to ensure a proper understanding and operation of the individual components. As oftentimes it is impossible to test optical components on their own, the end product of the T-I tests is a setup that can easily be extended into a TIL setup in later tests. Appendix G.1 will be used to test the Shack-Hartmann Wavefront Sensor (SHWFS). Appendix G.2 will then ensure the proper forming of a point spread function through the use of a CCD instead of the SHWFS. Next, in appendix G.3 and appendix G.4 the two deformable mirrors will be added and tested to ensure the poke matrices are calibrated correctly. If all the tests are successful, then the TIL testing can begin in appendix H.4. Otherwise, appendix H.3 will test all the optical equipment as was intended by the manufacturer to gain further insight into issues.

G.1 T-I.1: Test Shack-Hartmann Wavefront Sensor

The purpose of test T-I.1 is to ensure that the Shack-Hartmann sensor can be interfaced with and the expected wavefront is retrieved. The issue with just placing a laser diode in front of the SHWFS is that it may overexpose the sensor, as only around 5 % of the optical power is incident on on the SHWFS in the original kit table H.1. For this reason, two beam splitters are used in T-I.1, as shown in fig. G.1.

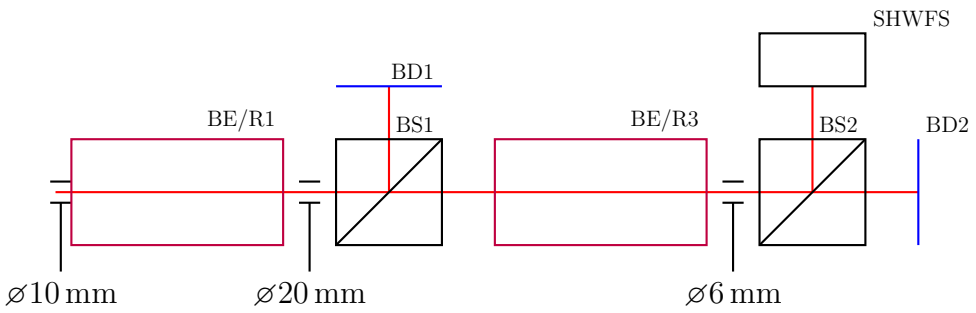

Figure G.1: T-I.1 Test setup

The beam exits subassembly 1 (fig. H.1b), and immediately enters subassembly 2 (BE/R1), where the ⌀10 mm, is expanded to ⌀20 mm. Next, it encounters BS1, where half of the beam optical power is incident on a beam dump, and the other half continous to a second beam

expander/reducer (BE/R3), where it is converted to a ⌀6 mm beam. After, 10 % of the optical power is reflected to the SHWFS and the other 90 % to a beam dump.

Adjusting the fine tuning mirrors in subassembly 1 should result in visible changes to the tip and tilt of the beam as measured by the SHWFS.

G.2 T-I.2: Test Camera Modules (x2)

The next test will be used to test the correct observation with the imaging sensor. The Shack-Hartmann sensor will be replaced with two imaging sensors. As the imaging sensors are just sensors with no lenses, L2 will be used to focus the light, as shown in fig. G.2.

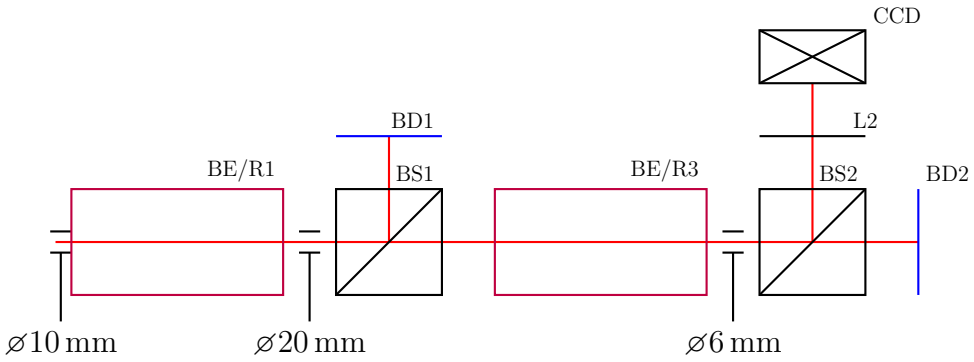

Figure G.2: T-I.2 Test setup

One would expect that using the fine tuning mirrors in subassembly 2 result in translational movement in the beam as observed by the imaging sensor.

G.3 T-I.3: Test Piezoelectric DM

In test T-I.3, the first actuator will be tested. The objective of this test is to introduce a known aberration using the deformable mirror, and observe it on the Shack-Hartmann sensor. This is possible as the mirrors are calibrated before hand by the manufacturer. The schematic of the test is shown in fig. G.3.

The test would be a success if the aberrations introduced to the deformable mirror can be observed by the wavefront sensor.

G.4 T-I.4: Test MEMS DM

Test T-I.4 is similar to the previous one, except that now a second deformable mirror is used, DM2 (MEMS deformable mirror). As with DM1, it comes calibrated from the manufacturer with known aberrations, hence it should be possible to observe these with the Shack-Hartmann sensor. DM1 should be in the flat configuration, where it behaves like a conventional flat mirror.

With this setup it will be possible to generate the simulated atmospheric turbulence, as will be used for the tests further on. DM2 and the SHWFS will operated in a closed loop, meaning a precise shape of the wavefront can be achieved. Starting from the phase screen at discrete time t, called ϕ_t, generated using the power law discussed in section 2.1.3. With the same theory as

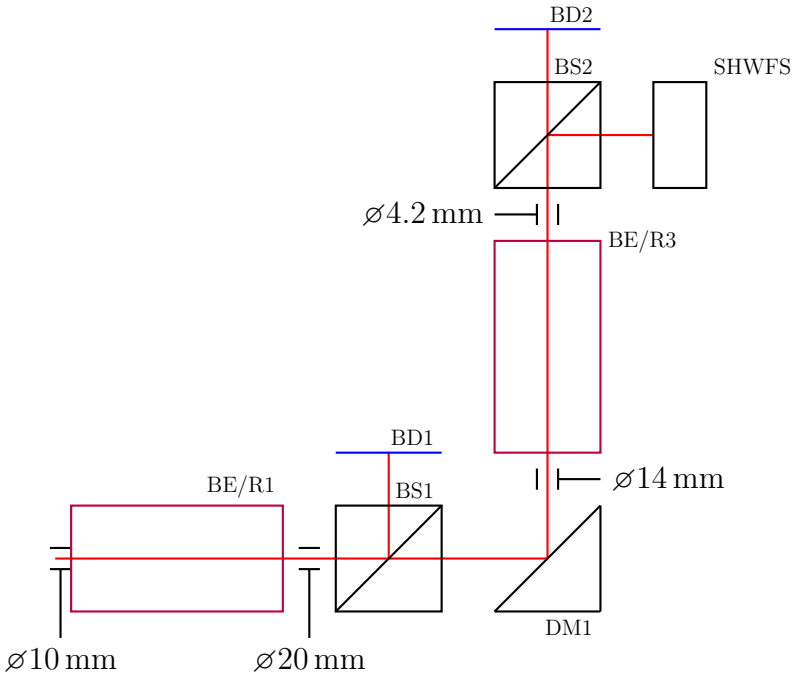

Figure G.3: T-I.3 Test setup

in section 2.2, one can iterate towards a control vector $\vec{u}_t$ solution that obtains the simulated phase screen of the atmosphere. Repeating the steps for multiple times with a small dt, results in a 2D array of $\vec{u}$ that can be used later on to simulate the atmosphere.

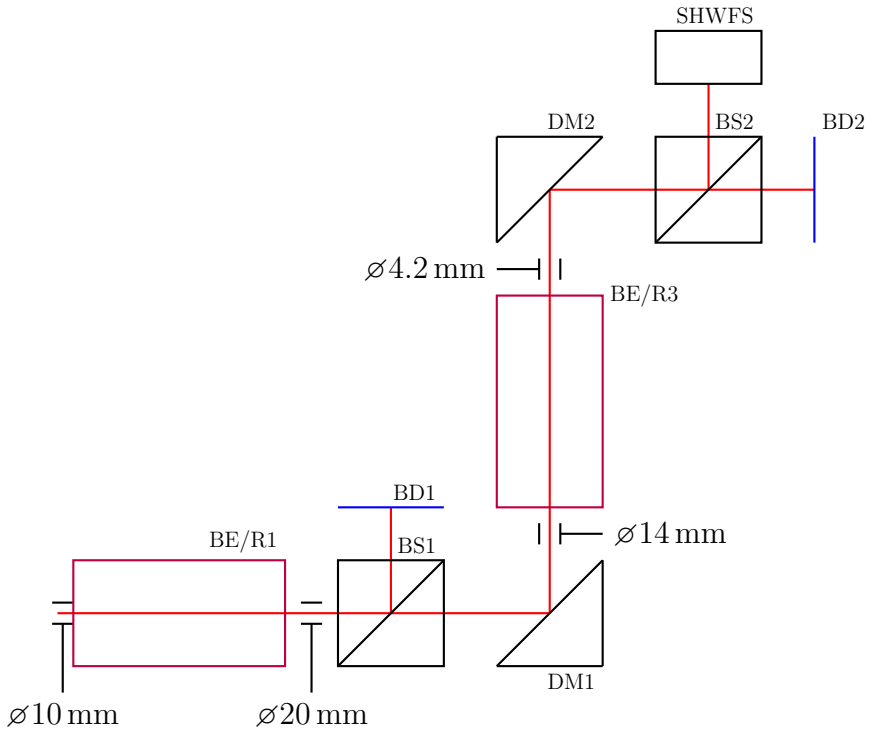

Figure G.4: T-I.4 Test setup

The same criteria as T-I.3 is used to evaluate a successful test. With the successful conclusion of T-I.4, and all the previous tests, one can move directly to appendix H.4. In addition, one requirement for this test is to re-obtain the input phase screen as a function of time from the physical setup. Otherwise, appendix H.3 can be used for debugging and understanding better

the shortfalls.

Appendix H

Hardware Experimental

In the previous chapter, the theory behind TIL was explored, including the relevant modeling techniques and hardware. With a good overview of the theory and relevant literature, it is now possible to design a test setup that can be used to test the different controllers, as well as design the simulation environment.

This chapter is broken into two main sections: the first will propose a setup to test TIL in a practical setup, and the second will propose a simulation test setup.

As part of the practical TIL setup, the available hardware and notation is presented in appendix H.1. Then, the individual and subassemblies present in the kit will be tested as outlined in appendix H.2. Appendix H.3 will test the correct functioning of the basic adaptive optics kit, as was intended by the manufacturer. Thirdly, the TIL setup will be assembled in appendix H.4. Fourthly, as the TIL setup is built and functional at this point, the different controller designs will be tested in appendix H.5.

Then, as part of the simulation tool, subassembly and individual tests are proposed in section 3.1.

This work sets the groundwork needed to construct the full simulation of a TIL system, as shown in chapter 3.

H.1 Experimental Hardware Available and Schematic Notation

For the proceeding content, it is important to establish the available hardware, and also what means will be used to represent the hardware in schematics. The hardware choice is presented in appendix H.1.1, and the notation in appendix H.1.2.

H.1.1 Experimental Hardware

As was found in chapter 2, to test TIL in a lab one would require at least one Shack-Hartmann wavefront sensor and two deformable mirrors: one with high spatial accuracy and another with large stroke for simulating a turbulent atmosphere and beam compensation, respectively. As a Shack-Hartmann wavefront sensor and two deformable mirrors (piezoelectric and MEMS) are sold as part of one kit [1], they will make up the basic setup. Additional optical components

[1] https://www.thorlabs.com/thorproduct.cfm?partnumber=AOKWT1/M

that are not part of the kit will be acquired to modify the AO setup into a TIL setup.

The schematic provided by Thorlab's is shown in fig. H.1a, where a 635 nm laser diode acts as a point source. The kit itself is made of 7 sub-assemblies, each labeled in fig. H.1b. Some tuning mirrors allow for the beam to be precisely centered, after which it passes a stage for a sample. In a typical AO setup, one may pass hot air at a specified wind speed to test the robustness of a control algorithm. The beam exits the stage for sample with a diameter of approximately ⌀10 mm. All the above forms part of subassembly 1.

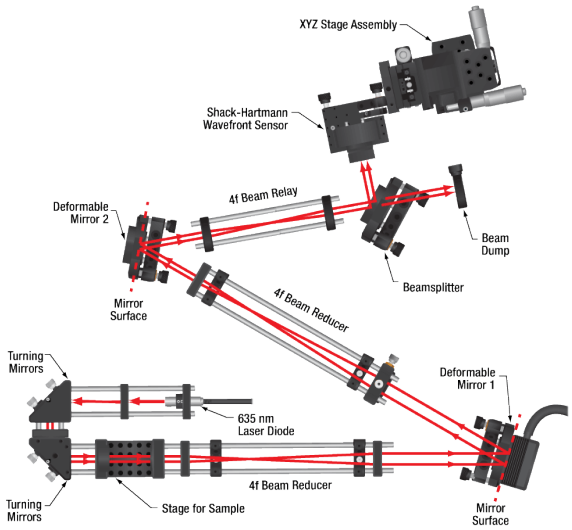

(a) Thorlabs' AOKWT1/M kit for adaptive optics. It includes two deformable mirrors (piezoelectric and MEMS) and a Shack-Hartmann wavefront sensor. Image from [33].

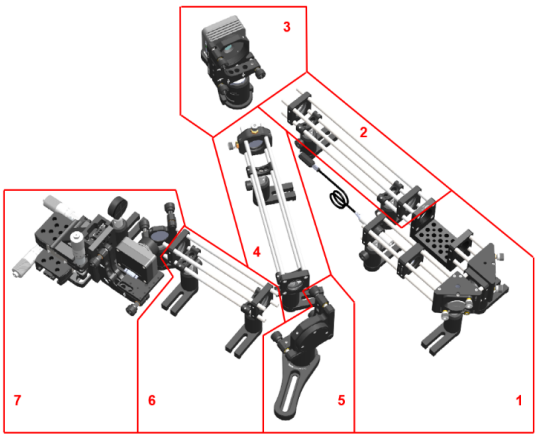

(b) Labeled AOKWT1/M kit with subassembly numbers. Image from [22].

Figure H.1: Thorlabs' AOKWT1/M kit acquired for this work.

The beam then passes through a beam expander, subassembly 2, where the beam is converted to a ⌀20 mm beam. The beam over-illuminates deformable mirror 1 (subassembly 3, ⌀14 mm pupil diameter), after which it passes through another reducer (subassembly 4, beam diameter ⌀4.2 mm). After subassembly 4, it reaches the second deformable mirror (MEMS, subassembly 5), and is reflected onto a relay (subassembly 6). In subassembly 7, the beam is split by a beam splitter, with 10 % being redirected into the Shack-Hartmann sensor, and the other 90 % onto a beam dump. With the appropriate controller, this setup is able to compensate turbulence induced in the stage for sample.

As the AO kit will be used as the basis for the TIL setup, it is important to note the key characteristics of each subassembly. For subassembly 2, the two lenses have focal lengths $f = 150\,\text{mm}$ and $f = 75\,\text{mm}$, meaning a magnification of 2:1. For subassembly 4, the ratio focal lengths are $f = 60\,\text{mm}$ and $f = 200\,\text{mm}$, meaning a ratio of 3:10. The relay optics have the same focal lengths and hence no magnification.

For the TIL setup, the available optical power to the Shack-Hartmann sensor is of concern. This is because a significantly increased path length and the scattering off the target. Hence, an optical power budget is used to estimate the available power to the Shack-Hartmann sensor in the AO kit, and it will be assumed that a similar energy input in a TIL application will result in satisfactory Signal to Noise ratio (S/N). The optical power budget is shown in table H.1. A loss value of 1 % is assumed for each lens, and each deformable mirror has an estimated loss of 2.5 %. Additionally, since the pupil of DM1 is overfilled, a only $\frac{14^2}{20^2} = 0.49$ of the energy is

retained, meaning a 51 % loss. The beam splitter incurs a further 90 % loss as it transmits that percent of light onto a beam dump.

Table H.1: Optical link budget for typical AO setup

Element Reference	**Optical Power Loss** (%)	**Power** (%)
Start	0.00	100.00
Beam expander L1	1.00	99.00
Beam expander L2	1.00	98.01
Overfilled pupil	51.00	48.02
DM1	2.50	46.82
Beam reducer L1	1.00	46.36
Beam reducer L2	1.00	45.89
DM2	2.50	44.75
Beam relay L1	1.00	44.30
Beam relay L2	1.00	43.85
Beamsplitter	90.00	4.39

As can be seen in table H.1, the available power to the Shack-Hartmann sensor is about 4.39 % from the total beam energy. As the optical power of the laser diode is 0.3 mW, 13.2 µW is incident on the Shack-Hartmann sensor [22]. This value and a similar optical budget will be used to determine the optical power required for the TIL setup.

H.1.2 Schematic Notation

When describing the tests that will be carried out, it is useful to draw schematics to display the desired layout, and the expected beam path and beam diameter. From appendix H.1.1 the type of hardware that will be used in this work is known. Figure H.2 shows the different schematics used throughout this work.

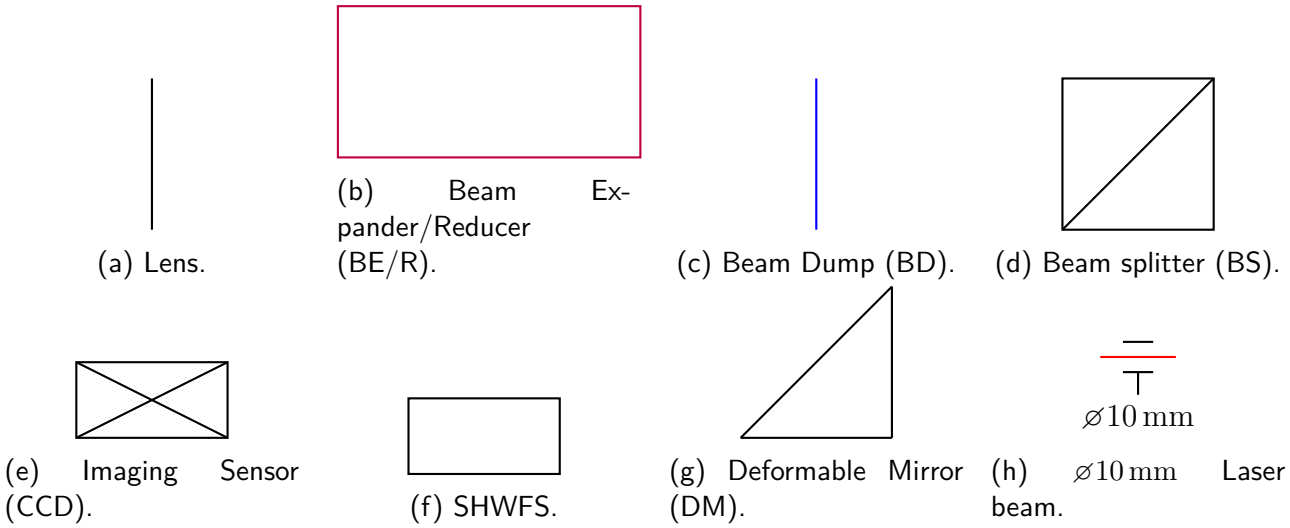

(a) Lens. (b) Beam Expander/Reducer (BE/R). (c) Beam Dump (BD). (d) Beam splitter (BS). (e) Imaging Sensor (CCD). (f) SHWFS. (g) Deformable Mirror (DM). (h) ø10 mm Laser beam.

Figure H.2: Schematic representations for different optical components used in this work.

Three things should be noted about the schematic notation used:

1. Each optical component is labeled in schematic drawings.

2. The schematic is used to illustrate the beam path, not physical dimensions or angles. Additionally, non-labeled non-optical components are used to secure the optical components.

3. The beam diameter only changes in a beam expander/reducer, or if a hard aperture is reached, for example the pupil diameter in a deformable mirror or Shack-Hartmann sensor.

With the schematic notation explained, the testing plan follows.

H.2 T-I: Individual and Subassembly Testing

The purposes of T-I tests are to ensure a proper understanding and operation of the individual components. The detailed steps are explained in appendix G. The if any of the T-I tests failed in an unexplained way, then T-II (appendix H.3) is used to fall back on the original test procedure.

H.3 T-II: Adaptive Optics Kit Testing

T-II is an optional test that can be used to debug the setup. As the manufacturer sells the kit to be used in this default configuration, one can find more support with this configuration. If it is needed, this test should be successfully performed, and T-I redone with the necessary modifications.

H.4 T-III: TIL Testing

Once the correct functioning of the forward path is established, it is now possible to test the return path. Tests in T-III will build the setup required to test the TIL application, including the limitations typically associated with TIL: speckle pattern and weak return signal. The testing begins with a basic target reflection test, where a target is illuminated and the return signal is registered in appendix H.4.1. In case the return signal is too weak, a higher power laser is used to increase the return signal intensity in appendix H.4.2. In appendices H.4.3 and H.4.4, the return signal is configured such that is registered by both an imaging sensor and the wavefront senor.

H.4.1 T-III.1: Test Target Reflection

In T-III.1, the target surface will be introduced to the setup, and the return signal will be imaged using an imaging sensor. Figure H.3a shows the forward path taken by the beam, up until it reaches the target surface (green, TAR). The key modifications from T-I.4 is the addition of L2, and which is used to focus the light reflected from DM2 to the target surface. To observe the point spread function on the surface, CCD2 is placed on the focal plane of L2, same as TAR.

Next, fig. H.3b shows the return path of the laser beam once it reflects off the target. Importantly, even though it is scattered off the target surface and not reflected, as an approximation

a reflection can be assumed, with the optical power loss to be estimated in subsequent tests. The return beam reflects and refracts off the optical components in the setup, eventually reflecting off BS1 and into L1 and into CCD1. CCD1 is expected to provide images of the return speckle.

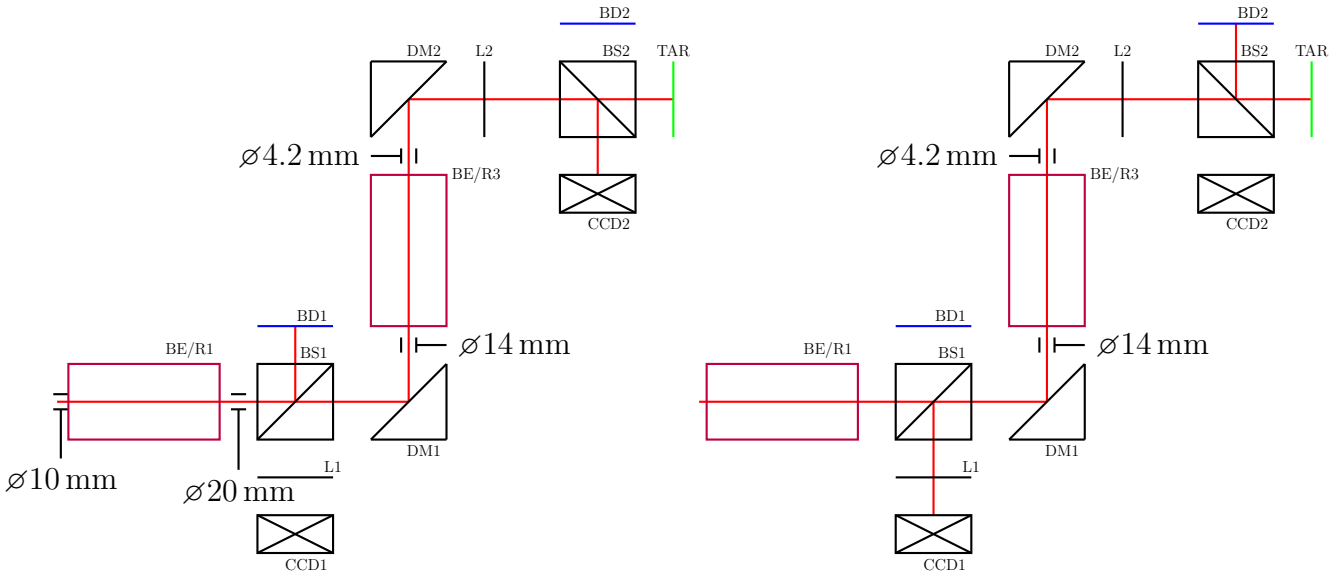

(a) Forward path of the laser beam.

(b) Return path of the laser beam.

Figure H.3: T-III.1 and T-III.2 Test Setup

A successful T-III.1 test would mean that CCD2 can faithful observe the projection pattern on TAR. Additionally, CCD1 should obtain a speckle return beam image, however the return field intensity is unknown (can be found using a power budget) and the sensitivity of the sensor is unknown. Hence, for this test, a return speckle may not be visible due to lacking optical power, which will be addressed in the following test. In the case the speckle pattern is observed, T-III.2 may not be necessary.

H.4.2 T-III.2: Test High Power Laser

Due to the significant power losses expected in T-III.1, it is conceivable that there is not sufficient power incident on the sensor to capture images. While exposure time could be increased, eventually a real time system would be preferred, and hence a higher optical power is critical. If necessary, T-III.2 will be used to install a high power laser (130 mW[2] versus 30 mW), which will increase the return field intensity and hence CCD1 should now register the return speckle pattern. Additionally, since the beam is square with dimensions $3.8 \times 3.8\text{mm}^2$, there will be no loss due to overfilled DM1.

A successful T-III.2 test would result in a speckle field return pattern being observed by CCD1.

H.4.3 T-III.3: Test Beamsplitter

T-III.3 will add beam splitter BS3 at the end of the return path, such that the Shack-Hartmann sensor can be used at the same time as CCD1. This will be beneficial as one will be able to observe the return field speckle, and also how the controller adapts to the speckle. The return path is shown in fig. H.4, with BE/R1 ommitted for compactness.

[2] https://www.lasertack.com/en/130mw-660nm-diode-laser-hl65051dg-hl6545mg/

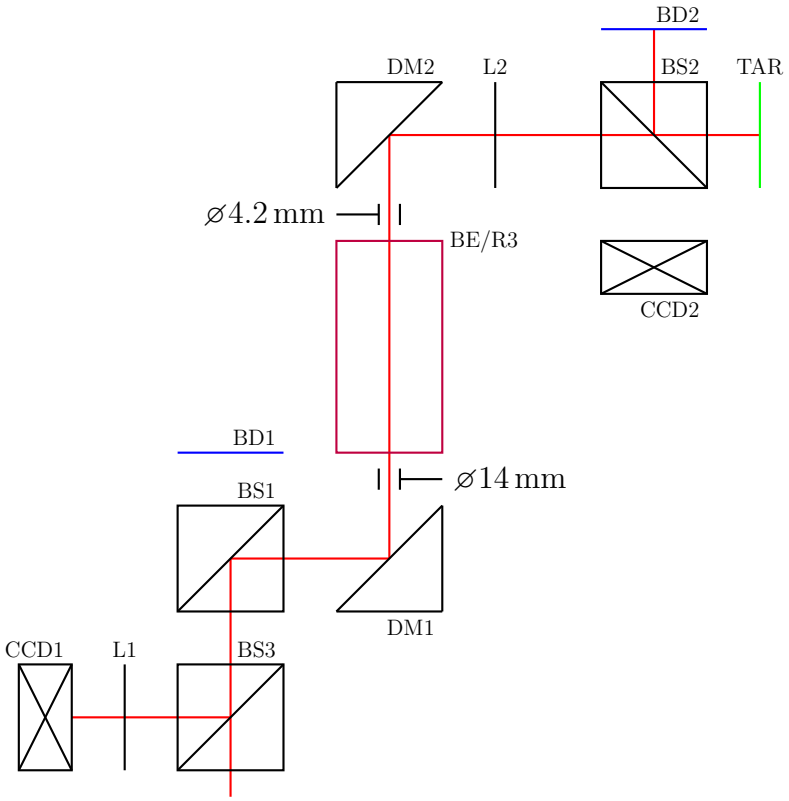

Figure H.4: T-III.3 Test setup just showing return path.

A successful T-III.3 test would result in a speckle field return pattern being observed by CCD1.

H.4.4 T-III.4: Complete Setup

Finally, the setup can be completed by adding the wavefront sensor at the transmission end of BS3. Because the beam diameter is ⌀14 mm, it needs to be converted to a smaller size for the wavefront sensor with a beam reducer (BE/R2). The complete schematic is shown in fig. H.5.

L2 can be slid back an forth, increasing the size of the projection on the target. This will vary the amount of speckle to the return optics, from which the controllers can optimize the beam projection, counteracting the introduced defocus.

The optical budget for the TIL system shown in fig. H.5 can be calculated and is shown in table H.2. The same optical power loss values are assumed for optical components as in table H.1. The power loss due to the target can be estimated using eq. (H.1), where D is the diameter of the optic (25.4 mm) and f is the focal length of the lens (200 mm). The derivation is shown in appendix F.

$$\begin{aligned} \%P_{loss} &= 100 \cdot (1 - \frac{D^2}{4 \cdot f^2}) \\ &\approx 99.60\% \end{aligned} \tag{H.1}$$

As calculated in table H.2, only 0.2314 % of the optical power remains for the Shack-Hartmann sensor. With an optical power of 0.3 mW, this would mean only 0.1 µW of optical power incident on the wavefront sensor. This is much less than the 13.2 µW, and hence is expected not to work. A 130 mW[2] laser would have 38.3 µW incident on the sensor, more than enough

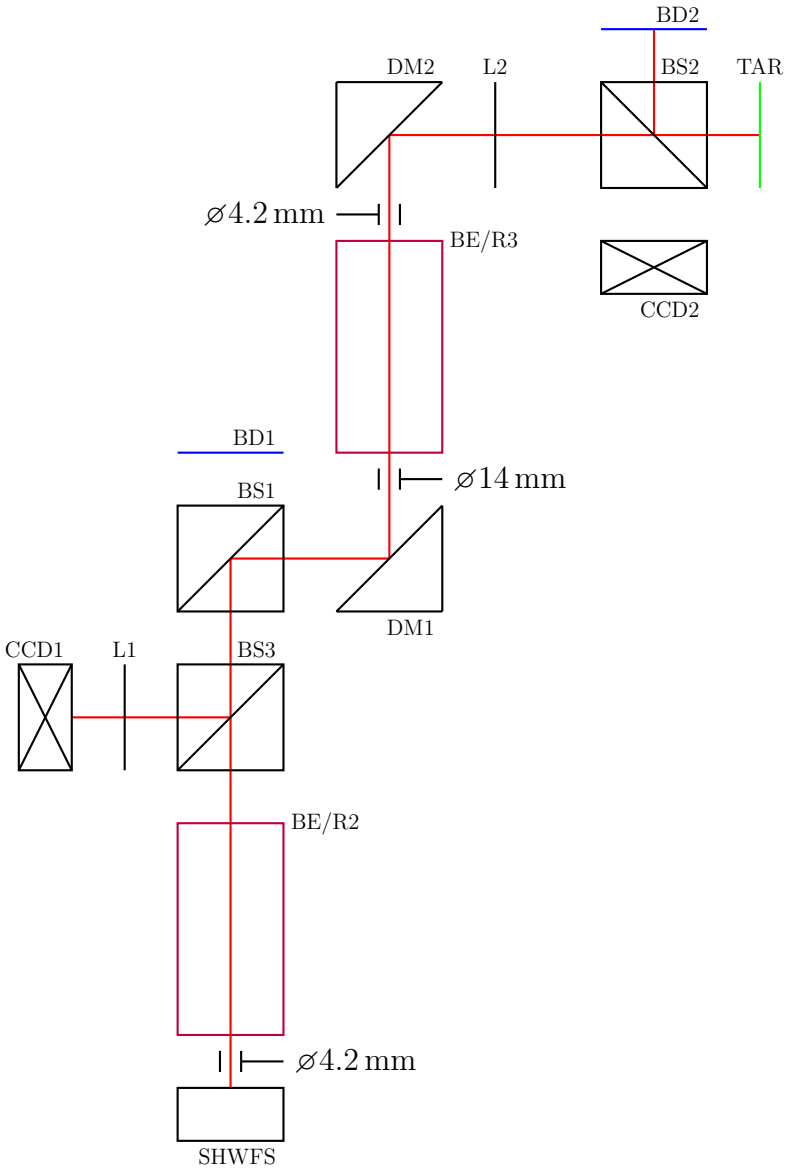

Figure H.5: T-III.4 Test setup

for a measurement. As the optical power of such a laser can be tuned[2], it will be used for the final setup.

The ray-traced final setup is shown in fig. H.6. The physical dimensions are correct, however the spacing between components are their angles are not realistic.

With the physical setup complete, it is now possible to test the various controller designs. The setup will require minimal changes to evaluate each controller.

H.5 T-IV: Controller Evaluations

With the experimental setup defined, it is now time to test the controllers. Firstly, the performance metric on which the controllers will be evaluated on will be defined in appendix H.5.1. Next, the four different controllers are polychromatic speckle (appendix H.5.2), speckle average (appendix H.5.3), temporal speckle (appendix H.5.4) and spatial speckle (appendix H.5.5). Each sub-section also covers the modifications required to test the controller.

Table H.2: Optical link budget for first TIL test. Note that all figures are rounded to 2 decimal points, except for the last entry because it presents the final results.

Element Reference	**Optical Power Loss** (%)	**Power** (%)
Start	0.00	100.00
Beam expander L1	1.00	99.00
Beam expander L2	1.00	98.01
BS1	50.00	49.01
Overfilled pupil	51.00	24.01
DM1	2.50	23.41
Beam expander L1	1.00	23.18
Beam expander L2	1.00	22.95
DM2	2.50	22.37
L2	1.00	22.15
BS2	19.00	19.93
Target	99.60	0.08
BS2	10.00	0.07
L2	1.00	0.07
DM2	2.50	0.07
Beam expander L2	1.00	0.07
Beam expander L1	1.00	0.07
DM1	2.50	0.07
BS1	50.00	0.03
BS3	10.00	0.03
Beam expander L1	1.00	0.03
Beam expander L2	1.00	0.0294

H.5.1 Performance Metric

The choice of performance metric is critical to successful testing. The two performance metric considered by [17] state Strehl ratio and power-in-a-bucket (PIB), which present the two major options. Strehl ratio is defined as the normalized peak intensity of the point spread function of an optical system [34]. What this means is that Strehl ratio is a ratio of true peak intensity over the ideal peak intensity. PIB is the total energy contained within a specified radius in the target plane.

While Strehl ratio may be useful in conventional optics applications, a better metric for TIL is power-in-a-bucket, as ultimately energy efficiency is the main objective in TIL (power delivery). Hence, the figure of merit chosen for this work is power in a bucket as collected by CCD2, as this represents a value proportional to the power delivery to the target.

There is one issue, however: PIB is an instantaneous measure. Different controllers may have different PIBs temporally. Hence, 4 metrics related to PIB will be used to evaluate the controllers:

1. Maximum PIB: Maximum PIB registered
2. Steady state PIB variance
3. Rise time: time required to first reach 90 % of the maximum PIB value from 10 %
4. Settling time: time required to stay with 5 % of maximal PIB

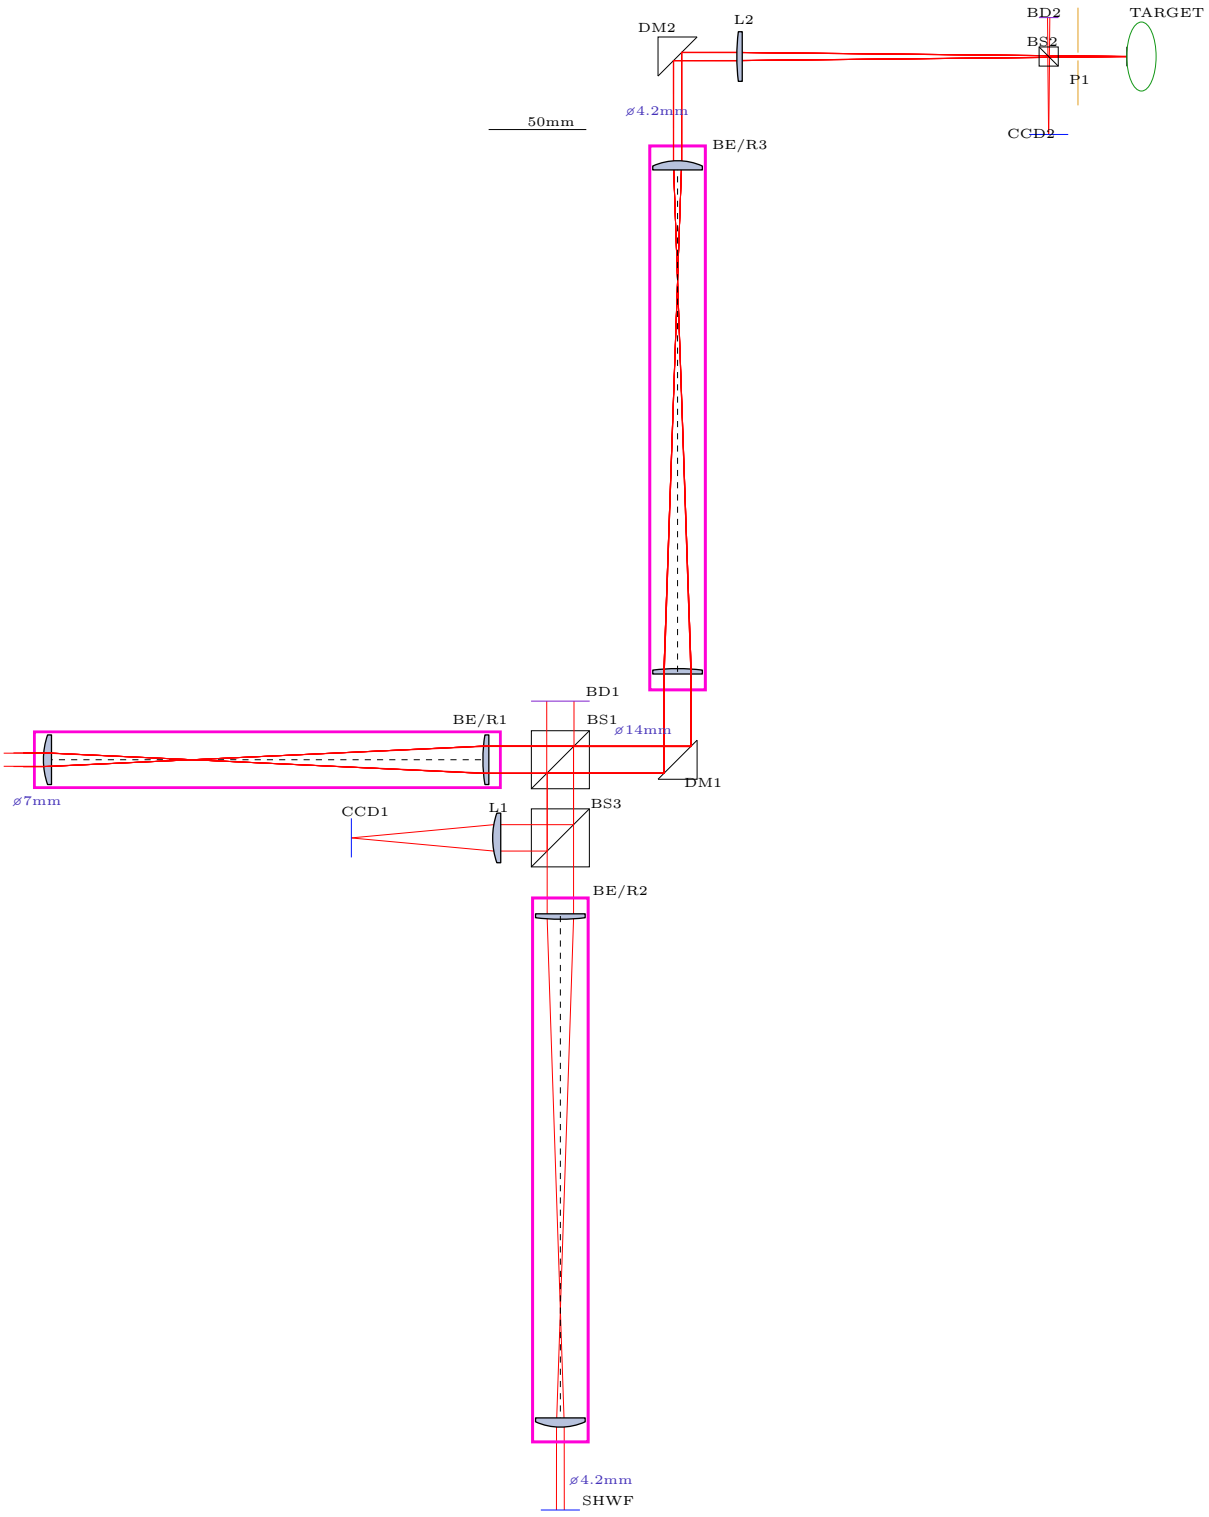

Figure H.6: Proposed final setup, with a $50\,\mathrm{mm}$ scale shown at the top. Red lines indicate raytraced path of beam, assuming perfect reflection at the target.

These metrics will be compared and contrasted for the various controllers to evaluate their performance.

H.5.2 T-IV.1: Test Polychromatic Speckle

The testing of the polychromatic speckle controller can be performed with no modifications to the physical setup. As this approach relies on the incoherence of beam, no special modifications need to be done on the software side of things, and the controller can just seek to obtain a planar wavefront at DM1.

The sequence flow diagram is shown in fig. H.7, with the block definitions present in the caption. Importantly, the control software comprises of two loops running in parallel with each other: the main control loop and the measurement loop. The measurement loop is responsible for taking an image and then setting a flag for data availability. The main control loop can then access the collected data, and update its information accordingly. As the measurement frequency is expected to be much lower than the control frequency, a separate loop exists where room is left extrapolate the controller while no new measurement is available. This loop runs at a constant frequency and hence has an adaptive wait function that pauses the execution of the program until the correct time to ensure a constant average frequency.

The main control loop will make use of a Kalman Filter, which combines measurement information with a dynamic model of the system to make an optimal state estimation [35].

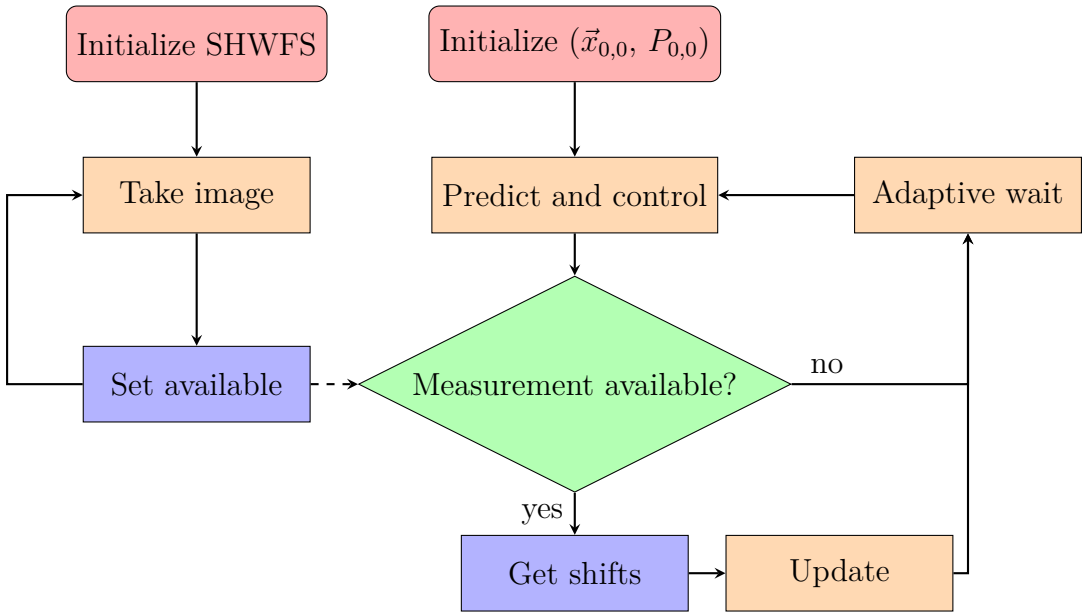

Figure H.7: Sequence diagram for polychromatic speckle with state estimation using a Kalman Filter. Red boxes indicate initialization steps, whereas orange ones processes, and blue synchronized processes, and green decisions. Synchronized processes can only run one at a time. Two cores work in parallel to close the loop (commonly referred to as multiprocessing). The first core initializes the Shack-Hartmann Wavefront Sensor, and then continuously takes and image, waits, and when ready sets the available flag. The other loop initializes with an estimate for the wavefront shifts and covariances, and then continuously predicts and controls the deformable mirror in a constant frequency loop. Whenever a measurement becomes available, the spot shift information is used to update the estimated shifts optimally.

Kalman Filter

To perform better, the control loop make use of a state estimator, than can predict the state of the system, and combine measurements with state dynamics optimally [35]. The basis of the

Kalman filter are two equations, the update and measurement equations. The general form of the update equation is shown in eq. (H.2), where the state variable x_{k-1} is propagated forward into $\vec{x}_k$ through the use of the state transition matrix (F_k) and the contribution of the control vector $\vec{u}_k$ through the matrix B_k. Additionally, due to some unmodeled effects the process noise term $\vec{w}_k$ is present, with distribution $\vec{w}_k \sim \mathcal{N}\left(\vec{0}, Q_k\right)$.

$$\vec{x}_k = F_k \cdot \vec{x}_{k-1} + B_k \cdot \vec{u}_k + \vec{w}_k \tag{H.2}$$

The second equation defining the Kalman filter is the measurement equation, as shown in its general form in eq. (H.3). The measurements performed by a sensor that produces $\vec{z}_k$, related to $\vec{x}_k$ through the matrix H_k. As in a real system there is some measurement noise, the term $\vec{v}_k$ is used to model it. It has distribution $\vec{v}_k \sim \mathcal{N}\left(\vec{0}, R_k\right)$.

$$\vec{z}_k = H_k \cdot \vec{x}_k + \vec{v}_k \tag{H.3}$$

There are 3 phases in a Kalman filter: initialization, prediction and updating. Initialization is performed when starting the system up, whereas predict is performed whenever necessary. Updating should be done as soon as a new measurement becomes available. The equations used to perform the initialization, prediction and update steps are shown in table H.3. E signifies the expected value and Var signifies the variance.

Step	**Process**
Initialization	$\hat{\vec{x}}_{0,0} = E\left[\vec{x}_0\right]$ $P_{0,0} = \mathrm{Var}\left(\vec{x}_0\right)$
Predict	$\hat{\vec{x}}_{k,k-1} = F_k \cdot \hat{\vec{x}}_{k-1,,k-1} + B_k \cdot \vec{u}_k$ $P_{k,k-1} = F_k \cdot P_{k-1,k-1} \cdot F_k^T + Q_k$
Update	$\tilde{\vec{y}}_k = \vec{z}_k - H_k \cdot \hat{\vec{x}}_{k,k-1}$ $S_k = H_k \cdot P_{k,k-1} \cdot H_k^T + R_k$ $K_k = P_{k,k-1} \cdot H_k^T \cdot S_k^{-1}$ $\hat{\vec{x}}_{k,k} = \hat{\vec{x}}_{k,k-1} + K_k \cdot \tilde{\vec{y}}_k$ $P_{k,k} = (I - K_k \cdot H_k) \cdot P_{k,k-1}$ $\tilde{\vec{y}}_{k,k} = \vec{z}_k - H_k \cdot \hat{\vec{x}}_{k,k}$

Table H.3: Summary of equations involved in Kalman Filter.

Polychromatic Algorithms

Within the context of the polychormatic speckle controller, the main control loop will perform the state estimation. Choosing the state vector to represent the spot shifts as measured by the wavefront sensor, one readily obtains eqs. (H.4) and (H.5).

$$F_k = I \tag{H.4}$$

$$H_k = I \tag{H.5}$$

Additionally, $B_k = \Phi^+$, some inverse[3] of the poke matrix. The covariance matrices can be calibrated in testing for optimal performance, by balancing the reliance on the predicted dynamics

[3] As Φ is not a square matrix in general, a pseudoinverse that satisfies the property $\Phi^+\Phi \approx I$ is sufficient. The pseudoinverse is non-unique and can be constructed in various ways.

versus the confidence in the measurements. The input to the deformable mirror is then shown in eq. (H.6), with $\vec{x}$ being the best estimate for the state, and $\vec{x}_{ref}$ the desired state (typically the zero vector).

$$\vec{u} = \Phi^{+} \cdot (\vec{x}_{ref} - \vec{x}) \tag{H.6}$$

With these equations it should be possible to compensate the beacon for optimal control.

H.5.3 T-IV.2: Test Speckle Average

The testing of the speckle average controller can be performed with no modifications to the physical setup. Dependent on the amount of speckle present in the image, this approach will use varying exposure time and slight tip/tilt movements to vary the speckle pattern. With enough speckle patterns realized, the wavefront sensor should be able to correct for the simulated atmospheric turbulence and compensate the beacon for optimal control. The control scheme is the same as in appendix H.5.2, except that the camera is initialized with a higher exposure time, and the target is rotated.

H.5.4 T-IV.3: Test Temporal Speckle

The testing of the temporal speckle controller can be performed with a photodetector in place of CCD1. Alternatively, for simplicity one may choose to sum the intensity of the image taken by CCD1, and use that as a metric instead. The sequence flow diagram is shown in fig. H.8. Note the simplicity of the algorithm as compared to fig. H.7, which is advantageous for software development.

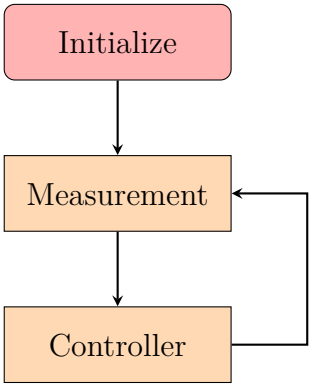

Figure H.8: Sequence diagram for temporal speckle. A measurement is performed, and is immediately fed into the controller.

The necessary hardware and software is first initialized. Once initialized, the first measurement is obtained, which is fed into a controller to minimize the objective function. The form of such as problem is as was state in eq. (2.22), or repeated below.

$$\min_{\vec{u}} f(\vec{u}) \tag{2.22}$$

Multiple methods exist to minimize such a function. These will be explored during testing, and the best performing algorithm will be chosen.

By rotating the target, theoretically enough speckle realizations will be present such that the beacon beam can be optimized for atmospheric turbulence.

H.5.5 T-IV.4: Test Spatial Speckle

The testing of the spatial speckle controller can be performed with no modifications to the physical setup. The same program outline as in appendix H.5.4 (fig. H.8).

Instead of using the SHWFS, CCD1 will be used to test the spatial speckle approach, which should theoretically compensate the beacon for optimal control.

Appendix I

Software Turbulence Verification

Parts of this work required simulating a wavefront propagating through a turbulent atmosphere. For $C_n^2 = 1 \times 10^{-13}\,\mathrm{m}^{-2/3}$, eq. (2.8) is plotted in red in fig. I.1, showing the analytical solution.

The simulated components where generated by creating a beam with 10 cm width. This was propagated to various distances through the turbulent atmosphere and a vacuum. The power in the bucket for 10 cm was extracted for both these wavefronts, and the ratio was stored. This process was repeated for 20 simulations. Each simulation is shown in fig. I.1 as a light and dashed gray line, with the average for each distance.

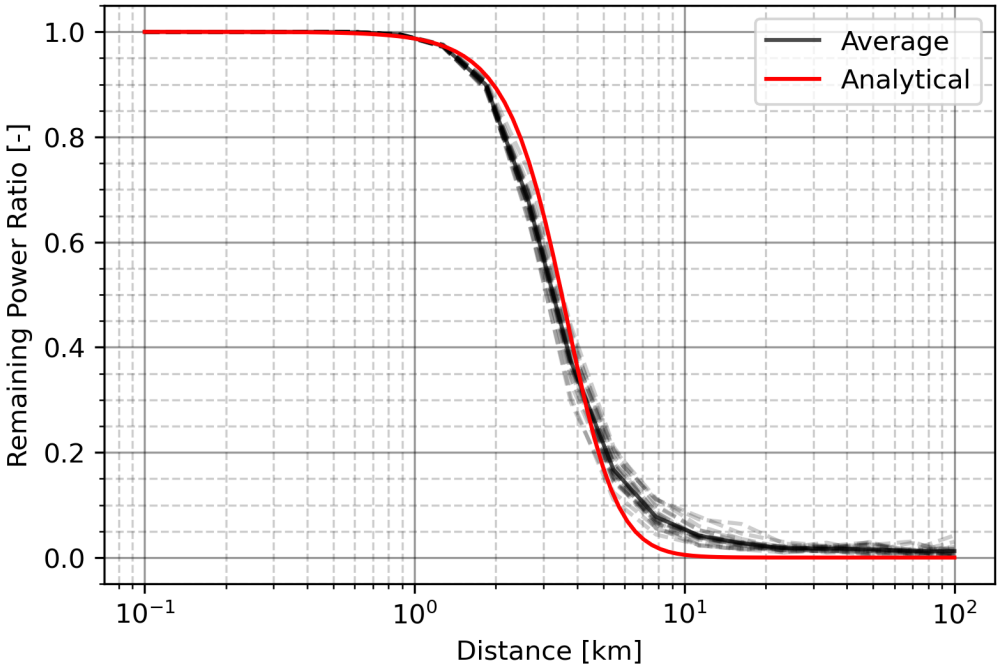

Figure I.1: Figure showing analytical versus simulated remaining power ratio for various distances at a $C_n^2 = 1 \times 10^{-13}\,\mathrm{m}^{-2/3}$ turbulence strength.

Figure I.1 shows a clear correlation between the analytical solution and the average of the simulations, indicating that the simulation method conforms to the expected statistics of the atmosphere.