Design and Setup of a Post-Compressor Adaptive Optics Loop at PHELIX

PhD Thesis by Jonas B. Ohland

Zur Erlangung des Grades eines Doktors der Naturwissenschaften (Dr. rer. nat.)
Genehmigte Dissertation von Jonas Benjamin Ohland aus Rüsselsheim
Tag der Einreichung: 07. Juni 2022, Tag der Prüfung: 13. Juli 2022

1. Gutachten: Prof. Dr. V. Bagnoud
2. Gutachten: Prof. Dr. M. Roth
Darmstadt

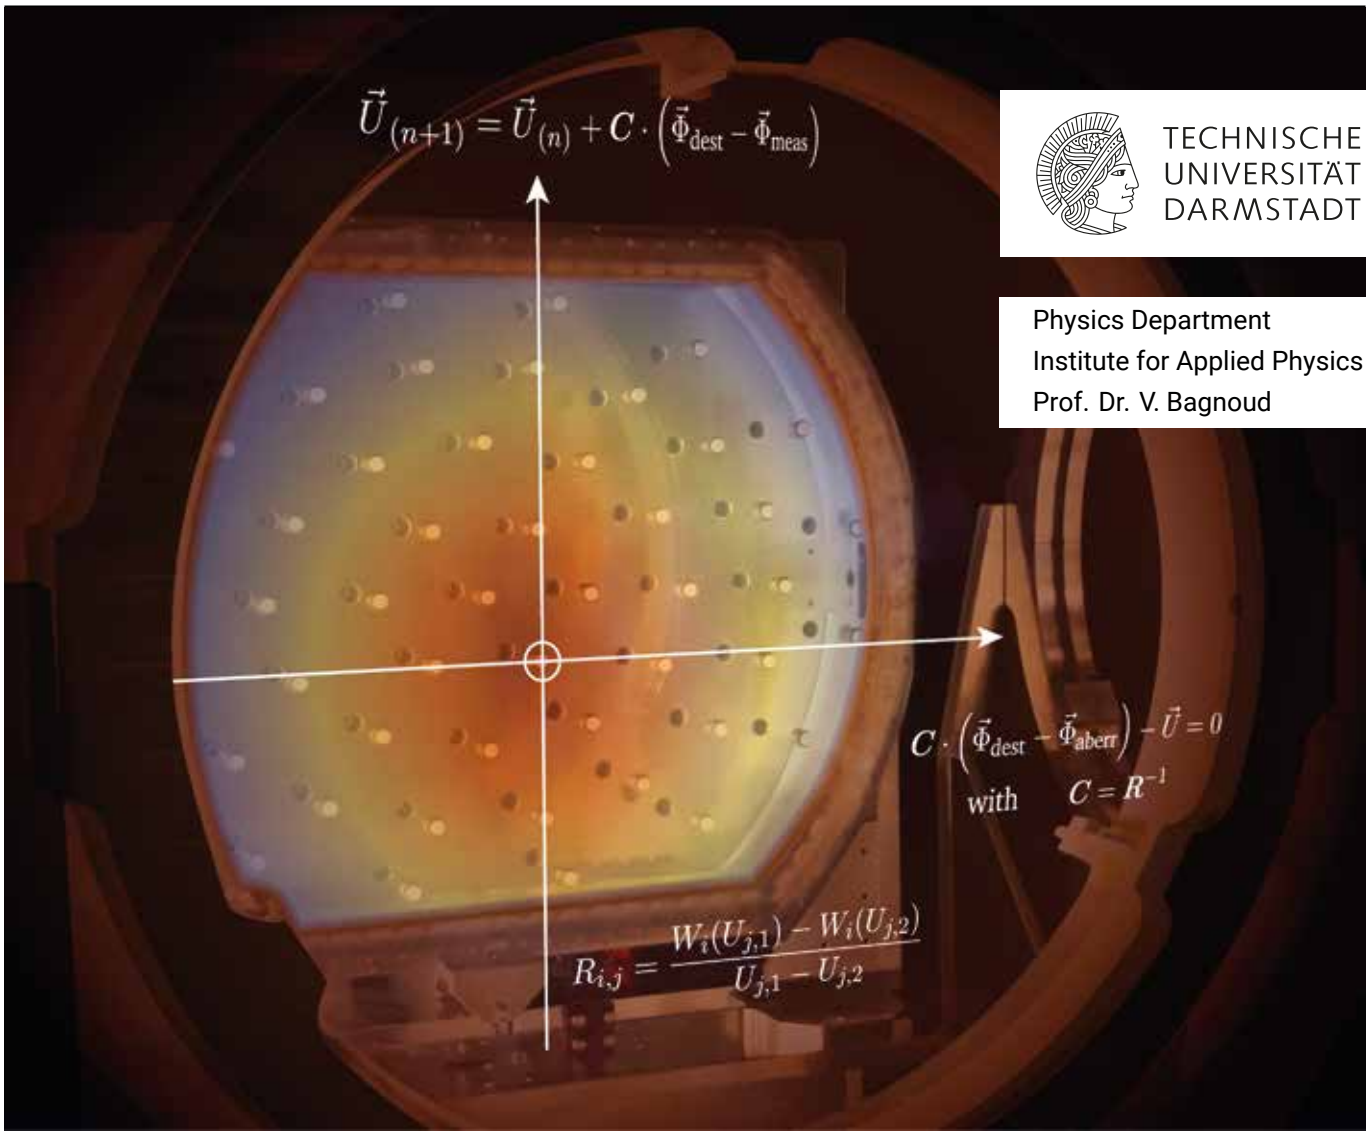

TECHNISCHE UNIVERSITÄT DARMSTADT

Physics Department
Institute for Applied Physics
Prof. Dr. V. Bagnoud

Design and Setup of a Post-Compressor Adaptive Optics Loop at PHELIX
PhD Thesis by Jonas B. Ohland

Submitted doctoral thesis by Jonas Benjamin Ohland

1. Review: Prof. Dr. V. Bagnoud
2. Review: Prof. Dr. M. Roth
Examiner: Prof. Dr. B. Drossel
Prof. Dr. T. Halfmann

Date of submission: 07. Juni 2022
Date of thesis defense: 13. Juli 2022

Darmstadt

Bitte zitieren Sie dieses Dokument als:
URN: urn:nbn:de:tuda-tuprints-219951
URL: https://tuprints.ulb.tu-darmstadt.de/21995

Dieses Dokument wird bereitgestellt von tuprints,
E-Publishing-Service der TU Darmstadt
http://tuprints.ulb.tu-darmstadt.de
tuprints@ulb.tu-darmstadt.de

Erklärungen laut Promotionsordnung

§8 Abs. 1 lit. c PromO

Ich versichere hiermit, dass die elektronische Version meiner Dissertation mit der schriftlichen Version übereinstimmt.

§8 Abs. 1 lit. d PromO

Ich versichere hiermit, dass zu einem vorherigen Zeitpunkt noch keine Promotion versucht wurde. In diesem Fall sind nähere Angaben über Zeitpunkt, Hochschule, Dissertationsthema und Ergebnis dieses Versuchs mitzuteilen.

§9 Abs. 1 PromO

Ich versichere hiermit, dass die vorliegende Dissertation selbstständig und nur unter Verwendung der angegebenen Quellen verfasst wurde.

§9 Abs. 2 PromO

Die Arbeit hat bisher noch nicht zu Prüfungszwecken gedient.

Darmstadt, 07. Juni 2022

J. B. Ohland

Kurzfassung

Das Ziel dieser Arbeit war es, erweitere Kontrolle über den Fokalpunkt des PHELIX Lasers zu erhalten. Dies sollte durch die Implementierung eine Regelschleife mit adaptiver Optik nach dem Pulskompressor des Lasers erreicht werden. Hiermit soll eine höhere Spitzenintensität auf dem Target erreichbar werden, welche eine kritische Rolle in Experimenten mit Laser-Plasma-Wechselwirkung spielt.
Im Rahmen dieser Arbeit habe ich einen ultrakompakten Wellenfrontsensor implementiert, mit dem die Wellenfront über die gesamte 28 cm messende Apertur von PHELIX auf einer Fläche von 2×1 Metern bei jeder möglichen Pulsenergie gemessen werden kann. Der Sensor wurde anschließend genutzt, um eine geschlossene Regelschleife mit einem kommerziellen verformbaren Spiegel zu betreiben.
Des Weiteren habe ich eine Routine zur Kalibration des Sensors entwickelt, welche es ermöglicht, die optimale erreichbare Wellenfront in der Targetkammer zu identifizieren. Mit dieser Routine war ich in der Lage, die Intensität auf dem Target zu verdreifachen, was einer Intensität von etwa $1{,}4{\cdot}10^{21}$ W/cm^2 entspricht. Dies war das erste Mal, dass bei PHELIX Intensitäten über 10^{21} W/cm^2 systematisch erreichbar waren.
Um einen qualitativen Nachweis für diese Intensitäten zu liefern, führten wir innerhalb der Gruppe ein Experiment durch, bei dem wir die Transmission des Laserpulses durch Polystyrol-Folien unterschiedlicher Dicke maßen. Das Ergebnis bestätigte die erfolgreiche Optimierung der Spitzenintensität.

Die Ergebnisse dieser Arbeit habe ich in zwei Publikationen veröffentlicht [1, 2].

Abstract

This work aimed to provide advanced control over the focal spot of the PHELIX laser by the means of implementing a post-compressor adaptive optics loop. The ultimate goal was to increase the peak intensity on the target, which is critical for a lot of experiments that are based on laser-plasma interaction.
Over the course of this thesis, I implemented an ultra-compact wavefront sensor, which could successfully measure the wavefront over the full 28 cm aperture of the PHELIX laser on a footprint of 2×1 meters for the full energy range of PHELIX. This sensor was used to run a closed control loop with a commercial deformable mirror.
Furthermore, I developed a calibration routine for the sensor, which enabled finding the optimum achivable wavefront at the focal spot in the target chamber. Together with this routine, I was able to triple the intensity on target up to $1.4{\cdot}10^{21}$ W/cm^2. This was the first time where an intensity of more than 10^{21} W/cm^2 could be reached systematically at PHELIX.
Lastly, we performed an experiment in the group which provided a qualitative prove of the on-shot intensity gain by observing the transmission of the laser pulse through polystyrene foils of various thicknesses. The result verified the successful optimization of the peak intensity.

I published the results of this work in two papers [1, 2].

Education

Master of Science (Physics)

05.04.2016 - 31.07.2018

Technical University of Darmstadt
Karolinenplatz 5
64289 Darmstadt
Germany

Thesis: Improvement of the Adaptive Optics at PHELIX for Operation at higher Repetition Rates and Investigation of Laguerre-Gaussian-like Modes
Grade: 1.0 (very good)
Final grade: 1.37 (very good)

Bachelor of Science (Physics)

01.10.2012 - 04.04.2016

Technical University of Darmstadt
Karolinenplatz 5
64289 Darmstadt
Germany

Thesis: Programming of an Evaluation-Software for VISAR-data
Grade: 1.3 (very good)
Final grade: 1.90 (very good)

General University Entrance Qualification

01.06.2012

Graf-Stauffenberg-Gymnasium
Bürgermeister-Lauck-Straße 24
65439 Flörsheim am Main
Germany

Specialized courses:

- Physics (written exam), 15 pt (very good)
- Mathematics (written exam), 15 pt (very good)
- Chemistry (voluntary, oral exam), 14 pt (very good)

Final grade: 1.40 (very good)

Contents

Kurzfassung v

Abstract vii

Education ix

1. Introduction 1

2. Adaptive Optics 3
2.1. History . . . 3
2.2. Wavefront Control . . . 6

3. Composition of Aberrations in High-Intensity Laser Systems 11
3.1. The Architecture of Modern High-Intensity Laser Systems . . . 11
3.2. Description of Aberrations . . . 19
3.3. Static Aberrations . . . 24
3.4. Dynamic Aberrations . . . 25
3.5. Chromatic Aberrations and Spatio Temporal Couplings . . . 28

4. Requirements on Adaptive Optics in High-Intensity Laser Systems 31
4.1. Critical Locations in CPA Lasers . . . 31
4.2. Pre-Compressor Challenges . . . 32
4.3. Post-Compressor Challenges . . . 35
4.4. Single-Shot-Related Challenges . . . 36

5. Implementation of Post-Compressor Adaptive Optics at PHELIX 41
5.1. Requirements . . . 41
5.2. Design of the Petawatt Target Area Sensor . . . 44
5.3. Calibration Routine . . . 70
5.4. Estimation of the On-Shot Intensity . . . 79

6. Experimental Proof of the On-Shot Intensity Gain 85
6.1. Concept and Expectations . . . 85
6.2. Experimental Setup . . . 91
6.3. Results and Discussion . . . 92
6.4. Possible Quantitative Intensity Proof using Optical Tunnel Ionization . . . 94

7. Summary, Outlook and Acknowledgments 97

8. List of Abbreviations and Bibliography 103

A. Available Technology **115**
A.1. Wavefront Metrology 115
A.2. Deformable Mirrors 126
A.3. Control Units 130

B. Diffraction Theory **137**
B.1. Electromagnetic Waves 137
B.2. Scalar Diffraction Integrals 138
B.3. Vectorial Diffraction Integrals 142
B.4. Diffraction of Polychromatic Light 148
B.5. Numerical Evaluation 148

1. Introduction

Since the invention of the laser in 1960 [3], the science community has strived to build pulsed laser systems with increasingly high peak intensities [4], which proved to be indispensable tools for a variety of applications. These include planetary science [5], where dense states of matter could be created by laser-induced shock compression in the laboratory, particle acceleration [6], where the laser induces strong electrical fields inside a plasma, which is able to accelerate ions or electrons on extremely short scales, or the use of secondary sources, such as laser-generated X-rays [7] or neutrons [8]. Furthermore, the recent advances in Inertial Confinement Fusion [9] as a potential green energy source made clear that a combination of these applications may be of economic interest in the future.
Some of these applications require high peak intensities, but at the very least, all applications rely on well known laser parameters. This includes the temporal characteristics of the pulses, such as the duration and contrast, and spatial beam properties, which ultimately determine how well the beam can be focused. From a facility point of view, controlling these parameters becomes very challenging. This is in part due to pulse powers exceeding terawatts or even petawatts, enforcing the use of large beam diameters in order to stay below the damage threshold of the optical components. After the final pulse amplifiers, beam diameters of tens of centimeters are commonly encountered [4], reaching up to over half a meter in some cases.
These beam sizes are much more susceptible for spatial phase defects, accumulating from alignment errors, heated amplifiers, air turbulence, optical imperfections and alike. These WaveFront (WF) aberrations ultimately deteriorate the focal spot on the target in an uncontrolled manner, which has a negative impact on the quality of experimental results and the performance of other applications. Specifically, WF aberrations reduce the so-called Strehl ratio, which is the ratio between the real peak intensity and the theoretical limit that could be reached in the absence of aberrations. This is unacceptable for applications which require the highest possible intensities on the target.
For that reason, a wide range of techniques to minimize and compensate these WF aberrations has been established in the field. This includes passive approaches on one hand, such as beam housing in order to reduce air turbulence, damped optical mounts to mitigate the impact of external vibrations and relay imaging, which counteracts the diffraction of the laser beam over the course of the beam transport. Active approaches, on the other hand, rely on moving or deforming elements in the beam path and strive to compensate for existing aberrations.
The most prominent approach here is Adaptive Optics (AO). Here, an active component, e.g. a Deformable Mirror (DM) is able to change the optical path over the aperture of the beam to counteract WF defects. The control happens using a WaveFront Sensor (WFS) that records the WF after the beam has passed the active component and generates control values, which are used to form a feedback loop. While AO can serve different purposes in High Intensity (HI) lasers, e.g. real-time beam stabilization or focal spot optimization [10], the possibilities are strongly tied to both the architecture of the AO system and the laser chain itself.
A great example for the latter is the Chirped Pulse Amplification (CPA) concept [11], awarded with the Nobel Prize in 2018, which strongly impacts the way AO has to be used in HI lasers. The CPA concept enables the highest peak powers that are available today by stretching the pulse temporally, which leads to a reduced power, amplifying it afterwards and recompressing the pulse under vacuum right before the target area. As the compression is affected by WF defects in the beam [12], the WF should be corrected

prior to the pulse compression stage (pre-compressor loop). In turn, all aberrations that are generated afterwards and negatively impact the quality of the focal spot should be corrected with a separate loop (post-compressor loop) [13]. As all modern HI lasers rely on CPA, post-compressor AO loops are generally needed. The construction of these, however, holds a particular set of challenges. This is mainly due to the fact that the surface of the active component (mostly DMs in this case) has to be imaged onto a WFS close to the target area. While imaging a beam with low power can be done with lenses, the high peak power of the compressed pulses make the use of conventional refractive optics impossible due to damage thresholds and nonlinear optical effects. Instead, reflective telescopes have to be used.

The main goal of this work was to implement such a post-compressor AO loop at the Petawatt High-Energy Laser for heavy Ion eXperiments (PHELIX) at the GSI Helmholtzzentrum für Schwerionenforschung GmbH (GSI) in Darmstadt, Germany, which is able to deliver pulse energies up to 200 J within 500 fs at a central WaveLength (WL) of 1053 nm. Prior to this work, PHELIX featured a pre-compressor loop only. The focal spot on the target, however, still suffered from aberrations due to compressor imperfections and defects of the focusing optic.

In this thesis, I will describe AO in general and then narrow down on the particularities of post-compressor loops step by step, which will lay the foundation for the implementation, test and experimental proof of the AO loop in the target area of PHELIX.
I will start by giving an introduction into AO in chapter 2. Here, I will summarize its history and discuss the core concepts of typical AO loops. In order to narrow down to HI lasers, I will discuss the core concepts of these laser systems in chapter 3, where I also describe the architecture of PHELIX. From this point, I will deduct the sources and characteristics of aberrations in HI lasers. This acts as a foundation for chapter 4, where I derive the requirements on AO in these laser systems. Here, the differences to AO in other fields, e.g. astronomy, will become clear.
With this background, I will finally present the heart of this work in chapter 5: the implementation of the post-compressor AO loop at PHELIX. I will list the requirements for this system, present the thought process and the associated design decisions and describe the technical implementation in great detail. This includes the WFS, the alignment scheme of its Off-Axis Parabolic mirror (OAP) telescope, as well as the calibration device and -routine. Furthermore, I will present test results of the setup, which I recorded at low pulse energies.
Lastly, I discuss the concept and results of an experimental campaign at PHELIX, which verified the on-shot intensity gain of this AO loop, in chapter 6.
Chapter 7 contains the summary of this work and some future perspectives.

Additionally, I recommend reading the appendix of this work: in appendix A, I review the AO technology that is commonly encountered in HI lasers or is, at least, of interest for certain applications. In appendix B, I summarize the diffraction theory that is relevant to the formation of the focal spot in HI lasers, including polychromatic and vectorial diffraction integrals, as well as their numerical evaluation. This chapter provides the knowledge that is necessary to realistically simulate the time-dependent intensity distribution on the target plane, and simultaneously shows why the exhaustive collection of the data necessary to do so is a challenging problem.

I published the results of my work in two separate papers: the first one was, a theoretical paper on an alignment scheme for OAP telescopes [1], as implemented in my setup. The second one included a summary of the technical implementation of the AO loop [2] and the results of the experimental verification of the intensity gain.

2. Adaptive Optics

In this chapter, I will cover the basics of Adaptive Optics (AO) in terms of history, viewed from the perspective of open science, a selection of popular applications, basic concepts and mathematical approaches.

2.1. History

2.1.1. The Challenges of Ground Based Astronomy

The origins of AO lie in the field of ground based astronomy. As a larger telescope aperture increases the angular resolution on the sky and the amount of collected light at the same time, scientists built increasingly large telescopes to view objects that were fainter, smaller and further away. This increase in size brought new problems that had not been faced before.
On one hand, there was the issue of gravity, which made large optics and constructs sag under their own weight, depending on the current orientation of the telescope. Changes in temperature suddenly got notable as well as the thermal expansion of the material accumulated with increasing size. Both these influences deteriorated the optical quality of the telescope itself.

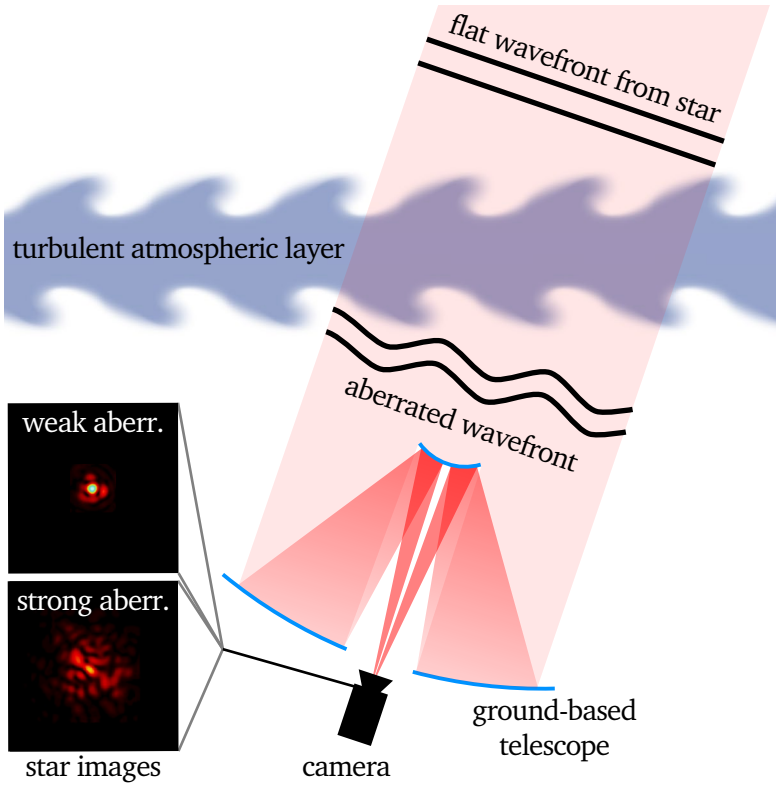

Figure 2.1.
Illustration of the problem of atmospheric turbulence in ground-based astronomy. Two simulated examples of a star image under weak and strong aberrations are shown on the left[1].

On the other hand, there was Earths atmosphere, through which the light of astronomical objects had to pass. Atmospheric layers of different temperatures and therefore densities generated rapidly evolving turbulence, which distorted the WaveFront (WF) of the incoming light, blurring out the images of the objects (see Fig. 2.1). This issue existed before, of course, but larger telescopes were impacted more than smaller ones as they captured more of the turbulence over their aperture, leading to higher spatial frequencies.

[1]The star images are public domain images from Wikimedia Commons, where I modified the color scale to be comparable with the rest of this work.

Both these effects made telescopes stay behind their expected performance significantly. The large costs that were associated with these devices made clear that scientists needed to come up with technological solutions to advance the possibilities of ground based observatories.
The two major innovations that tackled the two problems that I described above were active optics for the deformation of the telescope and, later on, adaptive optics for the atmospheric turbulence.

2.1.2. Active Optics

The first innovation was developed in the 1980s in the form of active optics [14]. This technology aims to continuously correct the shape of the mirrors in telescopes in order to compensate deformations introduced by gravity sag and temperature change.
Conceptually, the approach is quite similar to the one used in Adaptive Optics, which I will describe in section 2.2, where the latter is used to compensate the aberrations introduced in atmospheric turbulence [15]. While the borders between the two were blurred back then, the technological advances made their differentiation much easier over time. Generally, the update rate of active optics is smaller than one per second, which made building larger telescopes with thin, actively corrected primary mirrors possible even before the available computation speeds reached the regimes necessary for the compensation of atmospheric turbulence.

The definition of active optics in astronomy does, however, not transfer to High Intensity (HI) laser systems, as these do not suffer from dynamic gravitational sag. Therefore, I will refer to active optics in the context of HI lasers whenever optics are actively moved or deformed without a control loop instead. This is typically done to account for thermal or shot aberrations (see chapter 3). Using the terminology ”active optics” that way does therefore draw a suitable analog in terms of the user case of this technology.

2.1.3. Adaptive Optics

As mentioned before, AO in astronomy aims to compensate for rapidly evolving atmospheric aberrations, which is much more demanding in terms of update speed. The early history of AO is well documented by Becker [15].
The concept was first proposed by Babcock in 1953 [16]. His idea was to coat a mirror in the telescope with a thin oil film that had charges deposited on its surface. Electrodes on the rear side could then apply force on the oil surface, which would deform the film. The refraction of the light in the oil was supposed to compensate for aberrations introduced by the atmosphere. While the concept of this mirror was never realized, the proposal of Babcock remained significant as his description of how to implement a closed control loop was remarkably complete.
The first investigations on an actual implementation of AO were done in the well funded military context in a classified manner, which would only contribute to open science decades later when the progress was declassified. The first attempts in open science were done by Buffington, who published a paper in 1977, describing a one dimensional array of six mirror segments that were controlled by a stochastic feedback loop [17]. While being of little practical use, this demonstration showed the potential behind AO and inspired further developments: only a few years later the first implementation of a continuous Deformable Mirror (DM) with 21 actuators and a shearing interferometer as WaveFront Sensor (WFS) (see section A.1.3) was demonstrated for solar imaging [18]. While the results were promising, it showed the technical and financial hurdles that lied ahead in the path of AO.
It took another decade until significant increases in the number of actuators and control speed were made, which was driven by technological innovation in the manufacturing of DMs and slow grown in

computational power [19]. Additionally, around the same time in 1994, laser guide stars were first demonstrated, which greatly increased the use of AO as no natural, bright guiding stars in proximity to the observed object were needed anymore for the control loop [20].
This fueled the development of AO and, coupled with the exponential growth of computational capacities, triggered a cascade of ever advancing AO developments [21]. These included innovations in different ways: first, there were new conceptual schemes, like multiconjugate AO, which employed multiple DMs for different atmospheric altitudes [22] or multiobject AO, which measured the WF in different directions on the sky, enabling the correction for multiple objects at once. Second, the scaling of the hardware reached impressive dimensions, where AO nowadays can update DMs of thousands of actuators with rates of several kHz. Third, new control schemes became available, e.g. control in terms of beam modes, modal gains or predictive control (see section 2.2).
By now, AO is a mature field in astronomy and considered to be standard equipment for modern ground based telescopes.

2.1.4. Applications of Adaptive Optics

Over the course of this development, which was mainly driven by astronomy, AO found its way into a variety of other applications. While there are more complete and thorough reviews of this topic out there [23], I will list and briefly describe some of these here in order to underline the significance of this technological innovation.

Military applications, despite being a controversial topic, significantly contributed to the advances in the field of AO (e.g. concerning artificial guide stars). The development was fueled by the cold war, where laser defense systems against intercontinental nuclear missiles provided a promising alternative to other missile defense systems which had to be located in close proximity to the nuclear launching site [24]. Politically, placing a satellite equipped with a laser or a relay mirror in a low orbit was a much more sound way to install countermeasures than hiding submarines off the coastlines of the enemy in the face of a possible nuclear war. Here, AO was necessary to compensate for atmospheric turbulence and/or to counteract the WF degeneration of the chemical lasers which provided powers in the megawatt regime.

Laser processing is the discipline of altering the properties of matter via interaction with laser beams. The applications cover an enormous range from cutting over welding, material hardening by peening, surface cleaning and generation of nano-structures to three dimensional sub-surface processing. A lot of these applications are of high interest for the industry. Here, AO can be helpful in a wide range of scenarios, e.g. when the location of the focus is changed, the laser enters a transparent material or the system heats up. Furthermore, active shaping of the focal spot with AO can provide access to new applications or enhance existing ones, e.g. by generating several replica of the focal spot for faster processing [25].

Microscopy of the interior of living tissue is a discipline that relies on AO as well. Here, the objects of interest lie embedded within organic material. Contrary to telescopes which have to deal with aberrations introduced by the atmosphere, these microscopes receive light that has to pass through layers of living tissue with varying density, which can change over time. This deteriorates the spatial resolution of the microscope, which is why AO is used to counteract these aberrations. An example for this comes from the field of neurobiology, where the working of living neurons shall be observed down to synaptic resolution [26].

Ophthalmology is an application from biomedicine, which aims to provide a high-resolution, three dimensional image of the human retina. This diagnostic helps to identify a variety of diseases in the human eye and became an extremely valuable technology in medicine. The challenge here is that the light from the retina has to pass through the eye and exits the pupil before entering the diagnostic, which introduces aberrations on slow timescales (the kind that can be treated with glasses or contact lenses) and fast ones, which can be attributed to the dynamics of the living organism. These aberrations can be compensated using suitable adaptive optics. Solutions for this are commercially available nowadays [27].

Free-space communication is a method for information transfer that uses light propagating through space (ground-to-ground [28], space-to-ground [29] or even underwater [30]) rather than through optical fibers. This technology offers unprecedented data transfer rates compared to conventional wireless transmission technologies (e.g. radio waves) as the light frequency is much higher, while also outrunning fiber optics as the spatial properties of the beam can be used to encode additional information. One example for the latter is to use an overlap of different beam modes featuring Orbital Angular Momentum (OAM) where each mode represents an independent transfer channel. The receiver can then de-multiplex these by separating the beam modes [31]. This requires keeping a proper mode purity over long distances of propagation, which is the reason why AO is indispensable to compensate for turbulent aberrations.

Aside these examples, AO is of significant importance in the context of HI laser systems. This is the central concept of this work and I will go into detail on various aspects in the following chapters.

2.2. Wavefront Control

In this section, I will cover the basic conceptual and mathematical approaches to AO. The topics I cover here are not comprehensive and leave out a variety of concepts that do not apply to the field of HI laser systems, such as multiobject AO, but will support the understanding of requirements and hardware of the AO loop that I implemented in this work.

2.2.1. Active Components

As described in the history part of this chapter, one of the key components of any AO setup is the active component. This can be any piece of optic that is able to manipulate the phase of incident light in a spatially resolved manner, where the phase pattern can be actively changed. The most widespread approach in our context is to use a Deformable Mirror (DM) for this. Here, a mirror is equipped with a set of actuators on its rear side that can manipulate the shape of the mirror, thus actively influencing the shape of the reflected WF (see Fig. 2.2).
I give a list of available technology for AO in HI lasers in section A.2 of this work.

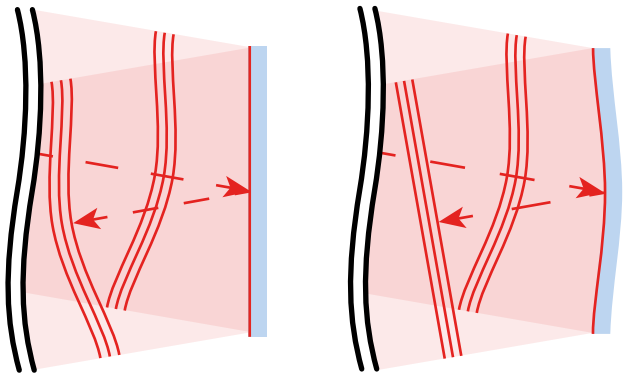

Figure 2.2.: Illustration of an aberrated WF being reflected on a flat mirror (left) and an adapted mirror (right).

2.2.2. Control Loop Setup

In order to achieve the desired WF, the active component (I will refer to these simply as DM from now on) needs to be controlled carefully. This naturally requires some sort of sensor that delivers feedback on the effect the DM has on the WF. Here, a broad range of different approaches is conceivable [21], but the most widely spread way is to use a WaveFront Sensor (WFS) that provides information about the WF on the surface of the DM.
Naturally, a Control Unit (CU) has to evaluate the data of the WFS, calculate the control values for the DM and send the latter to the hardware driver of the DM. This concept is shown schematically in Fig. 2.3. I discuss different technologies for WFSs and CUs in HI laser systems in the sections A.1 and A.3, respectively.

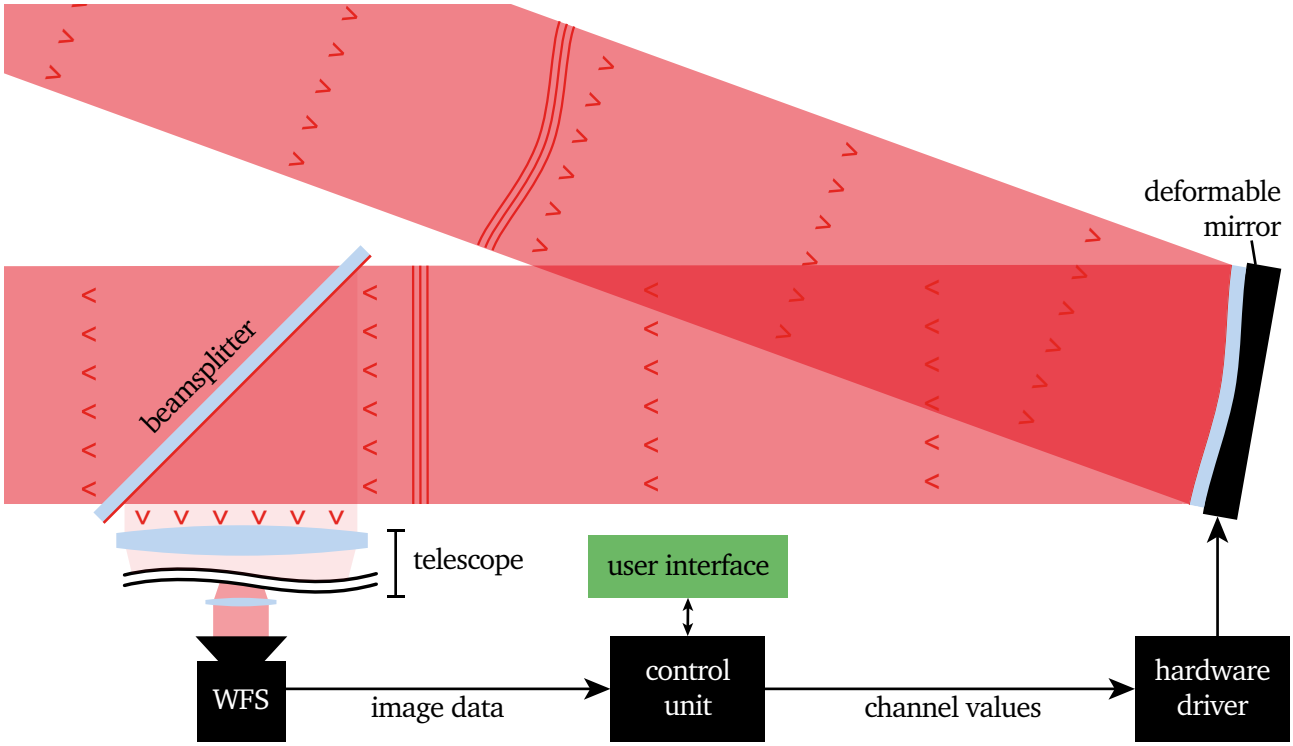

Figure 2.3.: Schematic of a simple AO setup for a laser beam with the main components. The laser beam enters from the top left.

An important concept for this setup to work is that the WFS measures the WF directly on the surface of the DM, as WF deformations translate into beam distortions over the propagation distance. A simple example would be a tilted WF which changes the direction of the beam propagation. Measuring the WF in another plane would mean that the beam as a whole shifts on the WFS, degrading the quality of the mirror response (see next section).
As the WFS naturally cannot sit on the surface of the DM itself, relay imaging has to be done: a telescope (shown in Fig. 2.3 on the bottom left) images the DM onto the WFS and thus reverses the diffraction that happens due to WF deformations.

2.2.3. Response and Control Matrix

In order to control the WF of the beam in a deterministic manner using a setup as illustrated in Fig. 2.3, the response matrix of the DM has to be known [32]. This matrix contains all response functions of the individual DM actuators. I will discuss the response recording and construction of this matrix in the following.
Given the WFS records the WF in a sampling manner, i.e. the phase of the beam is measured on an arbitrary distribution of discrete points indexed over i, then each actuator j of the DM has a unique response over these points. Specifically, one can measure the phase W_i at each WF sample while changing the control value U_j of each actuator in order to obtain the response $R_{i,j}$[2]:

$$R_{i,j} = \frac{W_i(U_{j,1}) - W_i(U_{j,2})}{U_{j,1} - U_{j,2}}. \tag{2.1}$$

Here, we meet the assumption that the impact of the control value U_j on the WF is linear, which is approximately the case for the deformations that are typically used in AO (no more than a few ten microns).
The values of $R_{i,j}$ can be written in the form of a matrix $\boldsymbol{R}$ (bold symbols represent matrices from here on), where the column at position j represent the WFs changes (written as a vector) for a voltage change of the corresponding actuator.
With this knowledge, any destination WF $\vec{\Phi}_{\text{dest}}$ can be approximated by applying the voltages $\vec{U}$ such that the expression

$$\vec{\Phi}_{\text{dest}} - \vec{\Phi}_{\text{aberr}} - \boldsymbol{R} \cdot \vec{U} \tag{2.2}$$

gets minimized. Here, $\vec{\Phi}_{\text{aberr}}$ corresponds to the aberrations that are already present in the system.
While the expression (2.2) can in principle be solved directly, this is a computationally demanding task. A much faster way of calculating the control values is to rephrase the problem such that

$$\boldsymbol{C} \cdot \left(\vec{\Phi}_{\text{dest}} - \vec{\Phi}_{\text{aberr}}\right) - \vec{U} = 0 \tag{2.3}$$
$$\text{with} \qquad \boldsymbol{C} = \boldsymbol{R}^{-1}.$$

The inverse $\boldsymbol{C}$ of the response matrix $\boldsymbol{R}$ is called the control matrix. The advantage of this is that the calculation of the control values $\vec{U}$ is now a matter of a single matrix-vector multiplication. At the same time, the capabilities of the DM are embedded in $\boldsymbol{C}$ such that only aberrations that can geometrically be compensated by the hardware have an influence on the control values, making an optimization (such as with the term (2.2)) unnecessary.
In reality, however, $\boldsymbol{R}$ is not necessarily quadratic and singular, which means that it cannot be inverted to form $\boldsymbol{C}$ in general. A feasible workaround is to approximate $\boldsymbol{C}$ with the pseudo-inverse of $\boldsymbol{R}$ using a Singular Value Decomposition (SVD) [34]. This operation takes any real matrix $\boldsymbol{M}$ of rank r and decomposes it into two unitary matrices $\boldsymbol{U}$ and $\boldsymbol{V}$ and a matrix S that contains only zeros except for the first r diagonal elements:

$$\boldsymbol{M} = \boldsymbol{U}\boldsymbol{S}\boldsymbol{V}^T \tag{2.4}$$

[2] While performing pokes on individual actuators is a straightforward way to record a response matrix, this may not be the best way to do so - especially not for DM types with highly localized influence functions for each actuator. An alternative is to use Hadamard-encoding for each poke to move half of the actuators each time, and decode the responses afterwards. That way, the signal is distributed over the whole aperture, improving the signal-to-noise ratio dramatically [33].

For a matrix $\boldsymbol{M}$ of the size MxN, $\boldsymbol{U}$ measures MxM, $\boldsymbol{S}$ MxN and $\boldsymbol{V}$ NxN elements. This way, the pseudoinverse of $\boldsymbol{R}$ can be calculated using equation 2.4:

$$\boldsymbol{C} = \boldsymbol{R}^{-1} \approx \boldsymbol{V}\boldsymbol{S}^{-1}\boldsymbol{U}^{T}, \tag{2.5}$$

where the right side of the equation is formed from the SVD of $\boldsymbol{R}$.
This decomposition offers even more insight into the characteristics of the control loop: the column vectors of $\boldsymbol{U}$ contains the eigenmodes of the DM, which scale with the singular values on the diagonal of $\boldsymbol{S}$. Bigger singular values indicate a larger deformation of the WF corresponding to this eigenmode for the same amount of actuator displacement.

2.2.4. Control Schemes

With the physical AO setup in place and the response matrix measured, there are a variety of options to control the WF of the system. In this section, I will briefly describe the three main families of control loops: open, closed and pseudo-open loops.

2.2.4.1. Open Loop Control

The most straightforward way to control the WF is to use equation (2.3) to directly forward a set of control values to the DM. Here, the vectors for the desired WF can simply be replaced by the desired WF shape and the DM will mimic this shape as close as mechanically possible with in the limits of its spatial resolution and dynamic range.
The disadvantage of this operation scheme is that one has to rely on the linearity of the system which is the underlying assumption behind using a control matrix. Furthermore, the aberrations that are already present in the system or even evolve over time are not automatically incorporated in the control commands.
Nonetheless this scheme can be useful e.g. when the stability of the loop is not given but one wants to check the influence of a certain type of aberration in the Zernike sense (see section 3) onto other parts of the beam. After the response matrix has been recorded, no WFS is needed anymore.

2.2.4.2. Closed Loop Control

The more accurate control scheme is the closed loop control, which is depicted in Fig. 2.3. Here, WFS continuously measures the WF and the CU compares the output to a desired destination. The control works different to equation (2.3), as only incremental steps are considered in this scheme. This leads to an iterative approach:

$$\vec{U}_{(n+1)} = \vec{U}_{(n)} + \boldsymbol{C} \cdot \left(\vec{\Phi}_{\text{dest}} - \vec{\Phi}_{\text{meas}}\right), \tag{2.6}$$

where $\vec{U}_n$ is the set of channel values that the DM receives in the nth step of the loop, $\vec{\Phi}_{\text{dest}}$ is the desired destination WF and $\vec{\Phi}_{\text{meas}}$ the WF that is currently measured by the WFS.
This scheme allows for some modifications to increase the stability of the loop. For instance, one can multiply the control matrix with a feedback-parameter $0 < g \leq 1$ to slow down the convergence. This is necessary for control loops that run faster than the time it takes from the moment a DM command is issued to the moment its influence is fully measured by the WFS.
Another way to increase the stability is to modify the control matrix $\boldsymbol{D}$ in terms of its modes. As I described in the last section, $\boldsymbol{C}$ is formed using the SVD of the response matrix $\vec{R}$. Here, the matrix $\boldsymbol{S}$

holds the singular values, which scale the modes of the DM. By artificially changing these values, the behavior of the loop adapts accordingly. A common practice is to set a number of the smallest values in $\boldsymbol{S}$, which correspond to the noisiest eigenmode measurements of the system, to zero, making the loop immune to overestimating noise. In a similar manner, the gain of individual modes can be changed to make the loop more stable, or set to zero to exclude them from the loop [21].

2.2.4.3. Pseudo-Open Loop Control

The last family of control schemes are the pseudo-open loops. This hybrid concept does not work incrementally as the closed loop, which I described in the last section, but on the WF prior to the DM in order to decouple the WF measurement from the control loop. Naturally, this does not work by simply placing a WFS prior to the DM as there is no response matrix available here. Instead, the control values $\vec{U}$ of the DM are multiplied with the response matrix $\boldsymbol{R}$ to calculate the current impact of the DM onto the WF, which is then added to the WF measurement. The DM is updated according to equation (2.3), hence the name ”pseudo-open”.
Having access to the raw, uncorrected WF offers new possibilities. Aside from enabling multiconjugate AO in astronomy where multiple WFSs and DMs act on one beam at the same time [35], any model requiring information on the original WF can be inserted into the loop. As an example, predictive models can be trained to estimate the evolution of the WF over the latency of the AO loop [33]. This can greatly increase the accuracy, mitigating the effects of loop latency.

Conclusion

In this introductory chapter, we have seen the history of AO with its rapid developments. These lead to a broad range of applications that benefit from AO nowadays. Furthermore, we learned the basic structure of simple AO setups and how they are operated.

Having this overview, I will now narrow down the problem to HI laser systems step by step. We will learn about the specific challenges that this field poses to AO, which is necessary to understand the concepts and design decisions of the practical part of my work.
To do so, I will first describe the spatial and temporal nature of aberrations that typically occur in HI laser systems in the next chapter. After that, I will discuss the challenges and requirements that arise from these circumstances in chapter 4. For the interested reader, I recommend reading the appendix A, in which I review the AO technology for HI lasers that is available nowadays, which greatly complements the knowledge needed for the main experimental part of this thesis: the implementation of a post-compressor loop at the Petawatt High-Energy Laser for heavy Ion eXperiments (PHELIX) laser, which I will describe in great detail in chapter 5.

3. Composition of Aberrations in High-Intensity Laser Systems

To understand the demands on Adaptive Optics (AO) systems in High Intensity (HI) lasers, one has to understand the nature of WaveFront (WF) errors that have to be compensated first. These errors are composed of static, dynamic and chromatic aberrations and are, contrary to other applications, not dominated by external sources but determined by the architecture of the laser system itself.
Therefore, I will start this chapter with a brief review-section on the architecture of modern HI laser systems and I will also describe the architecture of the Petawatt High-Energy Laser for heavy Ion eXperiments (PHELIX) laser, where this work took place. After that, I discuss different ways to describe and characterize aberrations. Finally, I will discuss the sources of static, dynamic and chromatic aberrations in HI laser systems in the sections 3.3, 3.4 and 3.5, respectively.

3.1. The Architecture of Modern High-Intensity Laser Systems

As the characteristics and amounts of aberrations in HI laser systems are determined by the laser architecture, it is necessary to be aware of the core concepts of modern HI laser setups. As each of these have been discussed in literature thoroughly, I will only provide a brief review here.

3.1.1. Laser Oscillator

The first laser was built by Maiman in 1960 [3]. His setup featured the core components that are necessary for any laser oscillator: an optical cavity or resonator, which allows light to oscillate in stable longitudinal and transversal modes, an active medium, also called gain medium, which is able to amplify certain WaveLengths (WLs) of light, and an energy pump which is able to generate a population inversion of the states inside the gain medium and thus enables the amplification process in the first place.
The traditional scheme of such an oscillator is shown in Fig. 3.1, where the resonator is depicted as two curved mirrors, the active medium as a volume in between and the energy pump as an abstract source of energy. Some of the light is transmitted through the output coupler of the resonator.
However, each of these components can be realized in very different ways. The resonator for example can also be a fiber with a regular pattern of optical density modulations, called a fiber Bragg grating, that acts as reflector. Examples for active media are gases, liquids, doped glasses or crystals and semiconductors. Examples for energy pumps are flash lamps, electrical discharges or even other lasers.

In the context of HI lasers, only resonators featuring stable modes are of interest. Here, their properties can be divided into two categories: spatial and temporal.
The spatial property that is of relevance for the beam quality is the transversal mode selectivity of the resonator, meaning that only certain beam modes can run stably in the oscillator. This generally means that the shape of the beam that leaves the oscillator is very pure. Changes in the resonator primarily lead to increased losses and therefore a reduced output energy, while at the worst changing the mode itself. High-order aberrations are usually negligible in this context.

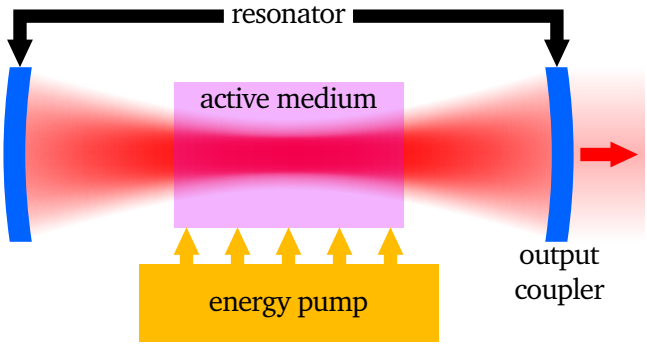

Figure 3.1.: Schematic of a laser oscillator. The running transversal oscillator mode is illustrated in red.

The temporal behavior of a laser oscillator is dominated by the amplification bandwidth of the gain medium and the longitudinal modes of the cavity (only standing waves). If no special measures are taken, the longitudinal modes that are amplified beyond the cavity losses are running in an uncorrelated manner, leading to a constant output of light from the cavity (see Fig. 3.2, left). This is of little use for HI lasers as one is interested to concentrate the energy into a time-interval as short as possible. A way to achieve this is to force the longitudinal modes to feature a linear phase correlation, which is called modelocking. Due to interference, this generates a single short pulse which is traveling through the resonator (see Fig. 3.2, right). This was first done by introducing a periodical modulation of the cavity losses, aligned with the round trip time of the resonator, using an Acousto-Optic Modulator (AOM) [36]. Aside from these active modelocking approaches, passive ones are also possible, e.g. by introducing a saturable absorber, which absorbs uncorrelated mode combinations more than intense, correlated ones, or relying on nonlinear optical effects (i.e. Kerr lensing) that only allow intense pulses to run in the resonator [37].
The light that is coupled out of a modelocked resonator is therefore a pulse train with a repetition rate well above megahertz. In HI lasers, a pulse picker has to isolate single pulses from the train (e.g. by using an AOM to deflect certain pulses) in order to amplify them to high energies.

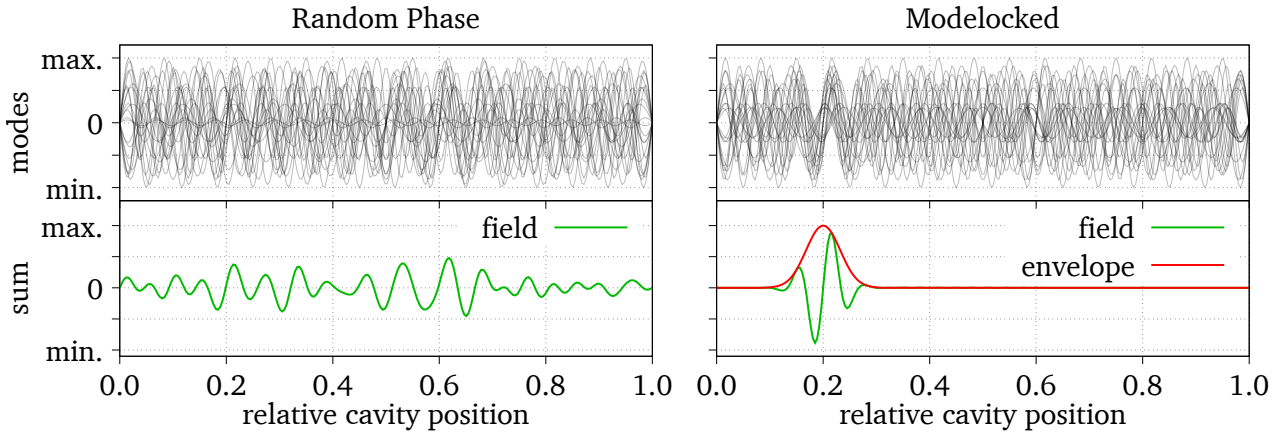

Figure 3.2.: Visualization of the effect of modelocking on the total field inside a cavity. The normalized mode fields are shown on the top, while the sum of all modes is shown on the bottom. The modes on the left side feature a random phase, while the modes on the right side are in a linear phase relation.

3.1.2. Master Oscillator Power Amplifier

A laser oscillator is strongly limited in energy and intensity due to damage thresholds and nonlinear optical effects. Aside from the fact that mode-locked cavities feature extremely short pulses (see next section), this is mainly due to the relatively small spatial mode inside the resonator, which is usually no larger than a few millimeters. Furthermore, only a small fraction of the energy is transmitted through the output coupler for every round-trip. In order to reach larger pulse energies after the oscillator, HI lasers employ the Master Oscillator Power Amplifier (MOPA) concept [38].
The idea of this is straight forward: the beam diameter is successively increased to stay below damage thresholds and keep nonlinear effects small. After each magnification, the beam passes through an active medium to gain energy.
The amplifiers can again take different shapes: fiber systems, as an example, may simply use active fibers for amplification. Freely propagating beams with small diameters on the other hand allow for the use of mode-selecting cavities, featuring a large number of passes through the medium while retaining a high beam quality. This is called a regenerative amplifier. Larger beam diameters may pass through rod-shaped media, while the largest amplifiers with apertures of up to 40 cm in diameter consist of doped glass slabs that are placed in the beam at the Brewster's angle in order to minimize reflection losses. Generally, larger beam diameters are more prone to accumulate aberrations than smaller ones. A schematic example for a MOPA system is shown in Fig. 3.3.

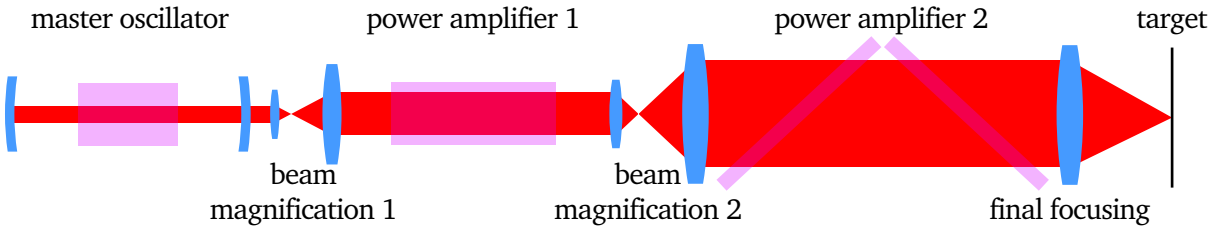

Figure 3.3.: Schematic of the MOPA concept, illustrated with a solid state oscillator, a rod- and a slab amplifier.

3.1.3. Chirped Pulse Amplification

The MOPA concept does allow for high energies, but is still only usable for long laser pulses. In 1985, Strickland and Mourou published a paper on a concept which would revolutionize HI lasers, enabling unprecedented peak powers by using pulses from mode-locked cavities [11]. This concept became popular as Chirped Pulse Amplification (CPA) and was awarded with the Nobel Prize in 2018.
CPA makes use of the linear phase correlation of the longitudinal modes in modelocked oscillators (see section 3.1.1). The idea is to add a quadratic spectral phase term to the pulse using a setup with a large dispersion, which effectively means that the pulse is spectrally smeared out in time - a process called pulse stretching. This generates a pulse with a much larger duration and therefore lower peak intensity, while maintaining a fixed correlation between the frequencies. This low-intensity pulse can then be amplified safely, until the quadratic spectral phase is removed using another setup which reverses the dispersion, generating the final, ultra-intense laser pulse. This is called pulse compression.
Today, all ultra HI laser systems worldwide rely on this scheme [4] and recently, intensities of more than $10^{23}\,\mathrm{W/cm^2}$ have been realized for the first time [13].
These systems usually rely on setups that are based on reflective optical gratings in order to achieve the dispersion that is necessary for damage-safe pulse amplification. The working of this is easy to understand when looking at a sketch of a common pulse compression setup, as shown in Fig. 3.4.

Here, two parallel pairs of gratings disperse the beam spatially first (gratings G1 and G2) and then re-overlap the frequencies again (gratings G3 and G4). Obviously, longer WLs feature a longer optical path than the shorter ones, which indicates the dispersion of the system. A larger distance between the parallel gratings increases the dispersion.
As the final grating witnesses the fully compressed pulse and therefore the peak laser power, these gratings have to be large, reaching meter-scale diameters nowadays.
Mathematically, one can obtain a dispersion with the opposite sign by building a compressor with a negative distance between the gratings. As this is obviously not possible straight away, a workaround has to be employed: using an imaging system, the second grating can be placed relative to the image of the first one, where a negative distance to this image can be achieved easily (see Fig. 3.5). The spatial overlap between the frequencies can then be restored by using a double-pass configuration, which saves space and components. This is not possible for pulse compressors due to the large beam diameter.

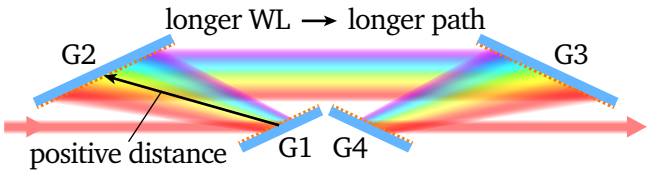

Figure 3.4.: Schematic of a grating-based pulse compressor with four gratings. The spatially dispersed beam is shown using a rainbow gradient, where the red side of the spectrum corresponds to longer WLs.

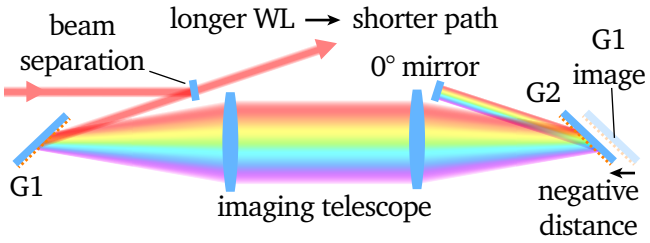

Figure 3.5.: Schematic of a grating-based pulse stretcher with refractive lenses for imaging. Two gratings are used in a double-pass configuration.

In theory, the setups of the pulse compressor and stretcher, as I discussed them here, could be exchanged. However, the stretcher setup requires more optical components as an imaging system is required, which makes the compressor setup clearly favorable for large beam diameters.

3.1.4. PHELIX Architecture

After I explained the core concepts of HI lasers, I will now spend a section to narrow down on the laser system where this work took place: the Petawatt High-Energy Laser for heavy Ion eXperiments (PHELIX) at the GSI Helmholtzzentrum für Schwerionenforschung GmbH (GSI) in Darmstadt, Germany [39]. Having a rough overview of the architecture of this system will greatly help to understand what can be achieved with a post-compressor AO loop and what decisions I made to achieve this.
In the following, I will go through the laser chain of PHELIX, briefly touching on the most important components, while placing emphasis on active and adaptive optics. The block diagram of PHELIX is shown in Fig. 3.6, which coarsely resembles the layout of the laser bay.

3.1.4.1. Frontends

At PHELIX, the term FrontEnd (FE) describes the functional blocks that provide laser pulses with sufficient energy to be passed to the PreAmplifier (PA). There are two FEs: the fs FE, providing pulses of less than 500 fs duration, readily stretched for amplification, and the ns FE, which delivers long pulses with a duration between 1 ns and 10 ns and an arbitrary envelope curve.

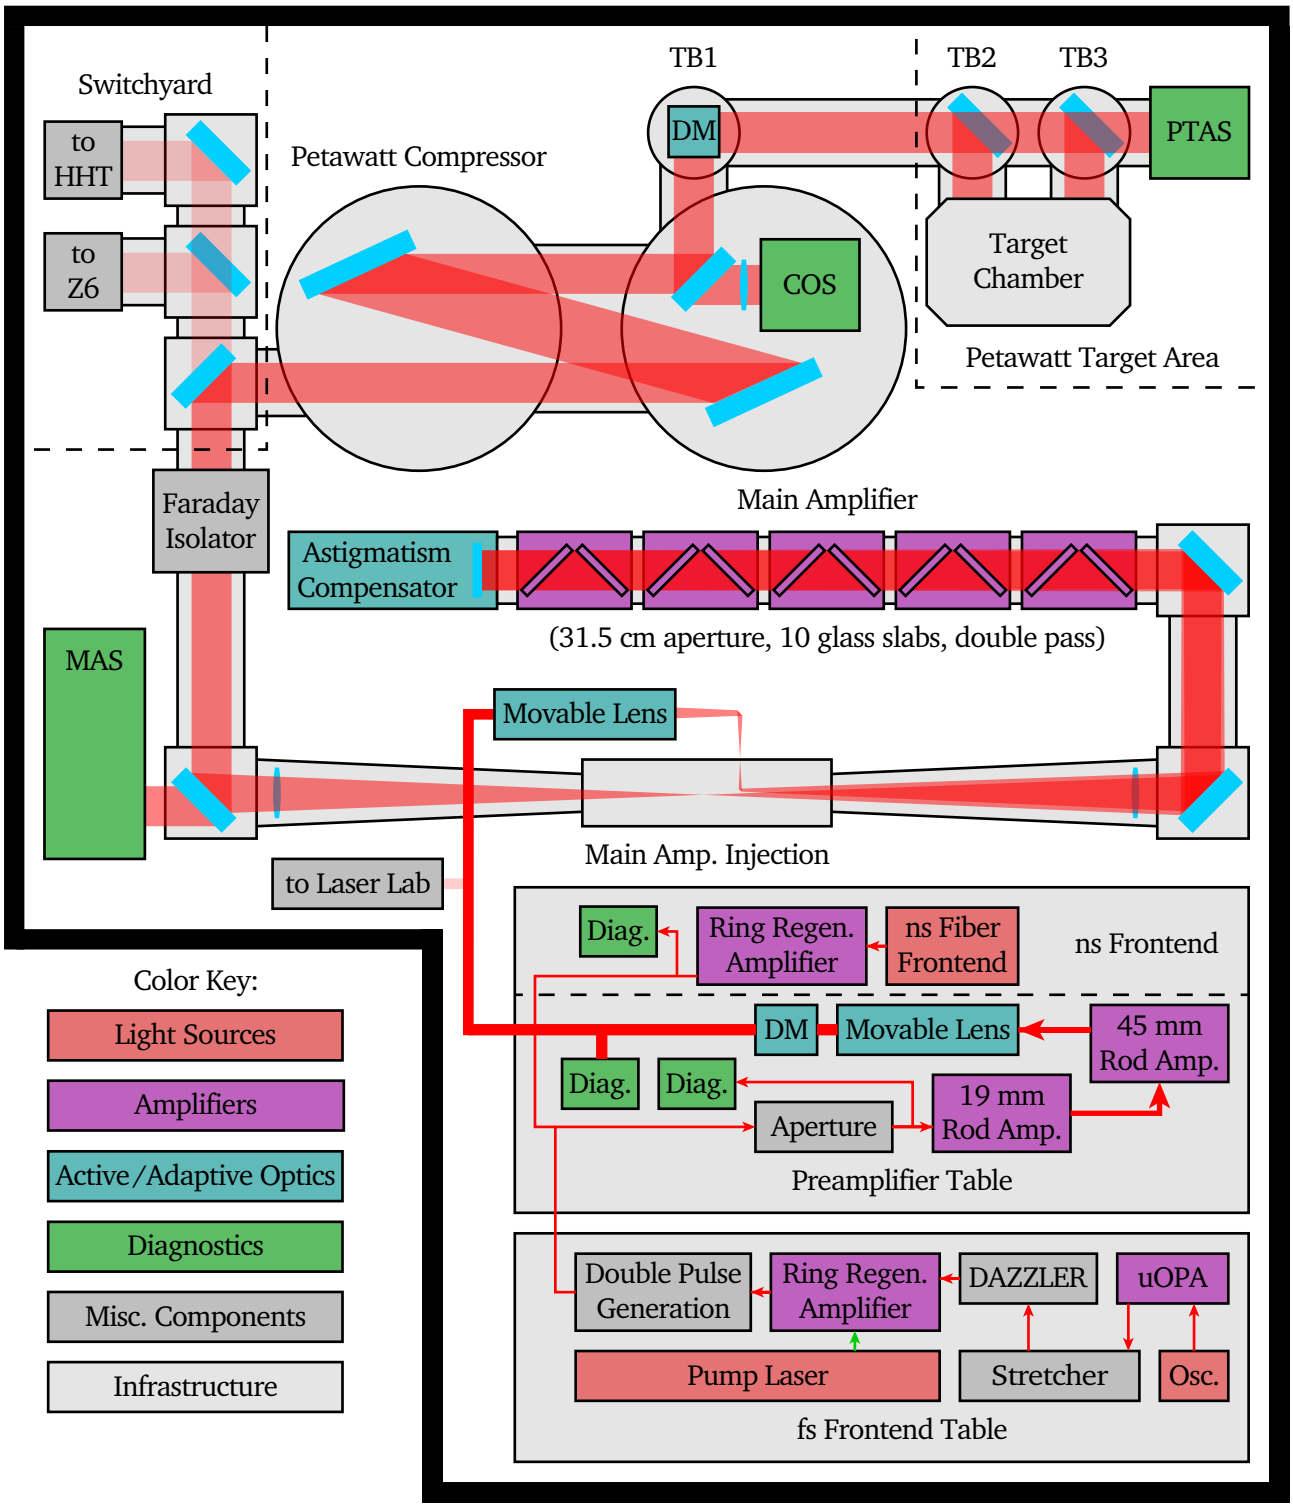

Figure 3.6.: A block diagram of the PHELIX laser system, where the blocks are arranged in a way that coarsely reflects the layout of the laser bay (not true to scale).

fs FE: The fs FE sits on a dedicated optical table (at the bottom of Fig. 3.6). The master oscillator is a Mira 900f Titanium-doped Sapphire (Ti:Sapphire) laser by COHERENT, INC. (USA) with a Synchrolock-AP that synchronizes the mode-locked pulses with the timing signal. The crystal is pumped by an external Continuous Wave (CW), frequency-doubled Neodymium-doped Yttrium Aluminum Garnet (Nd:YAG) laser. The cavity features a repetition rate of 72 MHz and delivers a pulse energy of 3 nJ at a duration of 130 fs Full Width Half Maximum (FWHM) due to the large amplification bandwidth of the Ti:Sapphire crystal. The cavity is tuned to a central WL of 1053.4 nm in order to match the Neodymium-doped Glass (Nd:Glass) amplifiers of PHELIX.
The first amplification stage is an integral part of the fs FE: a two-step ultrafast Optical Parametric Amplifier (uOPA) uses difference frequency generation in a nonlinear crystal to transfer energy from a pump pulse to the seed pulse to achieve a high gain while providing an excellent temporal contrast. This is the successor to the one-stage uOPA that was installed previously [40].
The pump pulse for the uOPA is generated from a preceding oscillator pulse: the pulse is coupled into an optical fiber, broadened via self-phase-modulation, stretched and then amplified in active, ytterbium-doped fibers to about 1 µJ. After that, the pulse is coupled into a dedicated regenerative Ytterbium-doped Yttrium Aluminum Garnet (Yb:YAG) amplifier with a 10 Hz repetition rate, followed by a pulse compressor. After frequency doubling, the pulse length is 1.4 ps at an energy of 11 mJ and a central WL of 515 nm. This pulse is then split to pump the two uOPA stages.
Here, the seed pulse is first amplified to 5 µJ in the first stage and to 270 µJ in the second stage, which corresponds to a gain of almost five orders of magnitude. This is sufficient to spare one regenerative amplifier, which greatly benefits the temporal characteristics of the PHELIX pulse [41].
After the uOPA, the pulse enters the stretcher of PHELIX and is stretched to 2.4 ns. The lowered intensity now allows for passing the pulse through an acousto-optic programmable dispersive filter (DAZZLER by FASTLITE, France), which modifies the spectral phase of the pulse to achieve the shortest pulse duration in the target chamber [42].
The last amplification stage of the fs FE is a ring-shaped regenerative Ti:Sapphire amplifier, pumped by a pulsed Nd:YAG laser (Powerlite 8000 by Continuum, nowadays a part of Amplitude, France). The output energy is 17 mJ with a beam waist of 2 mm. Note that as the last amplification stage of the fs FE is a regenerative amplifier, the beam quality will be generally good due to the mode selectivity of the system. Here, aberrations rather manifest in a loss of energy.
The last component of the fs FE is an optional double pulse generation stage, which can divide the pulse into two parts with an adjustable energy ratio via a polarizing beamsplitter and a waveplate. A delay stage sets the desired temporal relation between the two.

ns FE: The second FE of PHELIX is the ns FE. Due to the larger pulse length, the temporal concerns are much lower here, which makes the system simpler than the fs FE:
The light source is a fiber based pulse generator (ModBox FE-1053nm-AWG by iXblue, France). This device features an arbitrary waveform generator, which can deliver arbitrarily shaped pulses with a temporal resolution of 0.125 ns.
The pulse is then passed to a ring-shaped regenerative amplifier on the PA table (above fs FE table in Fig. 3.6) with a cavity length of 4 m, where the amplification is provided by a flashlamp pumped Nd:Glass amplifier at a repetition rate of 0.5 Hz.
The pulse length can be adjusted from 1 ns (lower limit due to damage thresholds) to 10 ns (upper limit due to the cavity length), where the output energy is 20 mJ on average, depending on the pulse length. The output waist is 0.7 mm.

3.1.4.2. Preamplifier

The second large block in the laser chain is the PA. Here, one of the properly magnified beams of the FEs is selected and passes through a serrated aperture, which, in conjunction with a following spatial filter, cuts out the desired part (up to 7.5 mm diameter, different shapes are possible) of the Gaussian beam. This is an important step as the final beam shall feature a high order super-Gaussian intensity distribution.
The overlap between the input beam and the aperture can be supervised by a ”Camera Box”. These measurement devices are standard at PHELIX and contain three cameras: one images the Near Field (NF) at the box entrance, one the WF at the same position and the last one observes the Far Field (FF). This means that the spatial beam quality can be fully assessed within one device.
After passing through the aperture, the beam gets magnified successively. At the end of the first telescope, it passes twice through a flash-lamp pumped Nd:Glass rod with 19 mm diameter. After the next magnification, a similar amplifier with 45 mm brings the beam energy up to a maximum of 6 J for pulses from the fs FE and 20 J for the longest pulses from the ns FE.
As the rods are pumped from all sides, they introduce some defocusing on-shot (see section 3.4), which can be compensated by another telescope with a lens on a linear stage after the amplifier. The beam now has a diameter of up to 6 cm, depending on the selected aperture.
Next, the beam gets reflected from a piezoelectric Deformable Mirror (DM) (see section A.2) with 31 actuators. Another Camera Box supervises the beam at the PA exit and can, if desired, form a closed loop with the DM.

At this point, the beam can either travel directly to the ”laser laboratory”, which is equipped with a pulse compressor and a target chamber, or pass on to the main amplifier for experiments that require higher pulse energies.

3.1.4.3. Main Amplifier

As the energy is in the joule range now, it can efficiently extract energy from the main amplifier. This block of PHELIX mainly consists of components from the dismantled NOVA laser facility [43], which was in operation until 1999 in the Lawrence Livermore National Laboratory in California, USA.
The beam gets injected into the main amplifier using another movable lens which precompensates the on-shot defocusing of the main amplifier. In the large telescope, the beam gets magnified to a diameter of up to 28 cm, depending on the aperture in the PA.
The amplifier itself features five flash-lamp pumped modules with two Nd:Glass slabs each and has a free aperture of 31.5 cm. The beam passes the slabs at the Brewster’s angle in order to minimize losses. The 0° mirror at the end of the amplifier chain is equipped with a set of stepper motors that can apply force to the substrate, making the mirror bend in order to compensate for the astigmatism that the main amplifier introduces on-shot. The beam travels back through the amplifier and passes the large telescope for relay imaging.
At the end of the telescope, the Main Amplifier Sensor (MAS) supervises the spatial beam quality, the spectrum and the energy of the amplified pulse. A WaveFront Sensor (WFS) here can control the DM in the PA to compensate for aberrations that are present in the system up to this point.
Finally, the beam passes through a large, pulsed Faraday isolator, which prevents light from being reflected back into the amplifier from the experiment.

3.1.4.4. Switchyard

At this point, the beam can be sent to different experiment areas:
The first option is the Z6 experiment cave, where the pulse can be compressed (fs FE) or frequency doubled (ns FE). Here, the beam can be combined with the ion beam from the Universal Linear Accelerator (UNILAC) of GSI.
The next option is to send the amplified pulses of the ns FE to the HHT cave via the novel 70 m beamline [44], where it is frequency doubled. Here, the beam can be combined with the ion beam of the heavy ion synchrotron *SIS100*. One major use of this is to generate X-rays with the laser to probe matter which is heated by the ion beam.
The last option remains inside the laser bay: the Petawatt Target Area (PTA) is the point of the facility where the highest intensities can be generated. As this is where this work happened, I will exclusively focus on this option in the following.

3.1.4.5. Petawatt Compressor

The petawatt compressor is shown on top center of Fig. 3.6. Here, the beam enters vacuum and remains to travel under this condition in order to prevent the pulse from ionizing gas, which can happen even in the NF after compression. Inside the vacuum vessel, two meter-scale optical gratings form a Treacy compressor [45]. This is similar to the compressor scheme shown in Fig. 3.4, but only G1 and G2 are present. Of course, this means that the beam is spatially separated in the NF after compression.
Naturally, this is not ideal for WF measurements, where only the spectrally weighted average can be measured, and therefore not for AO either. In the FF, however, the spatial dispersion will lead to a moderate pulse front tilt, which is tolerable in most cases.
Right after compression, the leakage light of the first turning mirror is used to supervise the beam compression in the COmpressor Sensor (COS). This beam sensor contains an autocorrelator to judge if the beam is well compressed or not, but also records the NF and FF of the beam.

3.1.4.6. Petawatt Target Area

The last subsystem of PHELIX is the Petawatt Target Area (PTA). Three turning boxes, named ”TB1” to ”TB3” (see Fig. 3.6, top), transport the beam towards the target chamber.
The large aperture DM which I used in the post-compressor AO loop was placed in TB1, while a leaky mirror in TB2 or TB3 reflected the light into the target chamber. The leakage light was used by the Petawatt Target Area Sensor (PTAS) to close the control loop with the DM.
Inside the target chamber, the beam gets focused with either a 45° Off-Axis Parabolic mirror (OAP) with a focal length of 400 mm or a 90° OAP with a focal length of 1.5 m, depending on the needs of the experiment.

3.1.5. Conclusion

In the last sections, I have summarized the main concepts of pulsed high-energy lasers in general and CPA systems in particular and provided an overview of PHELIX. The knowledge of the former is necessary when I discuss aberration sources in sections 3.3 to 3.5. The latter is crucial to understand the design decisions that I made when setting up the post-compressor AO loop at PHELIX, which I discuss in chapter 5.
In the next section, I will go over the most basic ways to describe aberrations, which I will pick up at several occasions in this thesis.

3.2. Description of Aberrations

When working with optical systems, a quantitative description of the WF is often required. This can be done in a variety of ways and the application determines the implementation. In this chapter, I will summarize the most common ways to quantify WFs and thus aberrations. First, I will discuss the modal decomposition. After that, I will describe some ways to statistically assess the quality of the beam, both in the spatial and the temporal domain.

3.2.1. Modal Decomposition

A common way to deal with WFs numerically is to decompose them into a basis which suits the application. In principle, such a basis can be any set of two-dimensional functions over the aperture where all members of this set are linearly independent. An example for such a set over a circular aperture are the well known Seidel aberrations [46, p. 236], which are constructed by a power series in polar coordinates and match the primary aberrations that often occur for misaligned systems.
Numerically, however, it is often useful to express the WF in an orthogonal basis over the aperture A, meaning that no single mode M_i contains contributions from any other mode M_j in the set. This can be phrased as

$$\iint_A M_i \cdot M_j \ d^2A = 0 \ \forall \ i, j \wedge i \neq j.$$

As an example, any eigenmode decomposition, e.g. from a DM (see section 2.2.3), automatically fulfills this condition.
The decomposition of a WF W into the basis-coefficients a_j then works as

$$a_j = \frac{\iint_A W \cdot M_j \ d^2A}{\iint_A M_j \cdot M_j \ d^2A}.$$

If the denominator equals unity for all indices, the basis is called orthonormal, which further simplifies the calculation of a_j.
The reconstruction of W then is

$$W = \sum_j a_j \cdot M_j.$$

Mathematically, orthogonal bases can be found for any aperture shape [47], which is often useful for WF reconstruction from its gradients. However, there are bases which are particularly useful for certain beam geometries.
The most popular of them are perhaps the Zernike polynomials [48, section VI], which form an orthogonal basis over the unit circle, and can therefore be applied for circular beams[1]. Similarly to the Seidel aberrations, the Zernike polynomials describe the primary aberrations in a system very well, meaning that only a handful of the lowest polynomials is often sufficient to describe the largest part of the aberrations of a misaligned optical system.
In polar coordinates, the Zernike polynomials can be expressed by the product of a polynomial $R_n^m(\rho)$

[1] Another popular basis for rectangular apertures are the Legendre polynomials, which are orthogonal and separable along the two axes of the aperture [49]. However I will only describe the Zernike polynomials here due to the circular beam of PHELIX.

in radial and a polynomial $A_m(\phi)$ in angular direction [50], where the non-negative integers m and n identify the polynomial:

$$Z_n^m(\rho, \phi) = R_n^m(\rho) \cdot A_m(\phi). \tag{3.1}$$

Here, ρ is the radial coordinate and ϕ the angle. For the indices, m needs to fulfill $m = n - 2k$ with the non-negative integer k, and $n \geq |m|$. The radial degree of the polynomial is n, while the angular degree is m.
The angular polynomial is defined as

$$A_m(\phi) = \begin{cases} m \geq 0 & \cos(m\phi) \\ m < 0 & \sin(m\phi) \end{cases}, \tag{3.2}$$

while the radial term is

$$R_n^m(\rho) = \sum_{k=0}^{\frac{n-|m|}{2}} (-1)^k \binom{n-k}{k} \binom{n-2k}{\frac{n-|m|}{2}-k} \rho^{n-2k}. \tag{3.3}$$

From equation 3.3 it can be shown that $R_n^m(1) = 1$, where the coefficients take integer values [50]. Furthermore, $|R_n^m(\rho)| \leq 1 \ \forall \ \rho \in [0, 1]$. This means that all polynomials are contained in $[-1, 1]$, as long as no other normalization scheme is applied. In this thesis, I will stick to the $|R_n^m(1)| = 1$ normalization. The first 15 Zernike polynomials are shown in Fig. 3.7, together with the expanded polynomial, the name and the associated Noll-index [51]. The latter is a convention that unifies the indices m and n into a single index j, which is a useful way to iterate through the polynomials.
The Noll-index is formed according to the rules

$$\begin{aligned} &j_\mathrm{u} < j_\mathrm{v} \ \forall \ n_\mathrm{u} < n_\mathrm{v}, \\ &j_\mathrm{u} < j_\mathrm{v} \ \forall \ |m_\mathrm{u}| < |m_\mathrm{v}| \wedge n_\mathrm{u} = n_\mathrm{v} \\ \text{and} \quad &j \ \text{is} \begin{cases} \text{even} & \text{if } m > 0 \text{ or} \\ \text{odd} & \text{if } m < 0. \end{cases} \end{aligned}$$

3.2.2. Statistical Analysis

While the modal composition as depicted above is extremely useful to describe the shape of the WF and perform numerical calculations, a more general description is preferred sometimes as well. Statistical evaluations of a single or multiple WFs can be useful e.g. for comparison purposes or to estimate the expected beam quality of a hypothetical AO loop. In the next sections, I will discuss some established ways to describe aberrations statistically in the spatial and temporal domain.

3.2.2.1. Spatial

Two common ways to characterize a WF in a single number that scales with the severity of the aberrations is to calculate the Peak-to-Valley (PtV) or the Root Mean Square (RMS) of the WF. The first one is intuitive in the sense that one immediately understands the amplitude of the largest phase deformations in the WF. However, it does not necessarily grant insight into the impact of the aberrations into the beam as the aberrations could in principle be zero for all except one sample on the WFS.

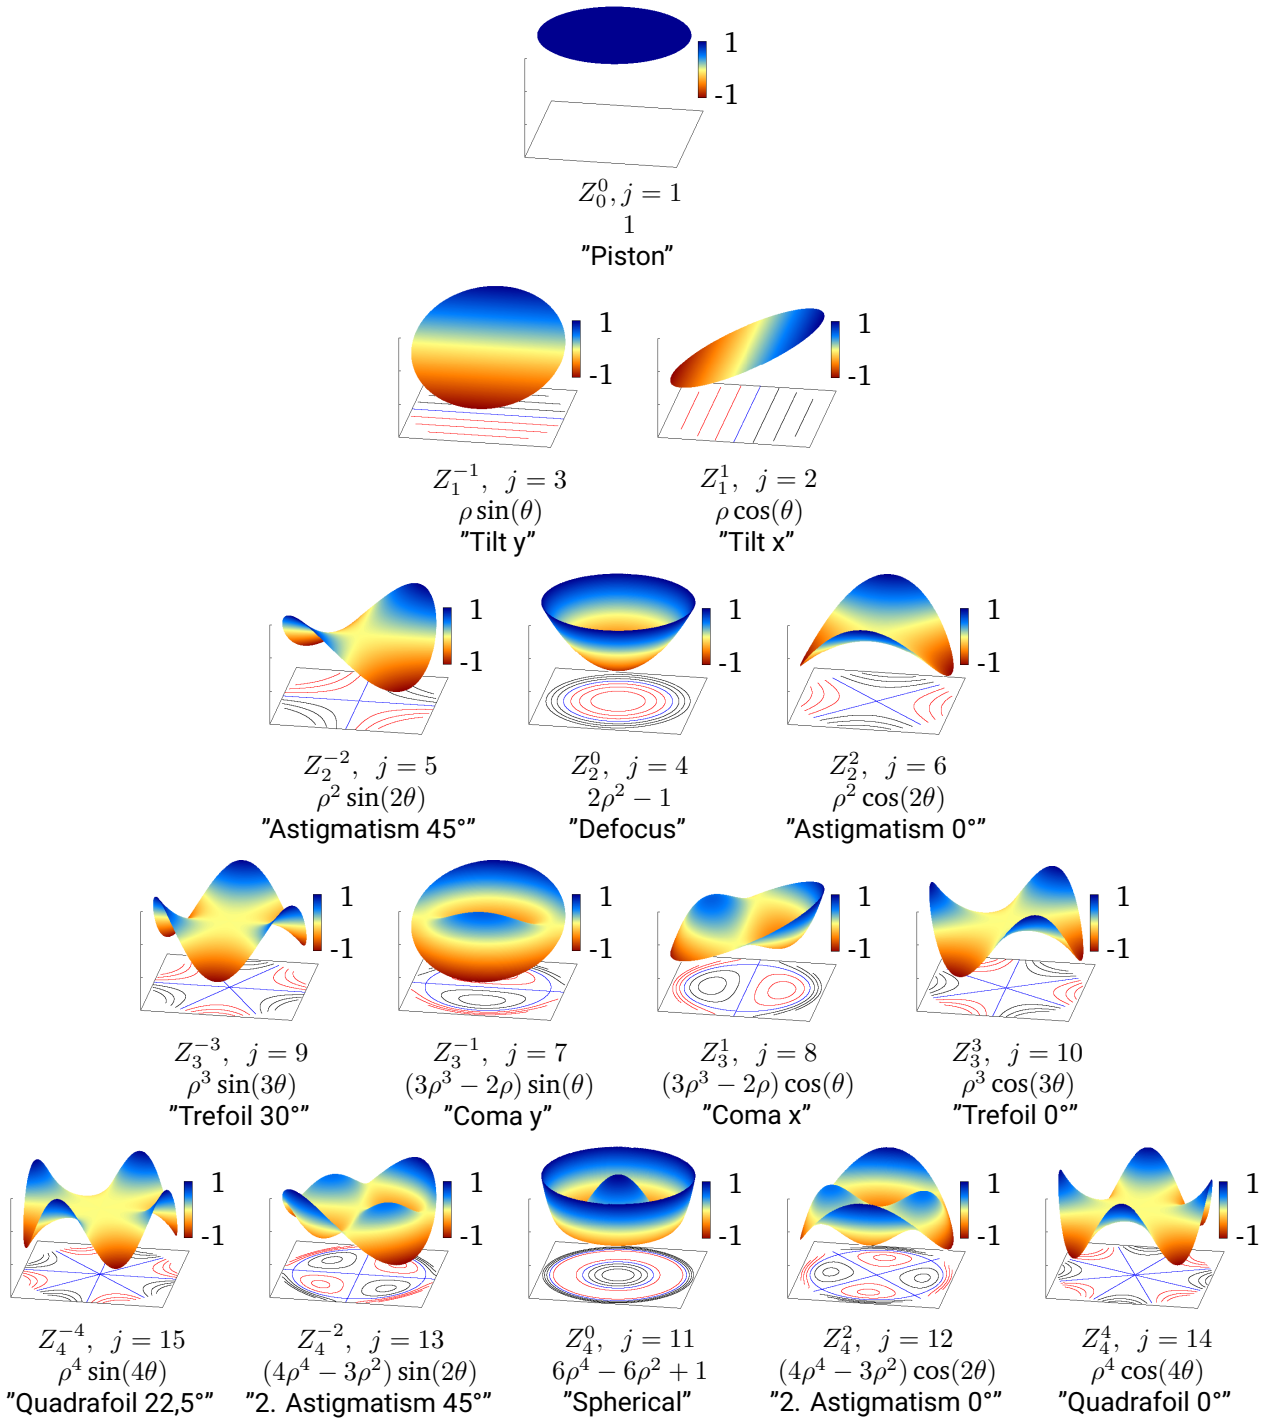

Figure 3.7.: The first 15 Zernike polynomials, normalized to $|R_n^m(1)| = 1$, including the Noll index j and the aberration name.

The intensity-weighted RMS is more suitable for this, as it is calculated by integrating over the whole aperture A, using the local intensity I:

$$\text{RMS}(W) = \sqrt{\frac{\iint_A I \cdot W^2 \ d^2A}{\iint_A I \ d^2A}}.$$

While the qualitative nature of the aberrations still gets lost in the process, it has been shown that the Strehl ratio Sr can be estimated directly from the RMS for small aberrations, independently from the manifestation of the aberrations, via

$$Sr(W) \approx 1 - \left(\frac{2\pi}{\lambda}\right)^2 \text{RMS}(W)^2, \tag{3.4}$$

where λ is the WL of the light [46, section 9.1.3]. This formula stems directly from the Fourier transform of the NF that calculates the FF in the case of Fraunhofer diffraction (see appendix B.2.3), where the exponential term in the transform integral was Taylor-expanded to the second degree. This indicates the limitations of the approximation: as soon as third order terms cannot be neglected anymore, deviations will occur. Note that equation (3.4) can be used for inhomogeneously illuminated pupils, as long as the intensity-weighted RMS is used.

Another way to describe WFs statistically in the spatial domain is to express the mean energy density spectrum of the spatial frequencies. This is more useful for AO than the PtV and RMS as this spectrum contains information on how severe the aberrations are on each scale. This is obviously more interesting for dynamically changing aberrations than for single WFs and is widely used in astronomy to mathematically describe the aberrations that are produced by the air turbulence over the telescope aperture.
A widely used spectrum to describe the turbulence in air is the Kolmogorov spectrum, where the power density of each three-dimensional, spatial frequency of air density fluctuations is proportional to its wave number to the negative power of five thirds [23, section 2.1.2]. After integration over the optical path, the Kolmogorov spectrum $\Phi(K)$ for the Optical Path Difference (OPD) is obtained:

$$\Phi(K) = 0.033 \ C_n^2 \ K^{-11/3}, \tag{3.5}$$

where K is the wave number of the spatial frequency and C_n^2 is the so-called refractive index structure constant, which is a measure of the strength of the turbulence. Note that equation (3.5) is only valid within the ”inertial range” of the turbulence, which is limited by the smallest and largest granularity, called the inner and outer scale, respectively. The inner scale is set by the viscosity of the medium, which converts the kinetic energy into heat, while the outer scale simply accounts for the largest variations. There are adapted versions of the Kolmogorov spectrum which take into account the changes near these boundaries, e.g. the Von Karman spectrum which saturates at the outer scale.

3.2.2.2. Temporal

Concerning dynamic aberrations (see section 3.4), it is useful to know how the WF is evolving. This can be useful to understand the qualitative behavior of a beam, e.g. during the cooling phase of the amplifiers (see section 3.4.1), which can be easily assessed by recording the evolution of certain modes over time. Another way to describe the WF evolution is to look at the spectrum of the temporal variations. From here, different useful insights can be gained, e.g. how large the expected WF error will be on-shot due to the delay between the last AO control step and the shot itself (see section 4.4.2).

The logical way of generating such a spectrum is to record the WF W over the aperture A for a given time period and calculate the average fluctuation energy density spectrum $P_W(\nu)$ over the aperture via

$$P_W(\nu) = \frac{\iint_A \left\| \widetilde{W} \right\| \; d^2A}{\iint_A d^2A}. \tag{3.6}$$

Here, ν is the fluctuation frequency and the $\widetilde{\ }$ denotes the temporal Fourier transform. Due to the normalization to the surface area of the aperture, the spectrum can be compared between arbitrary recordings, for example to compare the stability of the beams in different laser facilities.
For a more differentiated view, one can apply filtering to W before entering the calculation. This opens the possibilities to observe the spectra of only a subset of the WF deformations, e.g. low- or high-order aberrations, a certain set of Zernike modes and so on.
Of course, equation (3.6) demands a recording of the full WF, which may strongly limit the temporal resolution due to the acquisition speed of the camera. If no dedicated high-speed equipment is available, the frame rate can often be drastically increased by only reading out a subset A_s of the aperture A (and therefore only a subset of the camera chip). However, this comes at the cost of losing information, which usually means that the WF cannot be reconstructed from the measured gradients anymore.
In order to still assess the frequency components of the WF fluctuations (under the assumption that the spectrum is constant over the full aperture), one can calculate the gradient jitter energy density spectrum instead:

$$P_{\Delta W}(\nu) = \frac{\iint_{A_s} \sqrt{\left\| \widetilde{\frac{\partial W}{\partial x}} \right\|^2 + \left\| \widetilde{\frac{\partial W}{\partial y}} \right\|^2} \; d^2A_s}{\iint_{A_s} d^2A_s \cdot \sqrt{\iint_A d^2A}}. \tag{3.7}$$

Here, the numerator provides the integrated fluctuations of the gradient. The first term in the denominator normalizes this value to the area of integration. The second term is the square root of the full aperture area, providing a measure of the diameter for a given aperture shape. This accounts for the fact that a smaller aperture naturally means proportionally larger gradients for the same amplitude of aberrations. Therefore, the gradient jitter energy density spectrum given by equation (3.7) remains comparable for beams featuring approximately equal aperture shapes despite of their size. I will give an example for such a spectrum in section 3.4.2.
Equation (3.7) offers less flexibility than equation (3.6) due to the fact that only a sub aperture is recorded and no reliable WF filtering can be done. At the same time, its output has to be treated with care when thinking about AO as the spectrum only represents the WF slope instead of the local piston, which is the value of interest in this context. On the other hand, it does not suffer from potential errors that can be made during the reconstruction of the WF from its measured gradients. This includes the uncertainty of the total WF piston, which cannot be known.
Nonetheless, if either of the two spectra is known, one can identify a saturation frequency, which is the maximum frequency in which the WF changes temporally.

Up to this point in this chapter, I have introduced the architecture of HI laser systems, as well as different ways to describe aberrations. This knowledge serves as foundation for the main topic of this chapter: the composition of the aberrations that are typically encountered in HI laser systems. In the following, I will discuss static, dynamic and chromatic aberrations, giving both possible causes as well as qualitative manifestations.

3.3. Static Aberrations

Static aberrations are characterized by being constant over an application-dependent patch of time. For different cases, this may span from a brief period in which a task is completed up to the whole lifetime of an optical system. In HI laser systems, aberrations are usually considered to be static if they persist over the course of an experiment.

3.3.1. Misalignment

One common source of static aberrations is the misalignment of optical components, e.g. lenses, relative to the optical axis of the beam. Theoretical work on this topic exists for a long time now and has been refined and specialized continuously. One of the cases that has been considered frequently is the misalignment of rotationally symmetric spherical optics, as this is encountered in the majority of optical applications. A theory worth mentioning here is the Nodal Aberration Theory, originating in the work by Thompson and Shack in 1980 [52], which describes the aberration field of optical systems in terms of the position of unrefracted rays. This theory has been expanded to freeform optics later on. For smooth optics, these aberrations are typically composed of low-order WF deformations.
One case which is relevant for this work is the effects of misalignment of a reflective off-axis parabolic telescope on the WF. I derived analytical expressions for the first order approximations [1] and summarize my findings in section 5.2.3.

3.3.2. Deformations

Another source of static aberration is the deformation of the optics. This is an issue which becomes more notable for larger optics, peaking in the necessity for active optics in large telescopes (see section 2.1.2). The causes for these static deformations can be very different, but two common causes are mechanical stress due to the mounting of the optics and the sag due to the sheer weight of the bulk material.
An example for the former is a Trefoil aberration that can be generated by a mirror mount with three contact points to the substrate. An example for the latter is the copper OAP that is often used for focusing in the PHELIX target chamber. The material is not ridgid enough to sufficiently support its own weight, introducing aberrations depending on the orientation of the OAP.
As for the misalignment, the order of the aberrations is usually low, but may increase with the size of the optics.

3.3.3. Imperfections

The last common source of static aberrations is the quality of the optics themselves. Deviations from the optimum may be present either due to the non-ideal choice of the optic, manufacturing defects or damages.
An example for non-ideal optics is the use of spherical lenses. These are easy to manufacture and therefore cheap, but can only approximate the ideal shape for common applications. The induced aberration depends on the setup, but commonly features large amounts of spherical aberrations (Z_4^0, see Fig. 3.7).
An example for possible manufacturing imperfections of large optics in HI lasers are the deformations on the large scale diffraction gratings in the pulse compressor [53], which are due to the very challenging manufacturing process of these optics. While low order aberrations like astigmatism are often encountered, high order variations can also be significant. This poses a challenge for beam correction.

Lastly, the manufacturing of aspherical optics, e.g. OAPs, is also likely to induce high order aberrations. An example for this is shown in Fig. 3.8. These are the OAPs which I used at PHELIX in the verification experiment of this work (see chapter 6). Depending on the material, the precision can be very different: in our example, the OAP made from copper features significantly larger aberrations than the one made from glass.

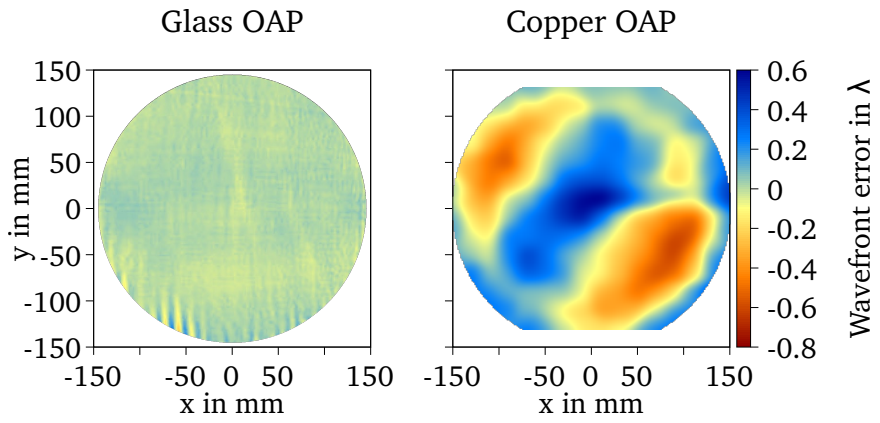

Figure 3.8.: The reflected WF errors of a glass OAP (left) and a copper OAP (right), which are used at PHELIX, provided by the manufacturers. Tip/tilt, defocusing and astigmatism were subtracted from the datasets, as these can be largely eliminated by proper alignment. Even though the data of the copper OAP were given with a low spatial resolution, the large amplitude of the surface imperfections immediately becomes apparent.

3.4. Dynamic Aberrations

Contrary to the examples above, dynamic aberations evolve temporally. The time constant of this evolution can range from slow drifts that introduce notable changes over the course of the operation up to high frequency jittering of hundreds of Hz or more, depending on the source. In the following, I will touch on four sources of dynamic aberrations that are often encountered in HI laser systems.

3.4.1. Temperature drift

Thermal expansion can make temperature drifts impact the properties of an optical system. Changes in temperature can originate from a non-ideal control of the air-conditioning in the laboratory or, in the case of telescopes, the evolution of the ambient temperature. These changes slowly transfer to the equipment and may cause optical surfaces to shift with respect to each other, lead to the deformation of tables etc., which effectively has an impact on the spacing and orientation of the optics. Similarly to misalignments, this causes primarily low-order aberrations, which follow the temperature curve of the environment.
In laser systems with high average power, the optics themselves may heat up over time and deform accordingly, causing notable changes in the beam quality. As an example, it has been observed that the compressor gratings of a high repetition rate CPA system with an average power of 25 W can heat up during operation and deteriorate the quality of the compressed pulse [54].
Furthermore, high power laser systems often incorporate a significant amount of heat into the optical

gain media during the pumping phase, which has to be dissipated afterwards. This cooldown process generates continuously changing aberrations, which asymptotically approach the baseline level (see section 4.2.1 for an example figure). Depending on the size, shape, medium and cooling scheme of the amplifier, the time until the aberrations decreased to an insignificant amount may vary from seconds to days, which means that the thermal load of one shot can affect the WF of the next shot significantly. Qualitatively, these aberrations are dominated by low order terms as well. This is due to the strive for a homogeneous or at the very least symmetrical distribution of pump energy over the aperture of the laser beam.

3.4.2. Air Turbulence

While the temperature drifts feature slow to moderate time constants, air turbulence can vary much faster: even though air has a refractive index close to one (as in vacuum), density fluctuations still change the optical path through the corresponding patch of space. These fluctuations can either be generated by a turbulent air flow or by temperature gradients, which cause convection. As discussed in section 3.2.2, the spatial spectrum of the WF caused by air turbulence can be described by a power law over the spatial frequency, e.g. by a Kolmogorov spectrum, with a given amplitude, indicating that smaller beams suffer less from air turbulence than larger ones. The temporal evolution, however, can be extremely different depending on the circumstances.
In astronomy, the common measure for the upper frequency limit is described in terms of the Greenwood frequency, which determines the minimum rejection bandwidth that an AO system has to feature in order to properly compensate the aberrations. The Greenwood frequency can be calculated by an integral over the optical path (ground to sky), where each infinitesimal layer of air contributes to the frequency by its specific refractive index structure constant C_n^2 and wind velocity. Here, the frequency can range from 20 Hz to over 600 Hz [23, chapter 2].
This definition is less useful for HI lasers as the turbulence induced aberations are rather dominated by the evolution of the turbulence itself than the wind speed over the aperture - HI lasers luckily don't operate under the plain sky. Here, no standard measure for the maximum temporal frequency has been established yet. One way to characterize this frequency is to use the fluctuation spectra given by equation (3.6) or (3.7), integrate over the frequency axis and look for a point of saturation.

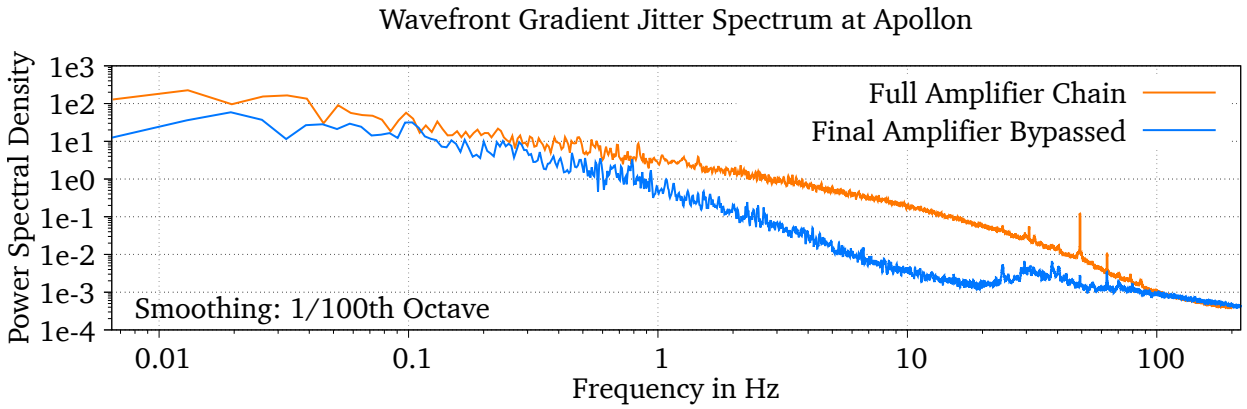

Figure 3.9.: The energy density of the temporal evolution spectrum of the WF gradient of the main beam at Apollon, measured for the beam passing all stages (orange) and all stages without the final amplifier (blue). The spectrum shows how strong the WF moves at each frequency.

An example for a spectrum, given by equation (3.7), can be seen in Fig. 3.9. Here, I recorded a narrow stripe of the WF in the main beamline of the Apollon laser system [55] with a repetition rate of slightly above 220 Hz for the full amplifier chain (orange curve) and a configuration bypassing the final amplifier (blue curve). Both spectra resemble a negative power law for the lower half of the spectrum, drop off at some point and reach a plateau in the highest frequencies. The difference between the curves indicates the contribution of the main amplifier at Apollon to the turbulence induced aberrations, which persists up to several tens of Hz. Some frequency spikes indicate vibrations in the system.
After integration (Fig. 3.10), the saturation of the spectra becomes more obvious. While the saturation is hard to pinpoint for the bypass configuration (blue curve) due to the lower signal-to-noise ratio, the curve for the full amplifier chain clearly reaches saturation (99 %) at approximately 70 Hz.

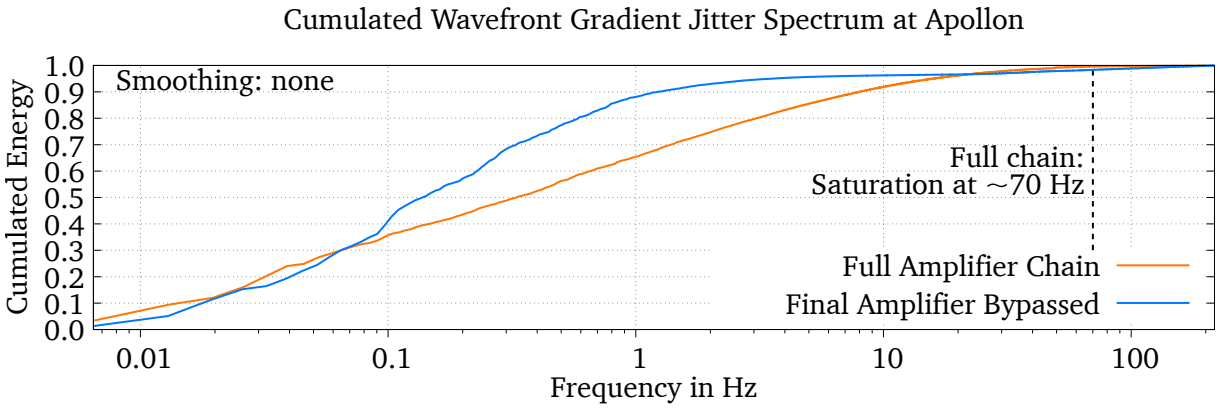

Figure 3.10.: The normalized integrals of the spectra shown in Fig. 3.9. The curve for the full amplifier chain saturates at approximately 70 Hz. The curve for the measurement where the final amplifier was bypassed does not saturate, which is due to the noise background of the measurement.

The turbulence at Apollon is a rather extreme case, which is mostly due to the bow-tie configuration of the main amplifier and the pump-beam geometry that make a proper housing hard to implement. Other laser systems may feature a lower saturation frequency. This does, however, not mean that there is no interference with the operation, even tough it is less challenging to AO. One example for this is the PHELIX laser. Here, the heat which accumulates in the main amplifier over the course of an experiment day causes convection in the gaseous nitrogen between the Nd:Glass slabs, which changes the WF over the course of several seconds. While being a lot slower than at Apollon, this is still a significant challenge when it comes to on-shot WF correction in the heated system (see section 4.4.2).

3.4.3. Vibration

In Fig. 3.9, spikes appear in the upper frequencies of the gradient jitter spectrum. The small width of the peaks indicate that the source of this jitter is a periodic movement, which can be described as a vibration in this frequency range. The cause for such features, aside from measurement noise, is then likely to be a mechanical component, coupling vibrations into a part of the optical setup. Examples for this are cooling fans, vacuum pumps and power supplies. Once the frequency of such a feature is known, the source can be searched out and decoupled from the optical setup to get rid of this feature. Other, non-oscillating sources of vibration can be traffic or construction sites, which generate vibrations

in the ground. These are generally harder to treat and active beam control can be a viable option to mitigate the effects on the beam. While vibrations mostly cause jitter in the beam pointing, this can be transferred to higher order aberrations by passing through lenses and other optics on a non-ideal path [23, chapter 2].

3.4.4. Shot Aberrations

A problem concerning dynamic aberrations which is specific to pulsed laser sources is the presence of shot aberrations. This term describes aberrations which are exclusively present in the systems during the shot, caused by the change in the index of refraction of active media when being pumped - sometimes, the phrase ”pump-induced aberrations” is used synonymously. As the WF deformations only appear on-shot, they have to be pre-compensated by active optics and/or AO, which requires accurate knowledge of the expected WF offset. At the same time, this generates the necessity to measure the WF in a laser system not only with a low energy alignment beam, but also on-shot, which poses a challenge on the design of suitable WF measurement systems (see section 4.4.1).

3.5. Chromatic Aberrations and Spatio Temporal Couplings

The last category of aberrations that are often encountered in HI laser systems are the ones where the WF is a function of the WL, i.e. each WL features a different WF. This can have a variety of effects on the beam quality and are potentially be difficult to treat, depending on their manifestation.
In general, these aberrations can be divided into two types. The first one can be characterized by a clear correlation between the frequencies, which I will call *chromatic aberrations* in the following. The second one does not feature such correlations and is therefore commonly much more complex and more difficult to compensate. I will refer to this type as *Spatio Temporal Couplings (STCs)*. Note that there is no clear naming convention on this in the field and I will only use these terms here to differentiate between the two.

3.5.1. Chromatic Aberrations

Chromatic aberrations can have a variety of sources and manifestations, but are commonly caused by either unwanted dispersion inside optical media or misalignment in dispersing setups, namely the stretcher or compressor of a CPA system [56].
An example for the former is chromatic focusing, caused by the frequency-dependence of the focal length of a lens. If no countermeasure is taken, this effect accumulates over the course of the amplifier chain, where the beam is repeatedly magnified using refractive telescopes. This causes a pulse-front curvature in the NF and deteriorates the focal spot in time and space [57]. Compensation strategies include the use of achromatic lenses or dedicated setups, where the beam repeatedly passes through a lens with a negative focal length in order to pre-compensate the chromatic focusing of the main beamline.
Dispersion can also generate chromatic aberrations when the beam passes through wedged optics, where one side of the beam is dispersed more than the other. This results in angular chirp, where the WFs of the different frequencies are tilted with respect to each other. This causes a pulse-front tilt in the NF and a spectral separation in the FF, but can easily be compensated by inserting an optical element with the opposite wedge.
Furthermore, misalignment in the stretcher or compressor can generate angular chirp as well, or even introduce spatial chirp. Here, the beam is spectrally separated in the NF, which can cause complications for AO (see last paragraph of this section).

3.5.2. Spatio Temporal Couplings

STCs do not feature such correlations, but can accumulate over the course of the beam transport of a HI laser nonetheless. The largest contributions are commonly generated in the stretcher and compressor of CPA lasers due to the spectral separation of the beam [56]. Here, imperfections of the optics are directly converted into STCs. This is especially problematic in the pulse compressor as the beam features a large diameter, which means that the aberrations which are commonly present on the large gratings [53] cause STCs of high spatial frequencies. The compensation of STCs is a challenging problem and still under investigation up to this day [58].

3.5.3. Chromatic Aberrations and STCs at PHELIX

Naturally, lasers with a larger bandwidth suffer more from chromatic aberrations and STCs. At PHELIX, the bandwidth is relatively small compared to titanium sapphire lasers which generate ultrashort pulses of tens of femtoseconds or less. The intensity reduction due to chromatic focusing was estimated to be about 30 % [57], while purely spatial WF aberrations reduce the intensity several times more and thus remain to be the limiting factor. The compensation of chromatic aberrations is therefore only of interest after the WF aberrations have been removed using a stable AO setup.
An exception to this is the compressor geometry of PHELIX, which generates a strong spectral separation in the outgoing NF (see section 3.1.4). This poses a problem for the AO loop itself as aberrations on the first grating of the compressor are spectrally smeared out during compression and cannot be compensated afterwards. There is no way to reverse this effect except for building a new pulse compressor, which is why I focused on compensating the mean aberrations on the beam instead.

Conclusion

In this chapter, I have discussed the nature of aberrations in HI laser systems. As the main aspect of this work is to compensate them using AO, I will build on this knowledge in the next chapter, deriving the challenges that AO has to face in HI lasers. This will be the last chapter on the fundamentals, after which I will move on to the work that I have done at PHELIX.

4. Requirements on Adaptive Optics in High-Intensity Laser Systems

In this chapter, I will cover some of the peculiarities of Adaptive Optics (AO) in Chirped Pulse Amplification (CPA) laser systems compared to AO in other applications, like astronomy. This will on one side help to understand why AO remains subject of ongoing investigations in such laser systems up to this day, mainly due to the diversity of laser facilities which renders the optimal solution for one system less useful for the other. On the other side, awareness of these challenges is important to understand the design decisions for the post compressor AO at Petawatt High-Energy Laser for heavy Ion eXperiments (PHELIX), described in section 5.2.
The structure of this chapter will build on the knowledge about the architecture of modern High Intensity (HI) laser systems from section 3.1 as the challenges for AO in these facilities are tightly linked to some of the core concepts like CPA. Therefore, I will start in section 4.1 by discussing the parts of the laser system where the aberrations in the beam should be minimized. We will see that it makes sense to have two AO loops in CPA laser systems, divided into a pre- and a post-compressor loop. Both of them feature unique challenges, which I will describe in sections 4.2 and 4.3. After that, I will summarize remaining common challenges related to the shot procedure in section 4.4.

4.1. Critical Locations in CPA Lasers

In general, the goal of WaveFront (WF) control in CPA lasers can be reduced to the minimization the aberrations in the focal plane. Commonly, this translates to approaching a spherical incoming wave in the target chamber and keeping the Spatio Temporal Couplings (STCs) at a minimum.
As discussed in section 3.5, STCs are generated in refractive optics independently from the incoming WF, which makes this part irrelevant for AO, or wherever the beam is spectrally separated while interacting with optics. The latter happens in the stretcher and the compressor of a CPA laser system. However, as higher order aberrations are easier accumulated for larger beams, the compressor is of much larger concern compared to the stretcher.
As discussed earlier, there are basically three ways in which STCs can be generated in grating-based pulse compressors: misalignment of the gratings, grating imperfections (either from manufacturing or heat deformations) and imperfections of the incoming WF. The first one is of no concern for AO but for the corresponding alignment procedure. Concerning the second one, Li et al. proposed to precompensate the STCs generated by grating imperfections [58] by adding a smaller compressor with identical geometry early in the system and deforming the gratings in an according manner.
The third source of STCs in the compressor is the beam quality of the WF entering the compressors. The fact that a beam which is not well collimated evolves pulse front tilt, a certain kind of STCs, is well documented in literature [59, 60]. This implies that the same happens for local WF curvatures, i.e. higher order aberrations, which leads to complex STCs over the full aperture [12] which cannot be corrected afterwards. This clearly shows that a good WF quality is crucial in order to mitigate STCs after the pulse compressor, making the compressor entrance one of the critical locations in CPA laser systems.

The other part of the low aberrations is the purely spatial component, i.e. the WF of the beam converging to the target plane. If one is interested in the highest possible intensity, this WF should be spherical. Any deviations from this alter the energy distribution in the focal plane and lead to a lowered peak intensity. Here, one encounters a conflict: as I discussed above, the WF at the compressor entrance should be as flat as possible in order to avoid STCs. However, the beam transport and focusing after the compressor do introduce aberrations as well. Pre-compensating these with an AO loop prior to pulse compression would come at the cost of adding STCs to the beam. This clearly shows that a CPA laser system should ideally run two independent AO loops: one for the compressor entrance and one for the WF in the target chamber [61]. An example is shown as a block diagram in Fig. 4.1. As these loops face different conditions, the according challenges are different as well, which I will discuss in the next two sections.

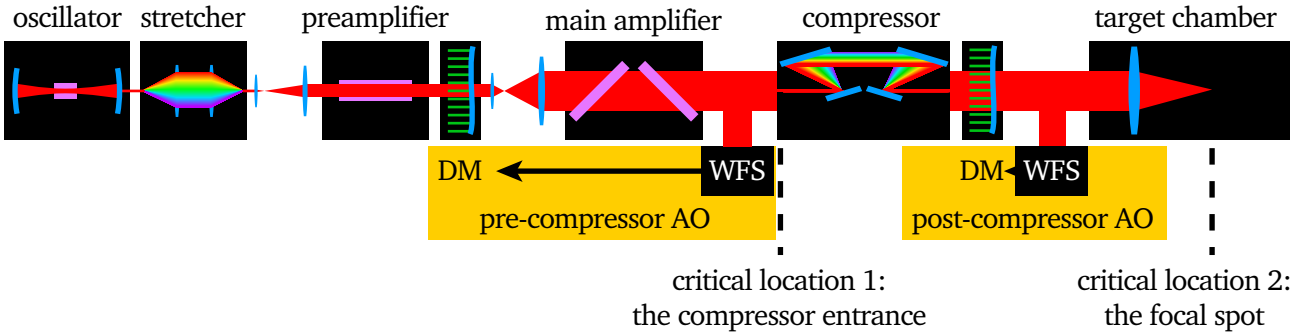

Figure 4.1.: Block diagram of a possible CPA laser chain, featuring a pre-compressor and a post-compressor AO loop.

4.2. Pre-Compressor Challenges

Following the laser architecture, I will first describe the challenges that a pre-compressor AO loop faces while optimizing the WF at the compressor entrance.

4.2.1. Heat Accumulation

Some HI laser systems feature components which gradually heat up during operation. The most relevant components are the amplifying media as the energy deposited by the pumps does not fully contribute to the amplification process. The exact amount of heating depends on a variety of factors such as the quantum efficiency of the medium, the pumping technology, the deposited energy and so on. At the same time, the impact on the transmitted WF quality depends on the host material of the medium in terms of the thermal refractive index, the thermal expansion coefficient and, of course, the characteristics of heat dissipation. The latter is more problematic for glass-based amplifiers as crystal lattices (e.g. sapphire or Yttrium Aluminum Garnet (YAG)) feature much higher thermal conductivity.
The nature of these aberrations additionally depend on the geometry of the amplifier and often includes defocusing (for rod amplifiers) and astigmatism (for slab amplifiers, where the beam passes a plano-parallel slab of material in the Brewster angle), but are mainly composed of low order deformations. The timescale on which thermal aberrations change is dominated by the cooling behavior and may span from seconds to hours or even days.
An example for the effects of heat on the transmitted WF of a rod amplifier is shown in Fig. 4.2 in terms of Zernike coefficients. Apparently, the amplitude of defocusing dominates the WF.

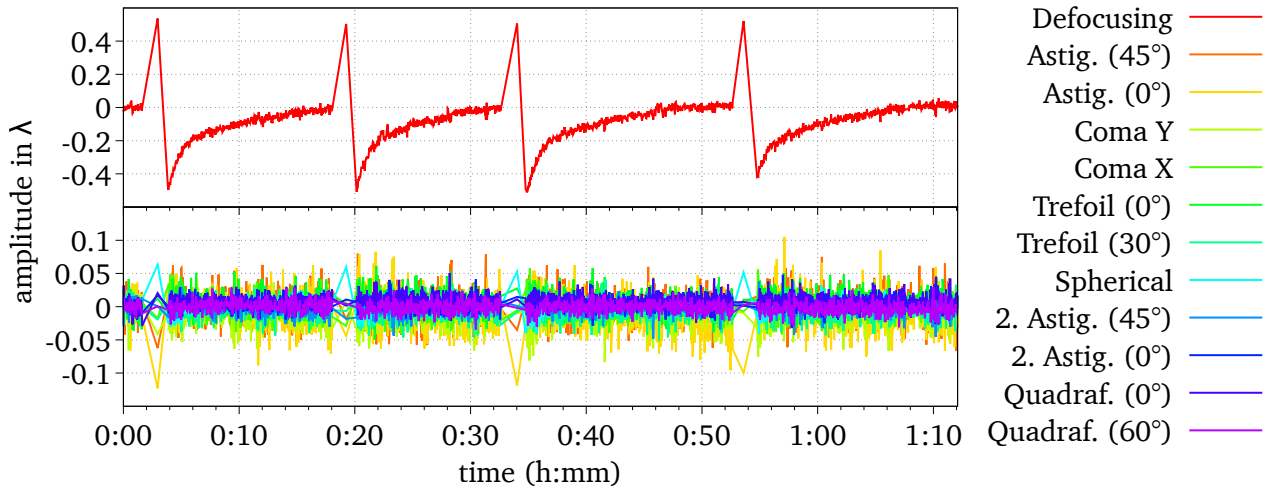

Figure 4.2.: The Zernike coefficients of the transmitted WF of the preamplifier at PHELIX over the course of four test shots. The top graph shows the strength of the defocusing, while the bottom graph contains the coefficients of the subsequent 11 polynomials.
The four spikes in the top graph are generated by the shot-aberrations, while the following exponential slope corresponds to the cooldown of the amplifier rods.

However, the smooth and predictable evolution of these aberrations does not pose a significant challenge for AO by itself. The issue is a different one: as the main point of a pre-compressor AO loop is to flatten the WF at the compressor entrance, the WaveFront Sensor (WFS) is located in close proximity to that point. Deformable Mirrors (DMs) of this diameter on the other hand come at significant costs, which is the reason why placing a much smaller DM earlier in the system, pre-compensating the aberrations of the larger amplifiers, makes sense from an economical point of view. In this case, dynamic variations between the DM and the WFS may negatively impact the validity of the response that was recorded to run the control loop (see chapter 2). For large amounts of thermal aberrations, this may even lead to unstable loops.
A strategy to counteract this is to remove low order aberrations with high amplitude from the beam, e.g. by moving motorized lenses.
Depending on the architecture of the amplifiers, another problem may arise from the heated components: the faces of the material may heat up ambient gas. This causes air turbulence, which gets worse with higher amplifier temperatures.

4.2.2. Air Turbulence

In most laser systems, the beam will propagate through turbulent air at some point during the pulse amplification. As discussed in section 3.4, this leads to chaotic aberrations which are in general unpredictable and can widely vary in spatial and temporal frequency. An example for the temporal spectrum of the average WF slope change over the aperture is shown in Fig. 4.3 on the left. This is the spectrum of the main beam of the Apollon laser facility in France [55] after passing all amplifiers and I describe how I recorded this data in more detail in section 3.4.2.
The spectrum roughly resembles the spectrum of a Brownian-like movement (random walk) for the lower half of the frequency axis. The integrated and normalized spectrum is shown on the right, where the energy saturation for the full amplifier chain is visible at about 70 Hz, indicating a cutoff frequency.

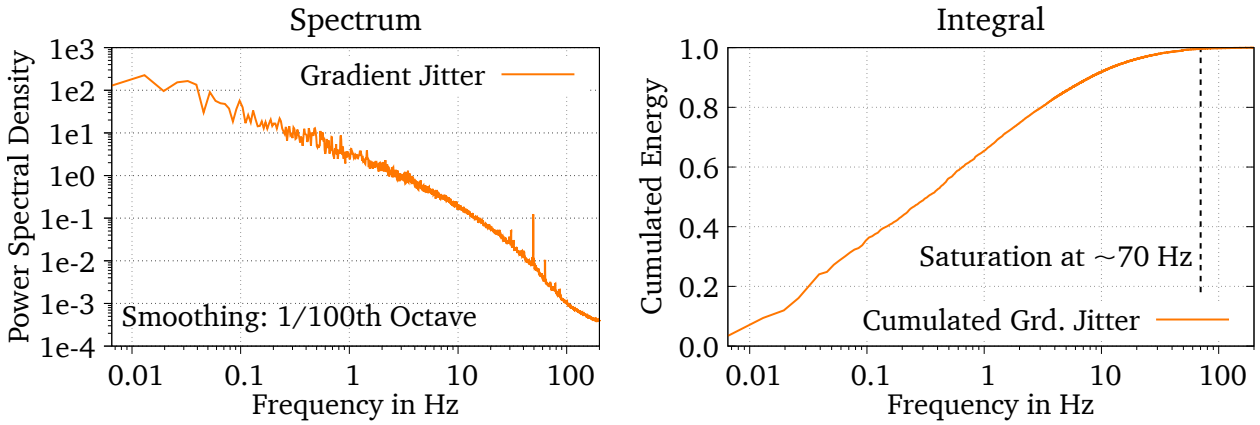

Figure 4.3.: Left: the temporal spectrum of the mean local WF slope change (Gradient Jitter) of the beam at the exit of the amplifier chain at Apollon. Right: the normalized integral over the spectrum to the left, where a clear saturation is visible. I explained the data and how I recorded them in more detail in section 3.4.

The challenges for AO that arise from this matter is twofold. The obvious issue is that a control loop that shall compensate turbulence has to run at least one order of magnitude faster than the timescale of the Optical Path Difference (OPD) fluctuations at any given position [62]. For the example mentioned previously, this would require almost kHz repetition rates. While this is generally possible, as real-time AO in astronomical telescopes demonstrate, HI laser systems are pulsed and do therefore not deliver a continuous signal to the WFS. This may be tolerable for systems with kHz repetition rates, but is not easily applicable to single-shot laser systems. On top of that, the shot sequence in these systems usually involves a time where no light reaches the sensors prior to the shot. I will go into detail on this topic in section 4.4.2.
Due to these reasons, aberrations caused by air turbulence are rarely treated in HI laser systems today, even though some efforts are made [63] or are in progress. Instead, passive methods, e.g. beam housing, are employed to keep the motion of air at a minimum.
However, turbulence-induced aberrations can pose another challenge for pre-compressor AO, even if one decides to omit real-time correction: as I discussed in the last section, laser systems tend to heat up with every shot, introducing slow varying thermal aberrations. While these can in principle be corrected with a slow AO loop, air turbulence can add a significant amount of noise to the WF measurement, leading to bad correction results or even destabilizing the loop.
This issue is extremely hard to circumvent: as the spectrum of turbulence-induced aberrations features large amounts of low frequencies (compare Fig. 3.9), averaging over a certain timespan will only reduce the noise of the WF measurement in part. The only other option would be to implement a real-time AO loop that feeds on a WF measurement of a Continuous Wave (CW) pilot beam instead of the pulsed main beam in order to compensate for the turbulence as well.

4.2.3. Beam Diameter

The established WFS technologies (see section A.1) that are evaluated by a computer are based on a digital camera, which means that their aperture size is limited and rarely exceeds 10 mm in diameter. As discussed in section 3.1, the beam size in HI lasers is scaled up successively in order to reduce the intensity and thus allow larger energy throughput. This implies that the beam needs to be demagnified

before entering the WFS.
The issue here is that one is ultimately interested in the aberrations of the large beam, which can easily become obscured by the aberrations of the demagnification and imaging system. When running a closed AO loop, these aberrations will be added inversely to the large beam.
The natural countermeasure is to perform a suitable calibration routine for the WFS, which is a challenge on its own. A suitable approach which has yet to be demonstrated in large scale beams is proposed by Dorrer et al. [64]. The method uses a large, apodized Hartmann mask in the beam to generate spots, which can be recorded using a regular, imaged Near Field (NF) camera. When shifting the mask along the beam propagation direction, the spots will be displaced. I will describe this method in section A.1.2 in more detail, including a schematic for better comprehension. However, the method has not been demonstrated on large scale beams up to this point and while it may be feasible, it requires a lot of space that may not be available in all laser systems.
This issue is somewhat analogous to the unavailability of the focal spot after the pulse compressor (see section 4.3.2).

4.3. Post-Compressor Challenges

The pulse compression stage in CPA lasers is commonly under vacuum as the light intensity drastically rises. This means that the issues that are associated with air turbulence, as described in the last section, do not occur here, potentially allowing for slow yet stable AO technologies. That being said, the beam stability at the compressor entrance does influence the beam after compression, which means that the disadvantages of an unstabilized beam do persist if no countermeasures were taken.
Aside from that, other problems arise here, which I will describe in the following sections.

4.3.1. High-Intensity Beam Transport

As the goal of pulse compression is to retrieve the highest possible power of the pulse by removing its quadratic spectral phase [11], the intensities in the NF rise to high values, often exceeding $10^{10}\,\mathrm{W/cm^2}$ on-shot in the NF. This generates a set of challenges to the beam transport.
Aside from using vacuum-compatible components, the main relevance for AO in these terms arise from the necessity of imaging the surface of the DM onto the WFS. At this intensity level, transmitting optics are not suitable anymore due to the intensity-dependent, non-linear index of refraction of the material. Passing the beam even through a single lens will deteriorate the beam significantly as WF deformations, self-focusing and self-phase-modulation will occur. This means that if relay-imaging of the main beam shall be done, e.g. due to a large distance from the DM to the target area, the imaging has to be done in an all-reflective manner. Up to this day, this is rarely done in HI laser facilities, even though it permits using stable post-compressor AO loops in some cases. The major reasons for this are the expenses for large, curved optics and the alignment which poses significant challenges. I have done extensive analytical considerations on the latter [1] and will discuss my findings in section 5.2.3.
But even if the distance from the DM to the WFS is small enough such that no relay imaging in the main beam is necessary, one has to stay wary about the high intensity in beam diagnostics as well: even the part of a laser pulse transmitted through a leaky mirror with 1 % transmission or less can result in a nonlinear OPD of more than one radian, which is comparable to the aberrations that AO shall remove from the beam. This indicates that thin mirrors, made of a material with a low non-linear refractive index, with a low transmission have to be used for post-compressor beam sampling.

4.3.2. Unavailability of the Focal Plane

Another challenge, similar to the beam diameter issue which I described in section 4.2.3, is the fact that the target focal spot cannot be observed on-shot. This is inherently different from astronomy, where a camera is located in the point of interest, giving direct feedback if the AO loop compensates aberrations well or not (see Fig. 4.4, left).

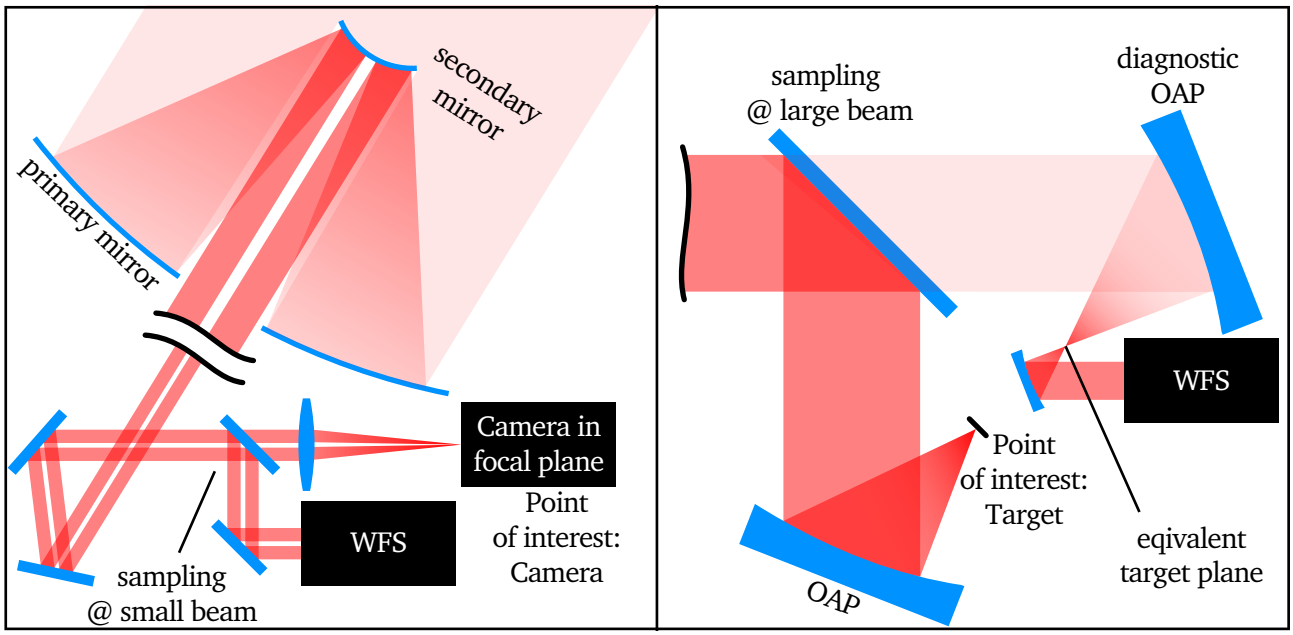

Figure 4.4.: Schematics of a WF diagnostic setup in an astronomical telescope (left) and in the target area of a HI laser system, featuring an equivalent target plane (right).

A widely used approach to mimic this is to set up a so-called ”equivalent target plane”, where a small energy fraction of the beam, sampled by a leaky mirror, is focused using an identical focusing optic, replicating the focal spot in the target chamber [65] (see Fig. 4.4, right). The main issue here lies in the Off-Axis Parabolic mirrors (OAPs), which are used for focusing almost exclusively nowadays. Each of these feature individual surface errors (much larger than for spherical optics due to the manufacturing process), limiting the comparability between the focus on target and in the equivalent target plane.
Furthermore, even if no OAPs are used, the beam diameter at the sampling position is large in HI laser systems due to the high beam intensity. This makes any successive imaging system prone to aberrations. In astronomy on the other hand, the small beam diameter at the sampling position allows for the use of optics with long focal lengths, which generally induce less aberrations.
This indicates that a reliable calibration routine of the AO system to the WF quality of the final focal spot is crucial as this area is unavailable on-shot.

4.4. Single-Shot-Related Challenges

Aside from the specific challenges for AO in CPA laser systems that arise prior to and after the pulse compression, there is a set of challenges linked to the way in which single shot lasers are operated. These do not apply for laser systems that run with repetition rates that are sufficiently high to run closed

AO loops only on the amplified laser shots, e.g. systems with kHz repetition rates.
As these points do apply to pre- and post-compressor loops, I will discuss these issues in this section without differentiating between the two.

4.4.1. Dynamic Range

Single-shot laser systems cannot run closed loops with on-shot WF measurements only, as the time span between two successive shots is too large to accurately predict the evolution of dynamic aberrations, e.g. the cooling of the amplifiers. This implies that the AO loop has to run prior to the shot, pre-compensating the aberrations that are expected on-shot.
Naturally, the pulse energies that propagate through the system while the amplifiers are inactive are several orders of magnitude lower than the ones obtained with pumped amplifiers. At PHELIX, for example, the alignment pulses deliver an energy of 400 µJ in the target chamber, while a full-energy shot can deliver up to 200 J, which is a factor of 5E+5 more.
For the beam correction, the closed loop has to run on the low-energy beam, while the on-shot WF must to be recorded in order to verify the validity of the predictions. This means that a suitable WFS needs to be equipped with an energy attenuation system that can cope with this dynamic range. A variety of approaches can be used here, ranging from simple reflective or absorptive filters and adjustable polarizers to using secondary internal reflexes of wedged mirror substrates [66] (see Fig. 4.5).

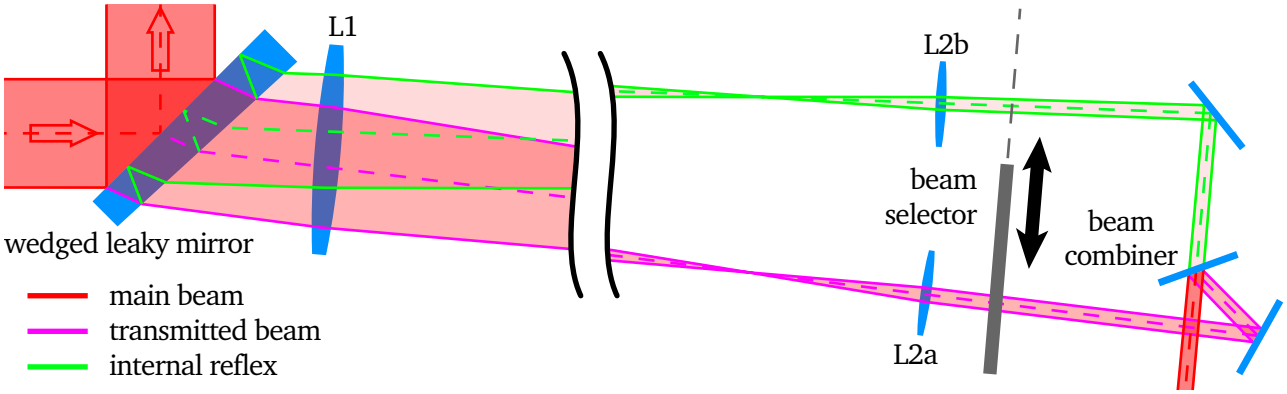

Figure 4.5.: Scheme of a possible use of internal reflexes in a wedged leaky mirror for energy attenuation. The rear surface of the mirror does not feature a highly reflective coating, making the internal reflex (green lines) much weaker than the primary transmitted beam (pink lines). After demagnification using a lens (L1), the two beams are separated and recollimated with separate lenses (L2a, L2b). A movable shutter then blocks either one or the other beam. After that, the beampaths are recombined with a beamsplitter after the same propagation distance. A very similar approach is used in the COmpressor Sensor (COS) at PHELIX.

The main challenge here is to design the attenuation system such that the transmitted WF stays identical despite the changing attenuation factor. While filters and polarizers are rather safe ways, splitting up the beampath has a high potential to feature different transmitted WFs and thus have to be considered with care.

4.4.2. Shot Delay

The next issue related to the delivery of WF-corrected single shots consists of the delay between the shot and the last WF measurement of the AO system. The origin of this delay may be different for each system. As an example, the Apollon laser system is able to continuously deliver one shot per minute. The unamplified beam cannot be measured in between as the energy attenuation measures are in place, which means that an AO system could only run on one WF measurement per minute as well.
In contrast to this, the PHELIX laser can only shoot every 90 minutes, which gives plenty of time to remove the energy attenuators and run the AO loop with 10 Hz repetition rate. However, the shot procedure involves moving all attenuators back into place, which renders the WF diagnostic blind for several minutes before the actual shot is delivered.
Both of these examples demonstrate that chaotic dynamic aberrations on a timescale of a minute or shorter (e.g. as shown in Fig. 4.6) cannot be compensated using an AO loop feeding on the main beam, which greatly differs from other AO applications, such as free-space communication, astronomy or retinal imaging. Solving this issue is challenging as the only way to circumvent it is to introduce a pilot beam to the system which can be separated from the main beam and thus may bypass the attenuation stages, continuously feeding the AO loop. To the best of my knowledge, this approach has not been implemented in HI laser systems yet.

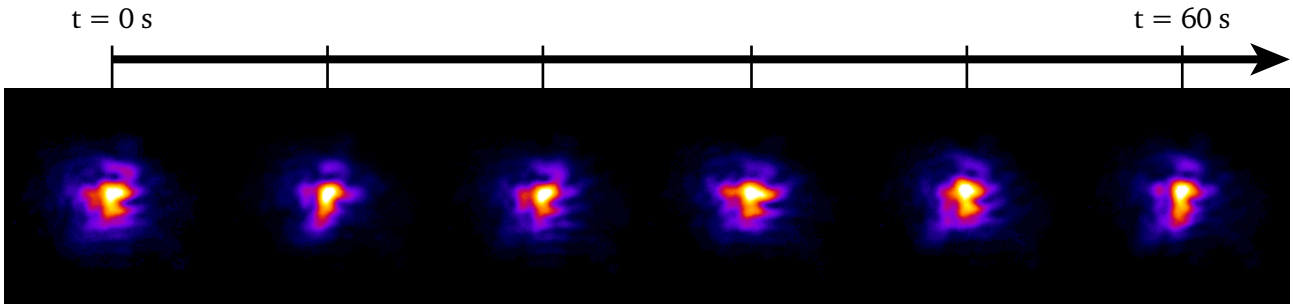

Figure 4.6.: Samples of the off-shot focal spot in the target chamber of PHELIX after a day of five shots, recorded over the course of one minute. The changes in the shape of fluence distribution demonstrate that the dynamic aberrations have a significant influence on the beam quality. At the same time, the shot delay at PHELIX is too long to compensate for these on-shot.

4.4.3. Shot Aberrations

The last challenge which is specific for AO in single-shot laser facilities compared to other applications is the occurrence of shot aberrations, which I already discussed in section 3.4. As the gain media of solid state laser amplifiers are pumped, their refractive index changes, adding a spatially dependent OPD and thus a WF error to the beam, which can reach a considerable amplitude. Countermeasures include moving lenses or bending mirrors, which reduce the requirements on the stroke of DMs.
Shot aberrations are consistent as the pumping geometry of the amplifiers does not change from shot to shot, which means they are theoretically easy to compensate. The challenge that once faces here is that the measurement needs to be precise in order to do so. By nature, shot aberrations can only be recorded on-shot, which means that the issues from attenuation and shot delay, which I described in the last two sections, hinder this measurement as well.
In that sense, shot aberrations are not the primary challenge but act as a second entry point for the shot-related errors generated by attenuation and shot delay.

Conclusion

In this chapter, I reviewed the specific challenges that AO has to face in HI laser systems. This knowledge is necessary to deduct if a certain approach to solve these issues is suitable or not, including the choice of AO technology. For the latter, I recommend reading appendix A, where I introduce a selection of relevant technological innovations that are nowadays commercially available. This greatly helps to understand the decisions I met when designing the post-compressor AO loop at PHELIX, which I will describe in detail in the next chapter.

5. Implementation of Post-Compressor Adaptive Optics at PHELIX

Prior to this work, the Adaptive Optics (AO) system at Petawatt High-Energy Laser for heavy Ion eXperiments (PHELIX) was limited to a pre-compressor loop, built from a bimorph piezoelectric mirror with 31 actuators at the exit of the preamplifier and a Shack-Hartmann Sensor (SHS) in the Main Amplifier Sensor (MAS) before the pulse compressor. This loop is still used to remove static aberrations at the beginning of an experiment day. The reference WaveFront (WF) is flat over the aperture of the SHS, while some low order aberrations may be added manually to account for the aberrations introduced in the diagnostics imaging path.
To that point, optimizing the focal spot meant precompensating aberrations that were introduced after the pulse compressor and thus entering the compressor with a deformed WF, which leads to Spatio Temporal Coupling (STC) [67] (see section 4.1). In order to avoid this, a post-compressor AO loop had to designed and tested, which was the main goal of this work.
In this chapter, I will first derive the requirements to the loop in section 5.1. In section 5.2, I describe the setup of the WF sensing system of the AO loop, called Petawatt Target Area Sensor (PTAS), that I designed and implemented in the course of this work and show results that demonstrate the feasibility of this system. Next, I outline the calibration routine in section 5.3, which we use to calibrate PTAS to the WF in the target chamber. In the end, I describe how I estimate the on-shot intensity from the calibration data and on-shot WF in section 5.4.

5.1. Requirements

Several requirements arise for the construction of a post-compressor loop at PHELIX, some of which are directly determined by the laser system itself. Generally, these look different for the three parts of an AO system: the active component (in this case: a Deformable Mirror (DM)), the control loop and the sensor.

5.1.1. Requirements on the Deformable Mirror

For the DM, the main goals of the desired correction routine determine what kind of technology (see section A.2) is best suited. At PHELIX, the requirements are:

1. **Geometry:** The beam diameter after the pulse compressor at PHELIX can reach up to 28 cm. As the DM shall replace one of the beam transport mirrors, the reflection angle has to be 45°, which means the active area has to be elliptical. On top of that, the actuator layout shall be optimized to generate Zernike polynomials with maximum efficiency.
2. **Spatial resolution:** As the focusing Off-Axis Parabolic mirrors (OAPs) at PHELIX are, in part, milled from copper, these introduce aberrations which may not be described well by low-order Zernike polynomials. As a result, a higher number of actuators is preferable compared to low order DMs as used in the preamplifier.

3. **Speed:** The pulse transport of a compressed laser pulse has to happen in vacuum, which means that the beam will not catch any additional dynamic aberrations. The control speed is therefore not a critical parameter in this case.
4. **Stability:** PHELIX is a Nd:Glass laser which is only able to deliver a shot every 90 minutes. Prior to each shot, the beam is available in a 10 Hz, low energy alignment mode, which serves the AO loop as control reference. The shot procedure involves switching off the light for several minutes while preparing the shot, in which no closed loop beam control can happen. Therefore, The surface of the mirror needs to remain stable during that time, making stability an important DM trait.
5. **Stroke:** In the target chamber, the focusing OAP can be used to compensate for Tip, Tilt, Astigmatism and Defocusing during alignment. These are the lowest order aberrations and thus usually the ones with the largest amplitude. Neglecting these, experience tells that any technology with an actuator stroke of 10 λwill be sufficient.
6. **Damage Threshold:** The mirror surface has to be able to withstand the intensities of the compressed pulse in the Near Field (NF). For a shot of 200 J and a circular profile of 28 cm, this corresponds to 325 mJ/cm^2 within 500 fs. Taking into account local intensity fluctuations of up to 100 % and a safety margin of 50 %, the damage threshold has to be above 1 J/cm^2 in normal direction, which rules out DMs with metallic coating or ultrathin membranes.
7. **Material:** The DM has to be built from materials that are vacuum compatible.
8. **Interface:** At PHELIX, a custom open source control software called Wavefront Optics Measurement and Beam Analysis Tool (WOMBAT) [68] is used. Compared to commercial software, this guarantees maximum flexibility and independence when problems are encountered. Therefore, the mirror has to have an interface that can be adressed without proprietary software.

These requirements lead to the decision to use a *MD-AME* DM by ISP system (see Fig. 5.1).
The purchased model *MD-280-E-53-HV-TipTilt* features 53 stepper motor actuators which are arranged for optimum Zernike polynomial compensation and can be operated up to 10 Hz. Once the control loop is switched off, they passively hold their position indefinitely, fulfilling the requirement for stability.
The surface of the DM is a glass membrane with suitable diameter and vacuum compatible High Reflectivity (HR) coating for the specified WF of 1053 nm at an Angle Of Incidence (AOI) of 45°.
The interfacing with the controller happens via a dynamically liked library, provided by ISP system, which can be added to custom software and opens a TCP connection to the device.
To aid the alignment of the DM, it came with a motorized Tip/Tilt mount, which can be controlled separately from the mirror actuators. This also provides another option for aligning the OAP in the target chamber.

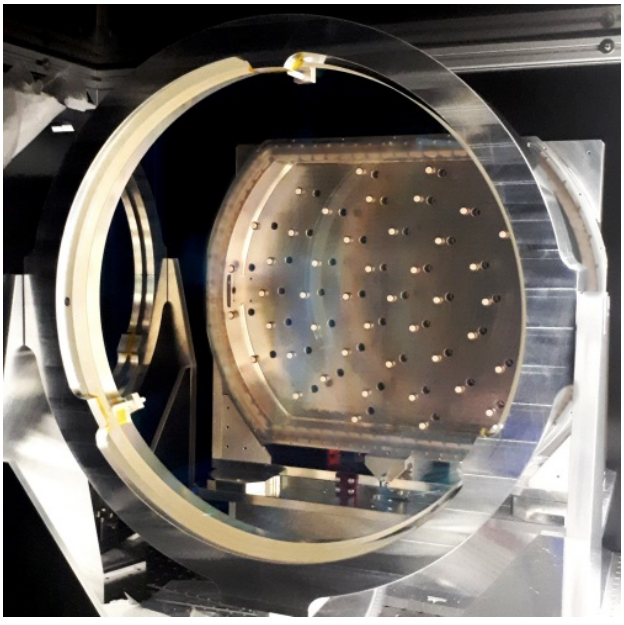

Figure 5.1.: Photo of the *MD-280-E-53-HV-TipTilt* DM during assembly at ISP system. The photo is kindly provided by the manufacturer.

5.1.2. Requirements on the Control Loop

As the speed requirement on the control loop is loose, a CPU based control system (i.e. a regular software running on the root complex of a computer) can be used in this application. While commercial solutions exist from most DM manufacturers, a custom software is already in use at PHELIX: the open source software WOMBAT [68] (see Fig. 5.2), written in NI LabVIEW (National Instruments, Texas).
WOMBAT uses asynchroneous data transfer between modules to guarantee smooth execution.

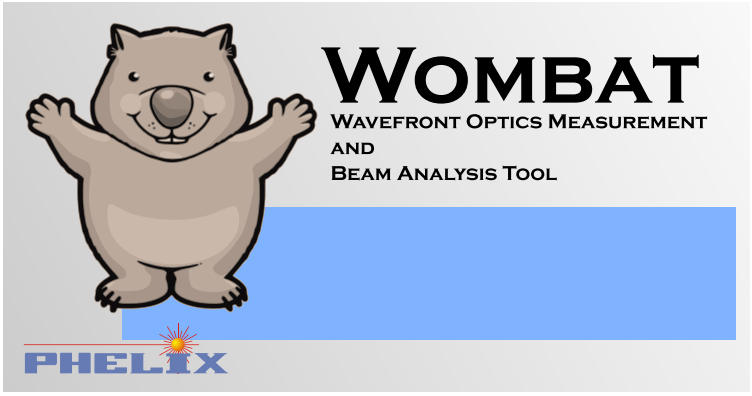

Figure 5.2.: The logo of WOMBAT

Using WOMBAT has several advantages. First, the new loop will be compatible with the rest of the facility, second, the operating crew is already familiar with the user interface and third, as the software is written in house, problems and feature requests can be handled in a flexible manner, which means that the implementation is independent from proprietary software.
All the required functionality is already implemented: in WOMBAT, cameras using the GigE interface can be read out and SHS images can be evaluated on the fly. Also, a general communication interface for DMs exists, which is used to record mirror responses and run a closed loop to match arbitrary target WFs, defined by an unlimited number of previously saved WF files and/or custom Zernike polynomials. The main disadvantage is that the libraries for the communication with the DM, provided by ISP system, cannot be used in LabVIEW. Therefore, I implemented a local server application in C++, which runs the libraries and takes TCP connections from any software application. Like this, LabVIEW can send commands to the server, which itself communicates with the DM (see section 5.2.6 for details and a schematic).

5.1.3. Requirements on the Sensor

The last component of the AO system is the sensor that shall be used to run the loop and estimate the intensity on the target. Here, the requirements are task- and facility oriented:

1. **Beam demagnification:** As for the DM, the full aperture of the incoming beam can reach up to 280 mm in diameter, while the aperture of the SHSs at PHELIX only measures 5 mm. Therefore, the demagnification factor has to be at least 56, where slightly larger values are acceptable. Values larger than 70 will reduce the beam size too much for the WF measurement.

2. **Imaging of the DM:** For a good stability, the WaveFront Sensor (WFS) has to be placed conjugate to the DM or, in other words, sit in an image plane of the DM surface. The sensor design thus has to take imaging into account.

3. **High dynamic range:** As discussed above, PHELIX is only able to shoot every 90 minutes. For this reason, the control loop has to pre-compensate the shot-aberrations using the low energy alignment beam. At the same time, the need to measure the on-shot WF in order to estimate the intensity on target, which means PTAS has to cover a dynamic range of about 5×10^5 (see section 5.2.5) and everything in between continuously.

4. **Appropriate transmitted WF:** As the beam has to be demagnified in order to fit onto the WFS, PTAS inherently cannot provide an absolute measurement of the WF due to aberrations introduced in the imaging system. However, a good accuracy is still desirable to allow a more intuitive estimation of the current performance of the laser system. Given that a Strehl ratio of 0.9 is acceptable, this corresponds to an approximate Root Mean Square (RMS) WF error of 52 mλ [69] (for pure astigmatism in the Zernike sense, this equals 286 mλ Peak-to-Valley (PtV)).

5. **Calibration routine:** Due to the fact that this sensor does not provide absolute WF measurement, a suitable calibration routine relative to the WF in the target chamber has to be implemented.

6. **Compact design:** At PHELIX, the pulse compressor is located only a couple of meters away from the target chamber, leaving little space for the sensor. Additionally, the target area is densely packed with equipment, local storage and radiation shielding (see Fig. 5.3), leaving an area of only 2 m^2. Therefore, the complete sensor system, including the demagnification of the beam, has to follow a very compact design.

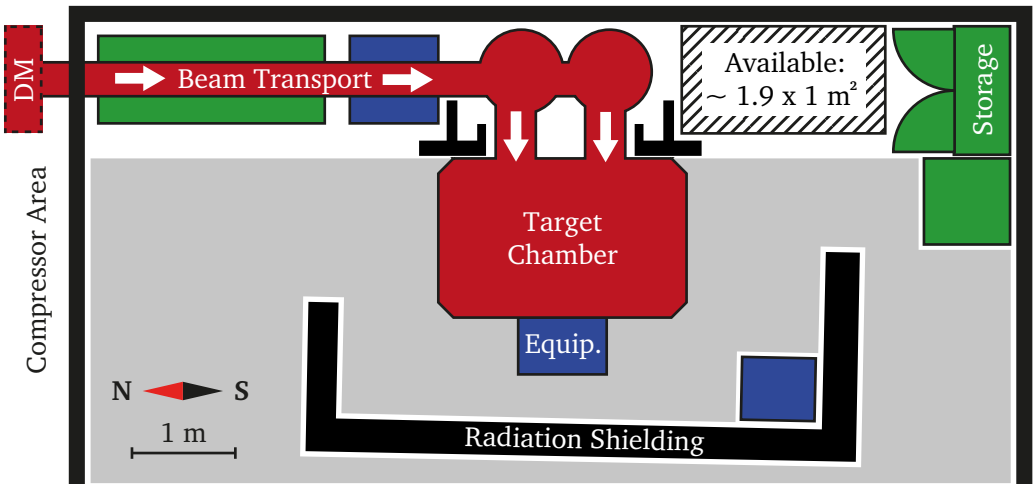

Figure 5.3.: Sketch of the PHELIX target area, true to scale, before the installation of PTAS. The DM will be placed at the compressor exit to the left - the position is indicated by a red box with dashed contour. As the light gray area is reserved for user equipment and personnel access, the available space left for PTAS is on the top right, indicated by the ruled area.

5.2. Design of the Petawatt Target Area Sensor

We published a paper about the design of PTAS and first experimental results before [2]. Here, I will describe the system in more detail, elaborating on considerations and decisions regarding the design. First, I will give a brief overview of PTAS, then I go into more detail on the imaging aspect, the chosen alignment strategy and the energy attenuation, respectively. In the end, I will give some results on the performance of PTAS.

5.2.1. Overview

A schematic of PTAS is shown in Fig. 5.4 in the same orientation as Fig. 5.3, also true to scale. The light leaked through the last turning mirror before the target chamber enters a vacuum vessel from the right, still featuring a beam diameter of 280 mm. In order to be compact, the main demagnification is done using a 8:1 OAP telescope, described in section 5.2.2. The alignment of such a short telescope is obviously critical, therefore we use an intermediate focal diagnostic (not shown in Fig. 5.4) for fast and repeatable alignment. I describe this technique more closely in section 5.2.4. Before leaving the vacuum, a Mach-Zehnder style attenuator, which I describe in section 5.2.5, optionally attenuates the beam by two orders of magnitude.
The rest of the setup follows a more conventional approach: the beam passes through two motorized, reflective neutral density filters for further attenuation and gets demagnified by a factor of 3.3 afterwards, using a regular Kepler telescope, composed of achromats (L1 and L2 in Fig. 5.4). After that, another continuous attenuator using a $\lambda/2$ waveplate and a single polarizer can be used for fine adjutement of the transmitted energy. Finally, another achromatic Kepler telescope (L3 and L4 in Fig. 5.4) demagnifies the beam by another factor of 2.4 and images the DM onto the entrance of a camera box. This is a closed device, developed at PHELIX, which contains a camera for the NF and Far Field (FF) of the beam respectively, as well as a custom SHS for WF measurements.

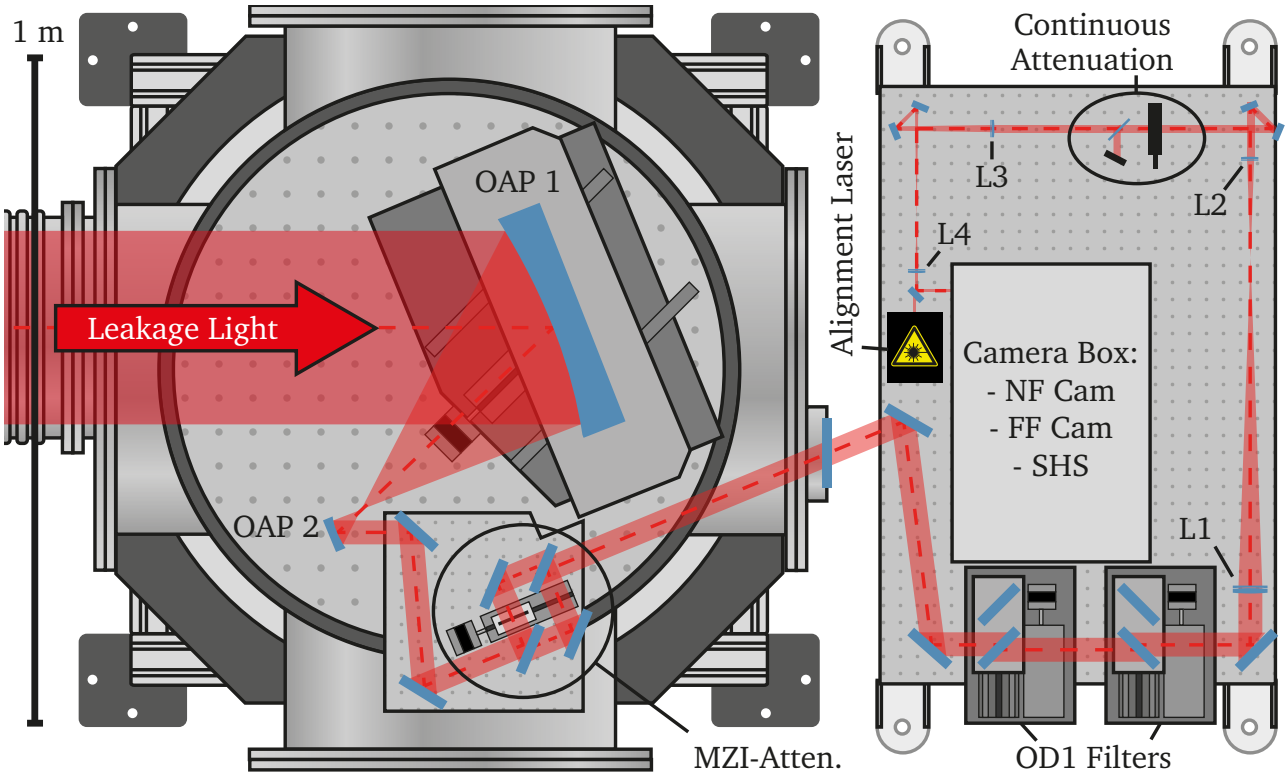

Figure 5.4.: Schematic of PTAS, true to scale. The intermediate focal diagnostic, described in section 5.2.4, is not shown for clarity. The other main components are labeled ("L" for lenses) and described in the text.

5.2.2. Imaging System

5.2.2.1. Reflective Imaging

At PTAS, several requirements are set to the imaging and demagnification setup: it has to be compact, introduce little aberrations, be achromatic and resistant to high intensities. The last requirement even holds true while only handling 1 % of the main beam energy as a demagnification of factor ten recreates the intensity of the main beam.
All these requirements lead to the conclusion that only reflective optics can be used in the first stage of the imaging system:

- The system has to be compact, which means that the F-number of the optics will be small. Generally, large angles are harder to deal with using refractive optics.
- The need for low aberrations require both reflective and refractive optics to be aspherical, where the costs greatly increase with the aperture size. Refractive optics, however, require more than one polished surface.
- As PHELIX is a pulsed laser system, the light is not monochromatic, creating the need for achromatic imaging. Again, the large angles make the manufacturing of achromatic lenses a lot harder and more expensive, while reflective optics are inherently achromatic.
- The high intensities that are expected here pose the biggest challenge for refractive optics, as the damage thresholds are lower than for reflective optics and nonlinear effects will happen inside the material on shot.

These points clearly show the need to use reflective optics for imaging after the pulse compressor of Chirped Pulse Amplification (CPA) lasers. In this case, the small F-numbers additionally enforce this requirement. The obvious solution for the imaging at PTAS is to use an OAP telescope. The main challenge is the sensitivity to alignment. With the right setup, however, it can be done in a fast and repeatable manner, which I discuss in section 5.2.4.

5.2.2.2. Imaging Design

The standard focusing optic at PHELIX is an OAP with an off-axis angle of 45° and an effective focal length of $f_1 = 400\,\text{mm}$, milled from copper or ground into a glass substrate. A telescope of this F-number would be very compact, as demanded for PTAS, while being compliant with the PHELIX standards at the same time.
For the second OAP as the counterpart in the telescope, the two main design constraints were the damage threshold and the producibility. As stated above, demagnifying the beam by a factor of 10 would recreate the main beam intensity, which should be the maximum in order to avoid laser induced surface damages. At the same time, such an OAP is very hard to manufacture due to the strong concave curvature of the substrate.
After a comprehensive survey, I found that an OAP with $f_2 = 50\,\text{mm}$ on a round 2” substrate posed a good tradeoff between size and accuracy. Our OAP was milled from copper by LT Ultra GmbH, Germany, featuring a surface error of less than 150 nm PtV and a roughness of less than 5 nm RMS. These imperfections are still relatively large and I will show that the alignment errors are therefore not the limiting factor for the transmitted beam quality of PTAS in section 5.2.4.
After this first demagnification of about a factor of eight, the beam diameter is only 35 mm. Thanks to the attenuation following the OAP telescope, the beam can now leave the vacuum vessel through a wedged window flange and be further imaged using off-the-shelf achromats on a breadboard. Due to the small beam diameter, this will not stress the space constraints anymore.

The second telescope is composed of one f = 500 mm and one f = 150 mm achromatic doublet (labeled ”L1” and ”L2” in Fig. 5.4) and the third one brings the beam diameter down to 4.4 mm using one f = 300 mm and one f = 125 mm achromatic doublet (labeled ”L3” and ”L4” in Fig. 5.4, respectively), imaging the surface of the DM onto the image plane of the camera box. The lenses are oriented such that the transmitted WF is optimal, while reducing the spatial resolution of the imaging.
The combination of all telescope reduce the beam size by a factor of 64.

5.2.3. Sensitivity to Misalignment

Despite the favorable properties of OAP telescopes as described in section 5.2.2, which ultimately lead to the decision of employing this technology in PTAS, these systems have the reputation to be notoriously hard to align with corresponding procedures being viewed as overly complicated. This statement holds true if interferometric alignment precision is needed [70].
Such precision is not a requirement for PTAS. Therefore, a different and simpler alignment procedure may be employed, as long as the expected precision is reasonable. In order to do quantified estimations, I did an extensive analytical consideration of the topic using a ray tracing approach before and developed an alignment scheme that is easy, repeatable and sufficiently exact for the application of OAP telescopes in High Intensity (HI) laser systems [1].
In this section, I reiterate the main findings about alignment errors in OAP telescopes, before I explain the alignment strategy in section 5.2.4 and estimate the expected WF errors of the PTAS telescope.

5.2.3.1. Description of an OAP Telescope

In order to determine the WF errors that are generated by the misalignment of an OAP telescope, I first have to describe the OAPs in a parametric manner.
Most commonly, OAPs are characterized by their properties in terms of the propagation of a chief ray, i.e. the so-called off-axis distance d_{oa}, the off-axis angle α_{oa} and the effective focal length f_{eff} (see Fig. 5.5).

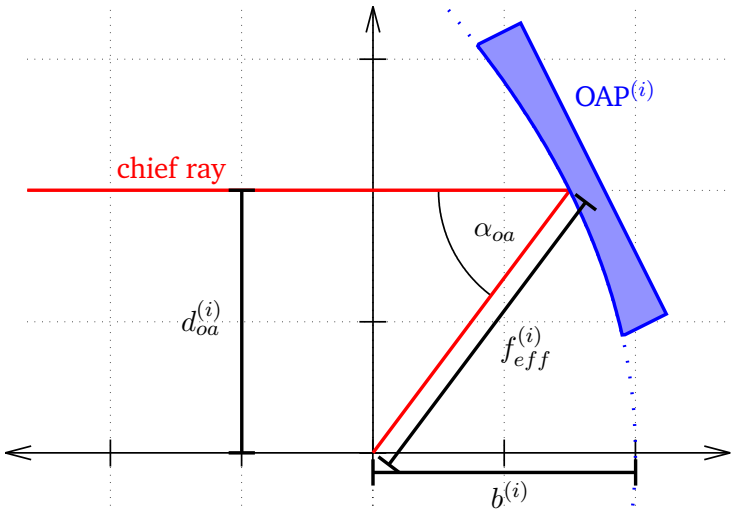

$f_{eff}^{(i)}$	Effective focal length
$d_{oa}^{(i)}$	Off-axis distance
α_{oa}	Off-axis angle
$b^{(i)}$	Parent focal length

Figure 5.5.: Sketch of an OAP with according parameters, where the superscript $^{(i)}$ will be used to differentiate between the first and second OAP of a telescope.

From these numbers, one can obtain a parametric representation of a parabolic surface $\vec{O}^{(i)}$ with the optical axis along the z axis and the focal point in the origin:

$$\vec{O}^{(i)}(u,v) = \begin{pmatrix} u \\ v \\ a^{(i)}(u^2+v^2)+b^{(i)} \end{pmatrix} \quad \text{with} \quad \begin{aligned} a^{(i)} &= -\frac{\tan(\alpha_{oa}/2)}{2d_{oa}^{(i)}} \\ b^{(i)} &= -\frac{1}{4a^{(i)}} \\ d_{oa}^{(i)} &= f_{eff}^{(i)} \cdot \sin(\alpha_{oa}). \end{aligned} \tag{5.1}$$

Here, the superscript $^{(i)}$ will indicate if the first or second OAP of the telescope is meant. Note that the off-axis angle α_{oa} has to be the same for both OAPs, as we will see in the next section, and thus does not feature this denotation.
Using this parametrization, one can analytically calculate the optical path of an incident ray, starting at a point $\vec{d}^{(0)}$, parallel to the optical axis along the propagation vector $\vec{k}$, through the telescope. The calculation consists of the following steps:

1. Calculation of the intersection point $\vec{d}^{(1)}$ between the incident ray and the first OAP
2. Calculation of the new direction of propagation $\vec{k}'$ after the first reflection by applying the Householder transformation [71] using the local normal vector of $\vec{O}^{(1)}$ at $\vec{d}^{(1)}$
3. Calculation of the intersection point $\vec{d}^{(2)}$ between the reflected ray and the second OAP
4. Calculation of the new direction of propagation $\vec{k}''$ after the second reflection analogous to the first reflection

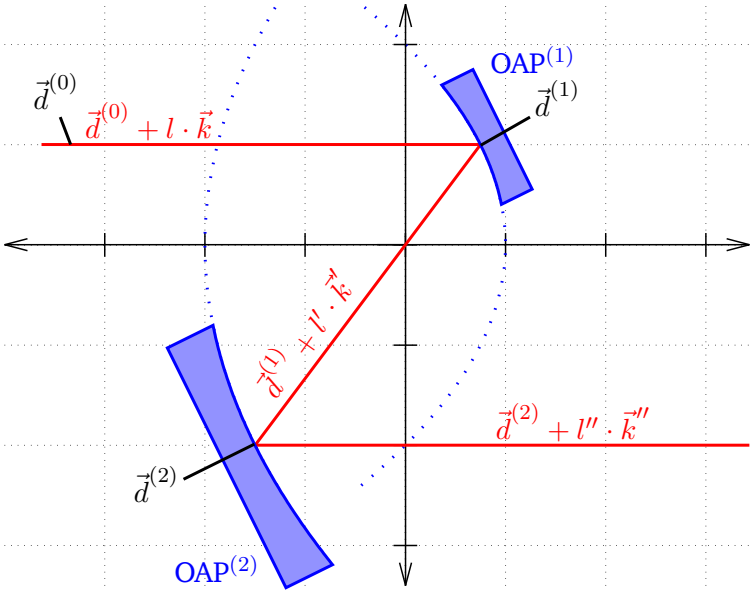

$\vec{d}^{(0)} + l \cdot \vec{k}$	Incoming ray
$\vec{d}^{(1)}$	First intersection
$d^{(1)}$	$\left\|\vec{d}^{(1)}\right\|$
$\vec{d}^{(1)} + l' \cdot \vec{k}'$	Intermediate ray
$\vec{d}^{(2)}$	Second intersection
$d^{(2)}$	$\left\|\vec{d}^{(2)}\right\|$
$\vec{d}^{(2)} + l'' \cdot \vec{k}''$	Outgoing ray
$\beta = -\frac{d^{(2)}}{d^{(1)}}$	Neg. magnification

Figure 5.6.: Sketch of an OAP telescope with a ray passing through. The corresponding parametrization is labeled and explained on the right.

From this point, one can introduce misalignment into the system and investigate its effects onto the outgoing ray. In the following section, I do so for translational misalignment, followed by rotational misalignment.

5.2.3.2. Translational Misalignment

Deviating from the ideally aligned OAP telescope, described in the previous section, any misalignment can be described by translating and rotating the second OAP, given that the incident beam propagates parallel to the optical axis of the first OAP. The simplest misalignment would be a translation of the second OAP by a translation vector $\vec{T}^{(2)}$ (see Fig. 5.7).

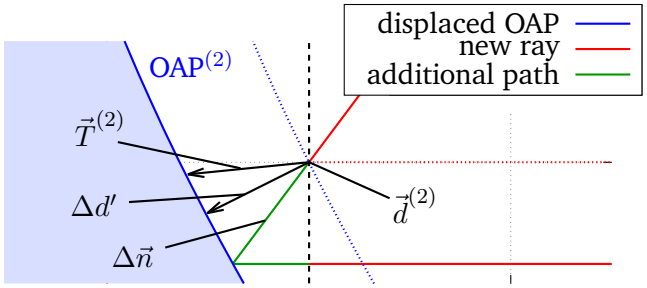

$\vec{T}^{(2)}$ Translation vector
$\Delta\vec{n}$ Surface normal translation
$\Delta\vec{d'}$ OPD between intersections

Figure 5.7.: Sketch of the displacement of the OAP at the region around the second intersection point $\vec{d}^{(2)}$

As long as the translation is small compared to the local curvature and not parallel to the surface at the original intersection point $\vec{d}^{(2)}$ of a ray, an optical path difference P (shown in green in Fig. 5.7) can be expressed in terms of a linear expression of $\vec{T}$:

$$P = \frac{\Delta\vec{d'} \cdot \vec{k'}}{\left\|\vec{k'}\right\|} - \Delta\vec{d'} \cdot \vec{e}_z = -\frac{1}{d^{(1)}} \begin{pmatrix} x \\ y \\ 2b^{(1)} \end{pmatrix} \cdot \vec{T}^{(2)} \tag{5.2}$$

with x and y being the lateral coordinates of incoming ray, i.e. the x and y coordinates of $\vec{d}^{(0)}$, $d^{(1)} = \|\vec{d}^{(1)}\| = -a^{(1)}r^2 + b^{(1)}$ and $r^2 = x^2 + y^2$.

Obviously, this expression only takes into account the linear evolution of the path difference. A detailed derivation of equation (5.2), together with a quantitative consideration on why higher order terms can be neglected in this user case, can be found in my corresponding publication [1].

Note that P does only depend on the characteristics describing the first OAP, which can be linked to an F-number via the extent of x and y and the effective focal length, but not on the magnification β. This implies that the WF error is the same for a fixed translation independently from the beam size and the magnification, requiring the same absolute alignment precision for a millimeter scale beam as for a beam with several tens of centimeters in diameter for OAP telescopes with the same F-number.

In order to gain some intuitive understanding of the OPD P introduced by translational misalignment, visualizing equation (5.2) is useful. Fig. 5.8 shows the OPD for the x, y and z components of the translation over a region of $1.5b^{(1)}$ around the origin. This has the advantage of eliminating the scale of the system and making the plot applicable for arbitrary sizes. The OPD itself (color coded) is given in relative units to the displacement and thus without unit. The actual OPD can be obtained by multiplying the relative values with the value of the corresponding translation.

As we can see, the spatial distribution of the OPD shown in Fig. 5.8 cannot directly be described by single polynomials in the Zernike or Seidel sense. This also holds true for the defocusing that can be observed for translation in z direction (right plot), as the OPD grows rather linearly with the distance to the center than quadratic.

However, Fig. 5.8 shows the OPD around the optical axis of the telescope. In a real telescope, however, the origin cannot be contained in the beam as the corresponding ray would be obscured by the second

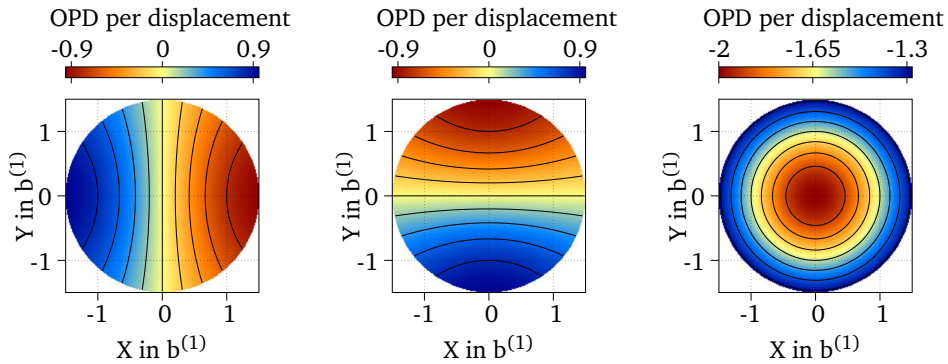

Figure 5.8.: OPD given by equation (5.2), plotted over the coordinates of the input rays in units of $b^{(1)}$, relative to the translation of the second parabola in x-, y- and z-direction (left to right). Note that the origin does not correspond to any beam center but to the symmetry axis of the undisturbed parabolic surfaces. Any beam of a real telescope cannot include the origin.

OAP before hitting the first one. For this reason, a similar set of plots is shown in Fig. 5.9, containing a circular subset of the plots shown Fig. 5.8, corresponding to a telescope with $\alpha_{oa} = 45°$ and F = 3. In this case, the translation is not split into x, y and z components, but into horizontal and vertical transversal and a longitudinal component, matching the chief ray between the two OAPs. For better comprehension, the piston and slope at the position of the chief ray are subtracted from the plots.

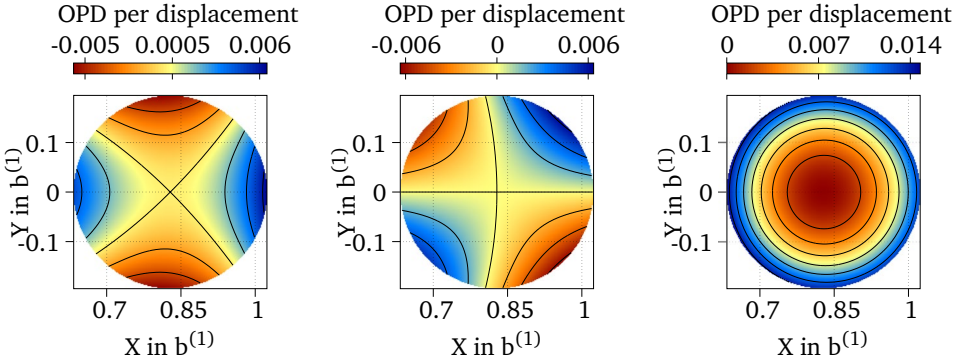

Figure 5.9.: OPD, given by equation (5.2), of a 45° telescope with an F-number of 3, plotted over the coordinates of the input rays in units of $b^{(1)}$, relative to the translation of the second parabola in horizontal, vertical and longitudinal direction (left to right). The piston and slope of the central ray were subtracted for better comprehension.

Viewing the OPD in this way, the resulting aberrations can be described more accurately in terms of Zernike polynomials, i.e. the WF errors for transversal translation are a compound of mainly astigmatism and coma (left and center), while longitudinal translation primarily features defocus. We also get an estimation for the scaling of the aberration amplitude relative to the displacement of the second OAP, which is between one and two orders of magnitude below unity. This implies that the translation has to be controlled in orders of WaveLengths (WLs) if one wants to keep the amplitude at a fraction of a WL.

5.2.3.3. Rotational Misalignment

Any misalignment of an OAP telescope can be decomposed into a translational misalignment, as described in the last section, and a rotation of one OAP around an arbitrary center. It therefore makes sense to choose this center appropriately to the subject. While choosing the intersection point of the chief ray may seem intuitive due to the mechanical mounting of the OAP, choosing the focal spot makes more sense in this scenario for two reasons:

First, one can clearly distinguish between the effects on the beam caused by rotation and translation in this case, as they are mathematically inherently different. I will elaborate on this below.
Second, choosing the center as anchor for the rotation matches my proposed alignment scheme very well, as it is based on spatially overlapping the focal spots of the two OAPs. This automatically minimizes the translation separated from the rotation around the focal spot. I will describe this approach in detail in section 5.2.4.
In the following, I will briefly report on my findings of the effects of rotational misalignment. Given that the in- and outgoing beams are well aligned to their corresponding OAP, which is generally the case when using my proposed alignment scheme (see section 5.2.4), either of the two OAPs can be rotated around an axis through the origin in the xy-plane, as rotations around the z-axis will not influence the beam due to the symmetry of the setup. I will use the angle ψ to describe the orientation of this axis, where $\psi = 0$ corresponds to the x-axis in positive direction and $\psi = \pi/2$ to the y-axis in negative direction.
In my considerations, I rotated the first OAP, which is shown schematically in Fig. 5.10. Here, the first OAP is rotated about a positive angle ω, while the rotation axis points away from the viewer.

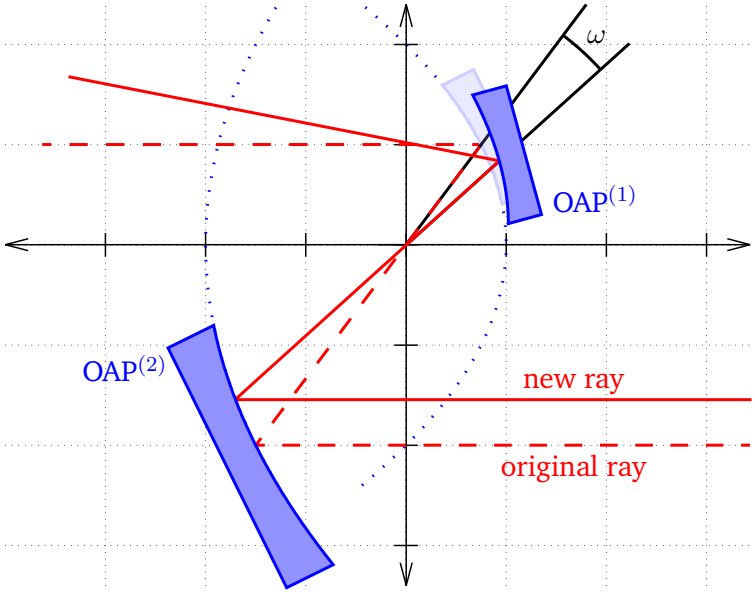

Figure 5.10.: Sketch of an OAP telescope with rotational misalignment of positive ω, where the rotation axis points into the plotting plane.

Due to the rotation around the origin and the given alignment of the first beam to the first OAP, all rays will pass through the origin. Therefore, outgoing rays will be displaced, but not deflected relative to the optical axis of the second OAP. This ultimately means that the rotational misalignment does not generate any OPD between the rays of the beam, but the NF becomes displaced and distorted instead. I derived an expression for this beam distortion using ray tracing with linear approximations [1]:

$$x_{out} \approx \beta_R(\psi)(x - \omega \sin\psi O_z^{(1)}), \tag{5.3}$$

$$y_{out} \approx \beta_R(\psi)(y - \omega \cos\psi O_z^{(1)}), \tag{5.4}$$

$$\text{with} \quad \beta_R(x, y, \psi) \approx \beta \frac{\omega(x \sin\psi + y \cos\psi) + 2a^{(1)} r^2}{2a^{(1)} \left(r^2 - 2(x \sin\psi + y \cos\psi)\omega O_z^{(1)}\right)}.$$

Here, $O_z^{(1)}$ corresponds to the z component of the parametrization of the first OAP, as given in equation (5.1), and ω has to be given in radians.
This approximation yields errors of less than 1 % of the displacement for the region of $0.2b^{(1)} \leq r_{in} \leq 3.7b^{(1)}$ for $\omega \leq 5°$, which covers the area of practically usable OAP telescopes and is shown in Fig. 5.11 on the left in bold lines.
Obviously, the distortion of the outgoing beam leads to areas of decreased intensity where the ray density is diluted and areas of boosted intensity where the rays contract. And while the distortion of the NF itself may not be an issue, depending on the application, locally increased intensity can pose a significant threat in HI lasers as damage thresholds my be reached faster than expected.
An analytical expression for the local intensity boost can be obtained using the derivatives of equations (5.3) and (5.4). However, the closed expression is rather complicated and therefore impractical to use, while the errors of the linear approximations get amplified in the process. I found that one can obtain a much simpler and more precise approximation of the local intensity using an ordinary quadratic function over $r = x\sin(\psi) + y\cos(\psi)$. A fit to numerical values yields

$$I(x, y, \psi) \approx \frac{ar^2 + br + 1}{\beta^2} \qquad \text{with} \quad a = \frac{\sqrt[4]{2}}{\sqrt{10}b^{(1)2}}\omega^2 \approx 0.376\frac{\omega^2}{b^{(1)2}}, \qquad b = \frac{\omega}{b^{(1)}}, \tag{5.5}$$

where ω is given in radians.
An example for the numerically calculated intensity boost is shown in Fig. 5.11 in the center. On the right, the relative error introduced by equation (5.5) is shown. Here, one can see that the error made by equation (5.5) is below 1 % for a region of $5b^{(1)}$ around the optical axis of the telescope, as long as the rotation ω is less than 5°.

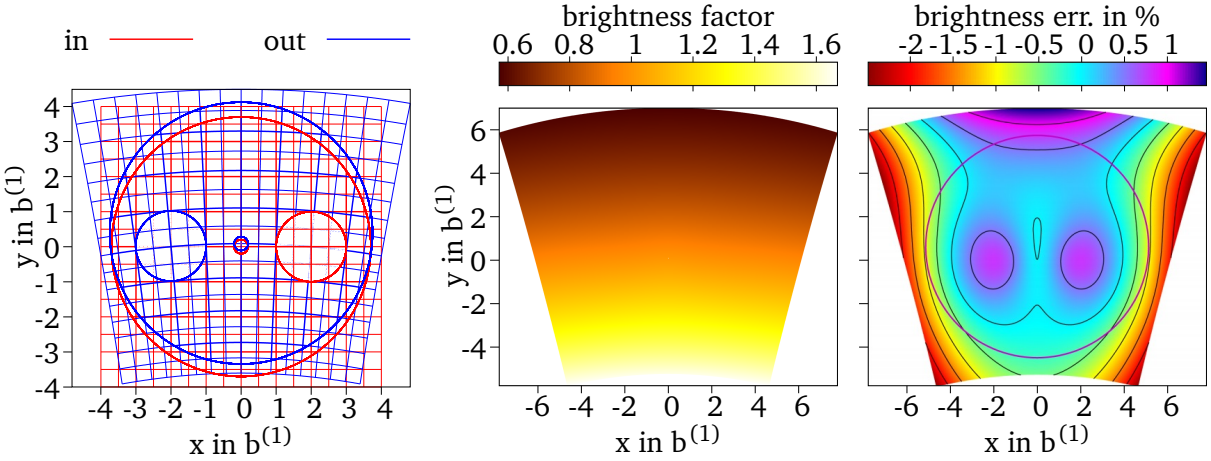

Figure 5.11.: Beam distortion example for a telescope with $\beta = -1$, $\omega = 5°$ and $\psi = 0°$.
Left: input coordinates (red) are translated to output coordinates (blue) for a $8b^{(1)} \times 8b^{(1)}$ region. The big and small circles indicates a region for $0.2b^{(1)} \leq r_{in} \leq 3.7b^{(1)}$, while the two medium circles mark the relevant area for a 90°, F=1 telescope.
Center: numerically calculated brightness factor at the telescope output for a $12b^{(1)} \times 12b^{(1)}$ region, taking into account the beam distortion.
Right: corresponding deviation of equation (5.5) in percent with contour lines at steps of 0.5 %, excluding 0 %. The purple line indicates the region around the origin of $r \leq 5b^{(1)}$ where the error is less than 1 %.

Assuming that a local intensity boost of 5 % is acceptable, as this is typically comparable to other inhomogeneities in large aperture beams, one can now estimate the precision that has to be reached in alignment in order to fulfill this criterion. To do so, one can insert the maximum value for r in equation (5.5) that is present in an OAP telescope. For a 90° telescope with F=1, this is $r_{max} = 3b^{(1)}$. In this case, equation (5.5) predicts an intensity boost of 5 % at $\omega \approx 0.94°$. This implies that the rotation generally has to be done with a precision of about one degree, which can easily be done while aligning the system geometrically and therefore does not require any specialized alignment procedure.

5.2.4. Alignment Strategy

As discussed in the last two sections, aligning the rotation of an OAP telescope generally is not an issue as the required precision is only in the order of one degree. Therefore, this can be ensured during the setup phase of the system using geometric approaches. The translational alignment on the other hand has to be treated more carefully, as the required precision is in the order of WLs.
A conceptually simple way to ensure small translational alignment errors is to spatially overlap the focal spots of both OAPs using an intermediate focal diagnostic. Aside from the main laser beam passing through the telescope which can be used to align the first OAP, this also requires a back-propagating laser beam for the second OAP. Furthermore, a fiducial in the focal plane has to precisely mark a fixed position in space at which the focal spots have to be located.
In this section, I will first estimate the alignment precision that can be reached, as this clearly indicates if this approach is suitable or not. We will find that the precision is easily sufficient for most applications in the context of HI beam transport, including PTAS.
After that, I will describe the technical implementation of this concept at PTAS.

5.2.4.1. Theoretical Considerations

In order to estimate the WF quality that can be reached by overlapping the focal spots, we need two things: first, a quantified measurement of the overall WF quality, depending on the misalignment, and second, an estimation of the expected maximum residual misalignment.

Wavefront quality over misalignment A widespread measurement for the quality of the WF is the WF RMS, which is the mean quadratic deviation of the WF from a reference surface. While different definitions may be used for the latter, depending on the application, we will use a plane that is best fitted to the WF, i.e. subtract piston and tilt before calculation the RMS with respect to zero. This is useful as neither piston, nor tilt change the intensity distribution in the focal plane aside from a lateral offset.
In Fig. 5.12, the WF RMS per translation is shown for horizontal, vertical and longitudinal misalignment of an OAP telescope, plotted over the F-number and the off-axis angle and therefore covering most conceivable OAP telescopes. Shown like this, three things become apparent: first, the RMS for horizontal and vertical translation look nearly identical, which is why I will treat the two as unified transversal misalignment from now on. Second, the RMS decreases with increasing F-number in all cases, which is intuitive as decreasing the F-number equals decreasing the aperture size of the system and therefore "cutting out" a smaller portion of the deformed WF. And third, the RMS generally increases with increasing off-axis angle. However, this is much more pronounced for transversal misalignment.
I calculated the values shown in Fig. 5.12 numerically using equation (5.2) as doing so analytically would only yield overly complicated expressions. The goal is now to approximate these values using an arbitrary analytical expression in order to deliver a closed-form estimation for the RMS caused by

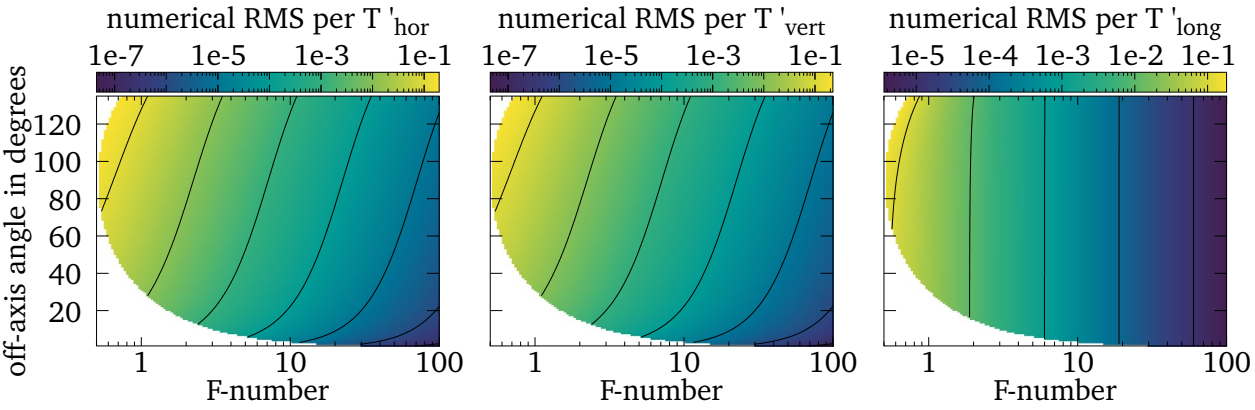

Figure 5.12.: The WF RMS relative to the translation amplitude of the second OAP, plotted over the setup parameters F and α_{oa}. The contour lines indicate orders of magnitude. The translation is not given in x, y and z direction but in horizontal, vertical and longitudinal translation relative to the chief ray between the OAPs.
Note that the RMS-per-displacement axis is again dimensionless as it has to be multiplied with the corresponding translation in units of length to reflect an actual RMS value.

residual misalignments. For the transversal misalignment, one can use

$$\delta_{\text{trans}}(F,\alpha) \approx 3.2 \times 10^{-3} \cdot \left(F^{-1.9} \cdot 10^{0.58\cdot\alpha} \cdot \arctan(6\alpha) + \frac{17 \cdot 2^{-2.5\cdot F}}{1+(\alpha-1.3)^2} \right), \tag{5.6}$$

while the RMS caused by rotational misalignment can be approximated by

$$\delta_{\text{long}}(F,\alpha) \approx \frac{35 \times 10^{-3}}{(F - 9 \times 10^{-3} \cdot 10^{\alpha/2})^2 + 65 \times 10^{-3}}. \tag{5.7}$$

One should keep in mind that these expressions are arbitrary, i.e. not motivated by a physical model, and only serve the purpose to provide a practical estimation. The error of these expressions with respect to the numerical values shown in Fig. 5.12 are shown in Fig. 5.13.
As we can see, the deviations raise sharply for small F-numbers and large off-axis angles. This region, however, covers OAP telescopes that are hardly practical to use as enormously large OAP with a very shallow AOI are needed for relatively small beam diameters. Therefore it is reasonable to exclude this area from the considerations and introduce a border of validity, over which the absolute error of the approximations may exceed 2.5e-3:

$$\alpha \leq 2.5 \log_{10}(F) + 1.9. \tag{5.8}$$

This border is indicated in both graphs of Fig. 5.13.

Maximum residual misalignment In the next step, we need to estimate the precision in translation that can be reached using only the intermediate focal spots as reference. Having these numbers, we can use equation (5.6) and (5.7) to estimate a corresponding WF RMS.
The precision for transversal misalignment is naturally limited by the size of the focal spot between the OAPs, which can be compared to the spatial resolution of an imaging system. For a homogeneously illuminated aperture, this corresponds to the diameter of the Airy disc, which is given by $D_f \approx 0.61\,\lambda/\text{NA}$ where λ is to the central laser WL and the Numerical Aperture (NA) of the focusing optic [72]. Naturally,

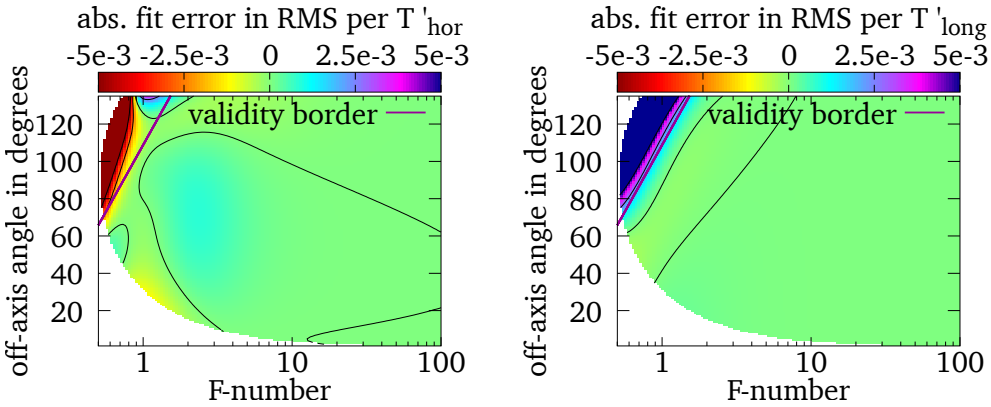

Figure 5.13.: Left: the absolute fit error of transversal translation, described by equation (5.6). Right: the absolute fit error of horizontal translation, described by equation (5.7).
The black contour lines indicated steps in 2.5e-3, starting at 0. The purple line indicates the threshold according to equation (5.8), above which the fit error may exceed 2.5×10^{-3}.

this only holds true for a beam with small aberrations. A suitable criterion for this is the one given by Marechal [73], where the WF RMS of the incoming beam has to be lower than $\lambda/14$. If this is not the case, D_f has to be increased accordingly.
Using this expression, we can derive an expression for the best guaranteed RMS Δ_{trans} that is caused by translational misalignment:

$$\begin{aligned} D_f(\lambda, F) &\approx \frac{0.61\,\lambda}{\text{NA}(F)} \qquad \text{with} \qquad \text{NA}(F) = \text{arccot}(2F), \\ \Rightarrow \Delta_{\text{trans}}(F, \alpha, \lambda) &\approx D_f(\lambda, F) \cdot \delta(F, \alpha) \\ &\approx \frac{2 \times 10^{-3}\lambda}{\text{arccot}(2F)} \cdot \left(F^{-1.9} \cdot 10^{0.58\cdot\alpha} \cdot \arctan(6\alpha) + \frac{17 \cdot 2^{-2.5\cdot F}}{1 + (\alpha - 1.3)^2} \right). \end{aligned} \tag{5.9}$$

Here, α has to be given in radians. Δ_{trans} is shown in Fig. 5.14 on the left for the allowed parameter range.
The same can now be done for longitudinal misalignment. In this case, the precision is not limited by the focal spot size, but by the Rayleigh-range z_R of the beam:

$$\begin{aligned} z_R(\lambda, F) &= \frac{\lambda}{\pi\ \text{arccot}(2F)^2}. \\ \Rightarrow \Delta_{\text{long}}(F, \alpha, \lambda) &\approx z_R(\lambda, F) \cdot \delta(F, \alpha) \\ &\approx \frac{12 \times 10^{-3}\lambda}{(\text{arccot}(2F)^2) \cdot \left((F - 9 \times 10^{-3} \cdot 10^{\alpha/2})^2 + 65 \times 10^{-3}\right)}. \end{aligned} \tag{5.10}$$

As for Δ_{trans}, α has to be given in radians. Δ_{long} is the best RMS that can be guaranteed to be achieved in longitudinal alignment for a specific setup and is shown in Fig. 5.14 on the right. Obviously, even the best setups only give a best guaranteed RMS of approximately 50 mλ.
However, longitudinal misalignment primarily causes defocusing of the beam, as shown in Fig. 5.9 on the right. While still undesirable, defocusing can easily be counteracted using conventional alignment methods, e.g. by using a shearing interferometer, or even post-compensated by another telescope.

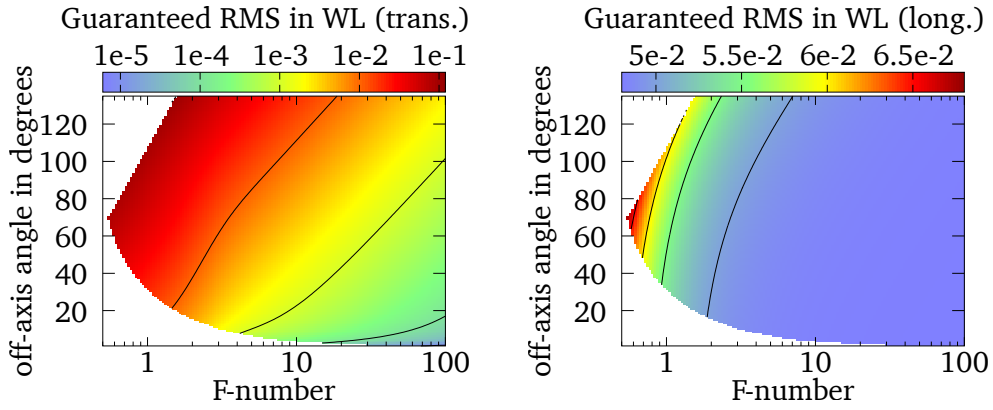

Figure 5.14.: The best RMS that is guaranteed to be achieved by the proposed alignment scheme in transversal directions (left) and longitudinal direction (right), given by equation (5.9) and (5.10) respectively. The black contour lines correspond to the ticks of the color palette.

Application to PTAS Now, as the WF errors can be estimated using equation (5.9) and (5.10), we can apply the parameters for the OAP telescope in PTAS: the incoming beam has a diameter of up to 280 mm, while the effective focal length of the first OAP is 400 mm. This yields an F-number of 1.43. Together with the off-axis angle of 45° (or $\pi/4$ radians, as required by the equations), we can calculate the best guaranteed WF RMS:
For the transversal misalignment, we get $\Delta_{\text{trans}} = 18.4\,\text{m}\lambda$. If this was composed out of pure astigmatism in the Zernike sense, this would correspond to a Strehl ratio of almost 0.99 for a homogeneously illuminated pupil, which is easily good enough for this application.
For longitudinal misalignment, we get $\Delta_{\text{long}} = 51.8\,\text{m}\lambda$. For pure defocusing, this would correspond to a Strehl ratio of about 0.9, which stresses the role of PTAS as a WFS. However, as discussed earlier, the defocusing can still easily be corrected using the SHS of PTAS and the main beam of PHELIX. Any focusing error left can then be compensated when aligning the OAP in the experiment chamber.
Therefore, we have shown that the proposed alignment procedure can in fact deliver a precision which is easily sufficient for this application.

5.2.4.2. Technical Realization

As discussed above, the critical part of PTAS concerning alignment is the OAP telescope in the vacuum vessel. The beam path outside of the chamber, however, can be aligned employing conventional methods: In this case, I used the back-propagating alignment laser, shown in Fig. 5.4. The beam is generated by collimating the light from a single-mode-fiber exit with a large F-number lens and therefore features low initial aberrations, and then injected through a beamsplitter next to the camera box. Using the WFS of the camera box and a 0° turning mirror, this beam can be used to align the optics on the breadboard in a double-pass configuration successively.
After that, the back-propagating alignment beam stayed in place in order to generate a focal spot of the small OAP of the OAP telescope in order to overlap it with the focal spot of the large OAP.
I achieve the spatial overlap by using an intermediate focal diagnostic, shown in Fig. 5.16, which I described previously [2] and will do so in more details here.
The alignment setup is in total composed of four components:

1. **A focal fiducial** marks the destination position of the focal spot. As micrometer precision is necessary, the fiducial has to hold up with this requirement. Therefore, I used an etched tungsten needle with an exceptionally sharp tip to achieve the required precision. An image of the neede in the final setup, including the focal spot, is shown in Fig. 5.15.
   The needle is mounted on a motorized linear stage to remove it from the focal area once the position has been marked or the alignment is completed. Both needle and stage are shown in Fig. 5.16 in the top view on the left and in different configurations in Fig. 5.17.

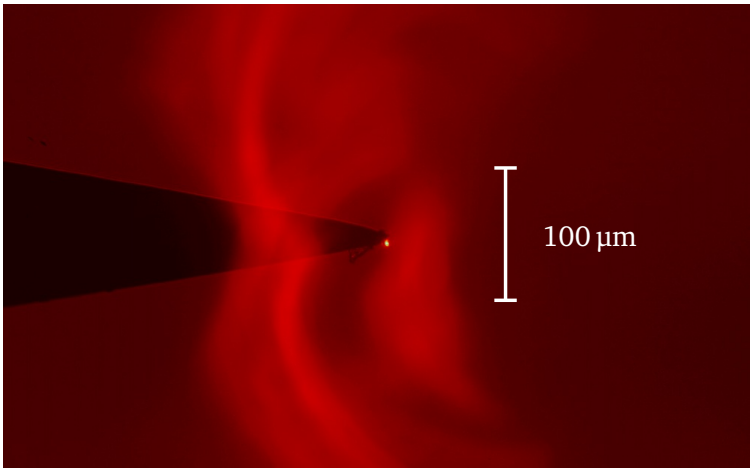

Figure 5.15.: Microscope view of the etched tungsten needle, marking the position of the focal spot focal spot (center). A bit of dust has stuck on its tip. The large caustics are created by backlighting which was used in order to picture the shadow of the needle for this illustration.

2. **Two 0.5" turning mirrors** are used to deflect either the main or the alignment beam after the intermediate focal spot 90° upwards. This is necessary due to the compactness of the telescope as the microscope is too large to directly insert it into the beam. Like this, one can observe the focal spot from above (see Fig. 5.16, on the right). The mirrors too can be moved in or out using a motorized linear stage. In total, three configurations are possible with the mirrors and the needle (see Fig. 5.17): operation, alignment of the first OAP and alignment of the second one.

3. **A high NA microscope** with a vacuum compatible camera can now observe the focal spot via the turning mirrors. The microscope has to be able to spatially resolve the focal spot in order to guarantee that the individual OAPs are well aligned. I used a "20× Plan Apo NIR" infinity-corrected objective (Mitutoyo Corp., Japan), combined with an $f = 200\,\mathrm{mm}$ plano-convex lens to achieve this, acquiring images using a 12-bit CMOS camera (Manta G-235B, Allied Vision Technologies GmbH, Germany). Cooling it under vacuum is significantly simplified as this board level camera is split into a sensor and a connector board, connected by a flexible flat cable. The major part of the heat is generated by the power supplies on the connector board, which can easily be mounted to a dedicated heat sink, leaving the sensor cool at all times.
   The microscope is mounted on a XYZ stage in order to center the view on the needle in both alignment configurations and bring the tip in focus.

4. **Motorized OAP mounts** provide the necessary degrees of freedom to align the focal spots to the needle tip. Each OAP needs five axes in total: XYZ translation and tip/tilt. The rotation around the optical axis of the mother parabola does not contribute to the alignment. While all five axes could be confined into a single mount for the small OAP, the large one does only feature tip, tilt and defocusing. The last two degrees for this one were supplemented using the tip and tilt axes of a large diameter turning mirror prior to the OAP.

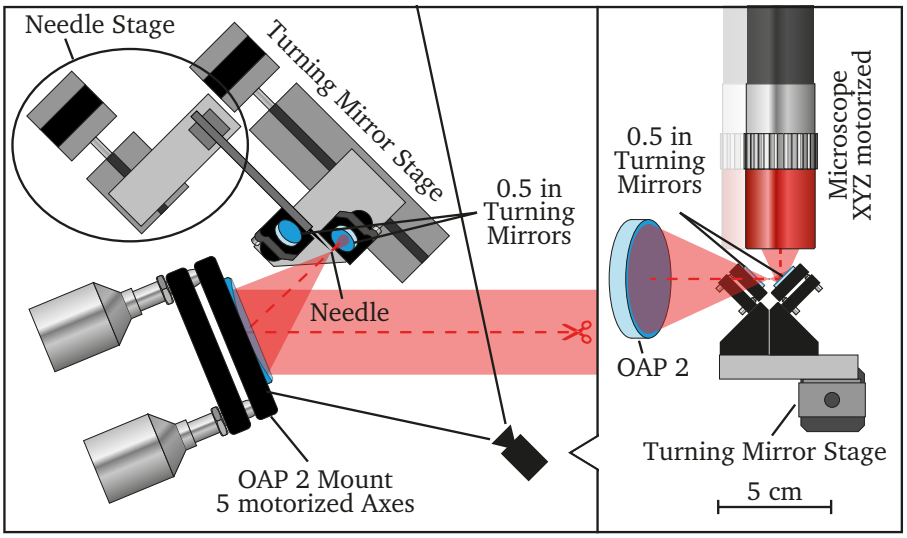

Figure 5.16.: Schematic of the intermediate focal diagnostic of PTAS in top view (left) and side view (right), true to scale, observing the backwards propagating alignment beam. The needle and the diagnostic mirrors can be moved in and out of the beam, while the microscope can be moved in three axes in order to change the viewing position. The other possible configurations of the diagnostic are shown in Fig. 5.17.

5.2.4.3. Alignment Procedure

Using this setup, my alignment procedure can now be performed within a time frame of less than ten minutes. In this section, I briefly describe the steps of this procedure, following my preceding work [1] yet adapting the description to the specificities of PTAS. These steps must be repeated for both of the OAPs of the telescope and in both cases, one should be wary about the laser source: a low power Continuous Wave (CW) beam should be used in order to avoid inflicting damage to the setup. Images of the focal spots, corresponding to these steps, are shown in Fig. 5.18.

0 Coarse alignment - Before the fine alignment procedure can be started, the telescope has to be coarsely aligned, including the needle marking the desired focal spot. This is only necessary when initially assembling the telescope or when major changes have been done. I describe this procedure in the supplementary document of our corresponding publication [1].

1 Insertion of fiducial and diagnostics into the beam - For the optimization of each OAP, one has to bring the needle, turning mirror and microscope into place to observe the focal spot.

2 Determination of target position - In order to observe the tip of the needle, one has to defocus the beam using the corresponding OAP until the needle is broadly backlit [Fig. 5.18, 2a)]. The exposure time of the camera or the power of the laser must be adapted accordingly. Now, the microscope can be moved until the tip of the needle is in the center of the image and sharply focused [Fig. 5.18, 2b)]. The tip indicates the target position where the focal spot shall be located, and has to be marked on the screen [Fig. 5.18, 2c)]. Now, the needle can be removed from the beam.

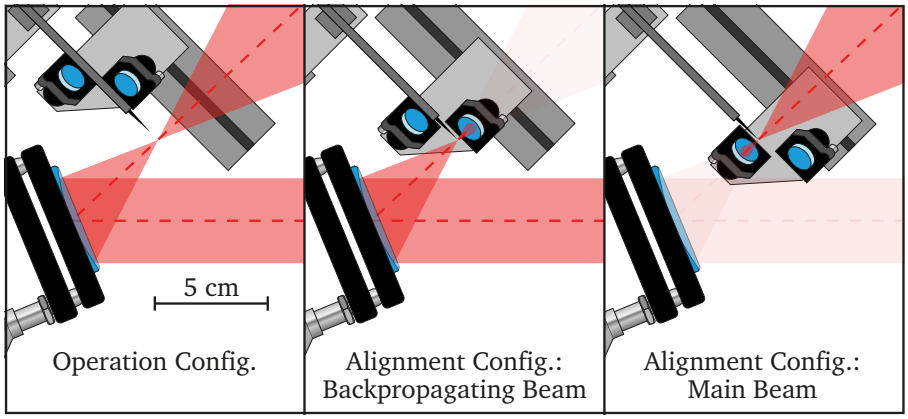

Figure 5.17.: Schematic of the possible configurations of the intermediate focal diagnostic of PTAS in top view. The microscope has to be moved in order to observe the focal spot, as shown in Fig. 5.16 on the right.

3 **Focal spot optimization** - In the next step, the focal spot must be optimized by minimizing oblique and normal astigmatism successively over several iterations. While doing this, one has to carefully monitor the beam power in order to avoid camera damage.
One starts the optimization loop by defocusing the beam until a sharp line appears on the camera [Fig. 5.18, 3], which is due to the remaining astigmatism of the beam. Now, iterating the following steps optimizes the focal spot:

   a) **Minimization of the oblique astigmatism** - Tilting the OAP vertically changes the amount of oblique astigmatism, aside from introducing tilt and defocusing. The latter two have to be continuously counteracted in order to keep the line sharp and in the center of the image. This works differently for the small and the large OAP of PTAS:
   Small OAP: As the small OAP features motorized translation in three axes, the defocusing and tilt introduced by tilting the OAP can counteracted by translating the OAP vertically and in focal direction.
   Large OAP: The large OAP cannot be translated vertically and horizontally. Therefore, the astigmatism is instead changed by tilting the preceding large turning mirror vertically, while keeping the line centered and in focus by tilting the OAP vertically in the opposite direction and translating it along the focal direction respectively.
   For both OAPs, this has to be repeated until the line is perfectly horizontal or vertical [Fig. 5.18, 3a), bottom]. If the line gets longer [Fig. 5.18, 31), top], one has to change the direction.

   b) **Minimization of the normal astigmatism** - This works analogous to the previous step, just by replacing vertical movements with horizontal ones. The goal is to shorten the line to a minimum. If the line starts to tilt [Fig. 5.18, 3b), bottom], one has to reiterate the minimization of the oblique astigmatism. If no oblique astigmatism is left, the line will shrink to the ideal focal spot of the corresponding beam [Fig. 5.18, ”3b), 4”].

4 **Placement of the focal spot to the marked position** - In the last step, the focal spot has to match the target position vertically and horizontally (the last axis is ensured by keeping the beam in focus). Again, this is slightly different for the two OAPs:
Small OAP: Here, the focal spot can easily be moved to the marked position by translating the OAP.

Large OAP: Moving the focal spot works exactly the other way around as optimizing it. The spot can be moved by tilting the OAP, while the introduced defocusing can be counteracted by moving the OAP in focal direction and the astigmatism can be minimized by tilting the preceding turning mirror.

5 **Diagnostic removal** - Once the telescope is fully aligned, the turning mirrors have to be removed from the beam path so the main beam can pass through the telescope. One should also make sure that the fiducial is far away from the focal region.

If the quality of the alignment beam is sufficient, i.e. one can observe the first diffraction ring around the focal spot, one can minimize the astigmatism even further between steps 4 and 5 by trying to homogenize the first dark ring around the focal spot [1]. In all other cases, the residual aberrations of the beam are larger than the WF errors caused by misalignment of each individual OAP, which makes further optimization obsolete.
This completes the alignment of PTAS. In the next section, I will discuss the energy attenuation that has to be used to switch between different energy levels of the laser system.

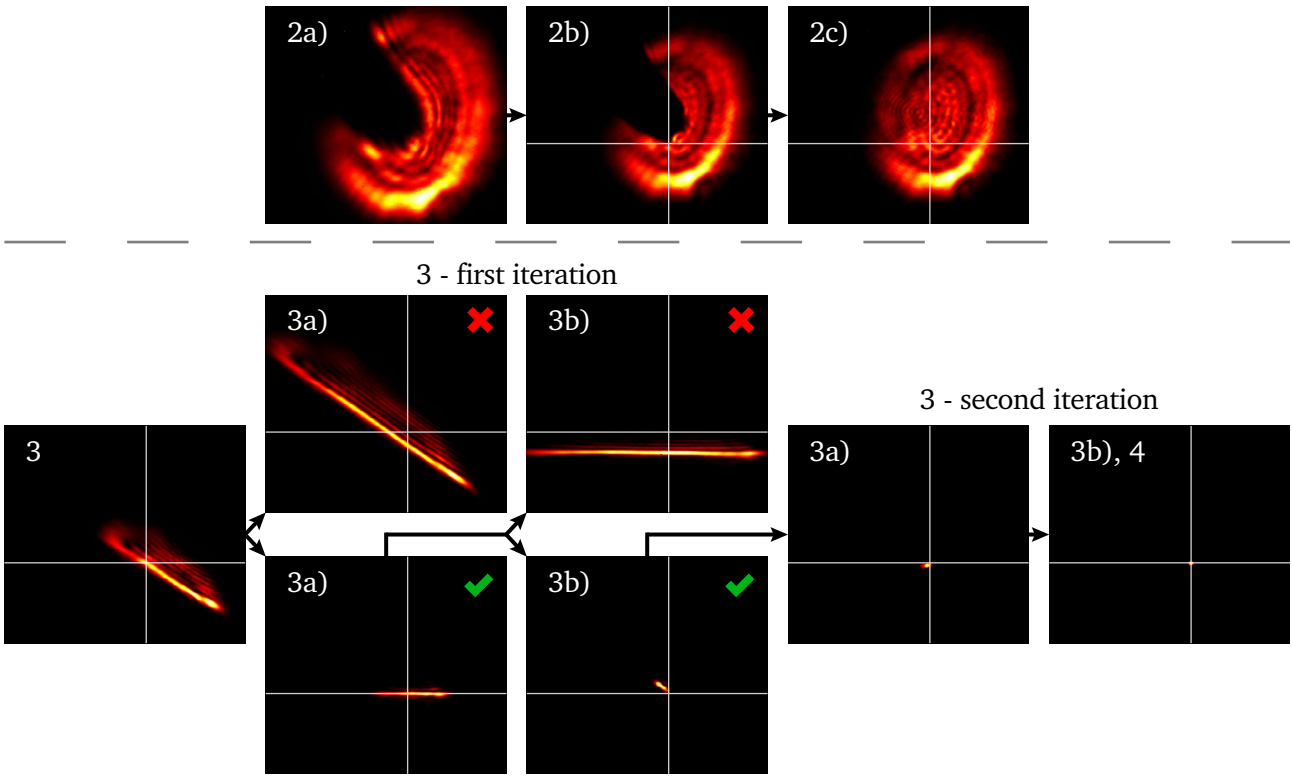

Figure 5.18.: Pictures of the intermediate focal spot taken during the alignment of the small OAP at PTAS. The pictures are labeled corresponding to the description of the alignment procedure above. For step three, two iterations are shown, while the first iteration also shows images of the focal spot where the axes are moved in the wrong direction (marked with a red cross). The last picture covers step 3b) of the second iteration and step four at the same time.

5.2.5. Energy Attenuation

As discussed in the requirement list in section 5.1.3, PTAS needs to cover a dynamic range over the input energy of about 5×10^5 and everything in between continuously. This number stems from the energy span between the 400 µJ for the 10 Hz alignment beam to the full energy shot of 200 J, where the architecture of PHELIX allows for shots with much less energy as well. Therefore, I had to implement several independent energy attenuation stages in PTAS [2]. I discussed this in the overview section already and will go into more detail here.
The first attenuation stage is a Mach–Zehnder Interferometer (MZI) setup right after the OAP telescope and still in vacuum. In Fig. 5.4 on page 45, it is depicted below the large OAP and labeled ”MZI-Atten.”. This structure uses a mirror with a reflectivity of 99.3 % (for s-polarization, as for all optics described in the following) to split a small fraction of the energy into the transmitted light. Here, the NF intensity in the volume of the beamsplitter is less then half of the intensity in the volume of the large leaky mirror before PTAS and therefore safe to transport. After that, the reflected and transmitted parts of the beam are recombined by another beamsplitter with a reflectivity of 40 %. A motorized shutter then blocks either one or the other beam, enabling a selection between transmissions of 59.58 % and 0.28 % respectively. This enables a discrete dynamic range of 2.33×10^2.
Using the 0.28 % option, even a full energy shot can safely be transferred out of the vacuum vessel to the in-air part of PTAS without risking nonlinear effects in beam windows or even damages on optics. At this point, it is worth noting that the two transmitting optics in the MZI setup, as well as the vacuum exit window are wedged in order to deflect double reflections in the glass. These are later removed by a PinHole (PH) of 2 mm diameter between the lenses ”L1” and ”L2” in the first in-air telescope.
Before the beam gets demagnified even further, two more attenuators, labeled ”OD1 Filters” in Fig. 5.4 on page 45 can each reduce the beam energy by one order of magnitude. This is done using two substrates on a linear stage for each filter: one being a beamsplitter with a reflectivity of 90 % under an AOI of 45°, thus transmitting 10 %, the other being a window with equal dimensions to replicate the beam displacement if no attenuation is desired. All surfaces but the beamsplitter are coated with an anti-reflective coating with a reflectivity of less than 0.2 %. This mitigates ghost images to an insignificant amount, even though the substrates are plane-parallel.
With these two filters, the dynamic range spans 4.33 orders of magnitude. The rest of the required 1.37 orders of magnitude can easily be covered by the last attenuator, which is composed of a motorized $\lambda/2$ wave plate to rotate the polarization, followed by a single reflective thin-film polarizer. This setup allows to change the attenuation continuously over almost two orders of magnitude before residual unpolarized light starts to significantly contribute to the signal. Even though this could be further improved by adding a second polarizer into the beampath before the wave plate, this is sufficient to continuously cover the full energy span of PHELIX by choosing the right combination of filter settings. With this, all critical components of PTAS are covered.

5.2.6. Test Results

Over the last sections, I have provided a detailed description of the requirements that are set for the design of PTAS (section 5.1.3), the thoughts and theoretical considerations that went into it (section 5.2.3) and the technical realization of the individual subsystems (section 5.2.2 and 5.2.5). In this section I present test results concerning PTAS and its ability to run a closed AO loop. Here, I consciously exclude aspects concerning the optimization of the focal spot as this is purely a matter of the calibration routine, as I will explain in section 5.3.
Specifically, I will present a qualitative comparison between FF measurements at PTAS and another existing post compressor sensor, a brief analysis of the imaging capabilities, a recorded response of the DM and the measurement of shot aberrations to demonstrate the capabilities of on-shot measurement.

5.2.6.1. Farfield Imaging Quality

As discussed in the requirements section (section 5.1.3), PTAS is not designed to provide an absolute WF measurement, which is hardly possible given the necessity for imaging systems. This also means that the camera in the FF, while in principle imaging the equivalent target plane from the focus of the OAP telescope, does by no means record a fluence distribution that accurately reflects the focus in the actual target plane.
Still, we are interested in a decent imaging quality as this greatly simplifies the estimation of the beam quality that enters the target chamber. Choosing a suitable reference for comparison is arguably a difficult task though as the focusing optic in the target chamber is replaced and realigned between experiments or even exchanged by a different optic.
At PHELIX, the COmpressor Sensor (COS) can provide a suitable reference. This sensor uses leakage light from the very first turning mirror after the last grating of the compressor (see Fig. 5.19) and provides valuable information about the compressor alignment and the state of the gratings.
Of course, this means that the COS cannot be used to run a closed loop with the DM, which has to be placed downstream of this location in order not to block the leakage light, which was the reason to implement PTAS in the first place. The fact that this sensor uses a high-F-number lens ($f = 6\,m$) to shrink the beam promises small aberrations caused by misalignment errors.

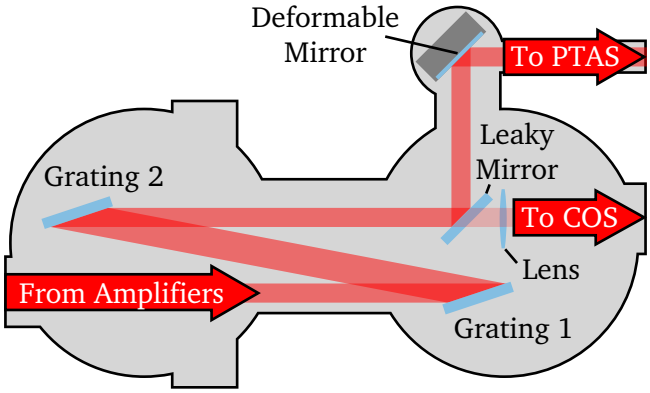

Figure 5.19.: Sketch of the PHELIX compressor, indicating the main optics and the directions to other parts of PHELIX.

Knowing this, a comparison between the FFs at COS and PTAS is a suitable way to judge if the FF at PTAS is of good quality - at least by qualitative measures. For this, no beam correction may be applied to stay close to the ground truth.
Figure 5.20 shows corresponding FFs, recorded at the same time and normalized. The fluence distributions show obvious similarities, indicating a coarsely similar imaging quality of both sensors. However, differences are also visible. These can be attributed to two main reasons:
First, residual differences in alignment may introduce low order aberrations in either of the two sensors, presumably composed of astigmatism and defocusing. These deviations may cause the differences in the shape of the brighter areas in the focal spot.

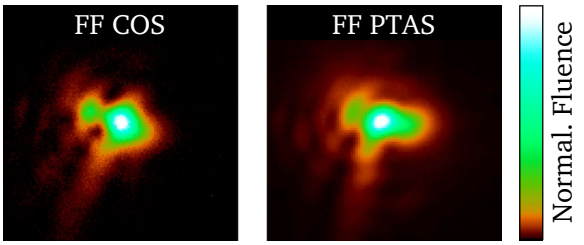

Figure 5.20.: A comparison between the FF recorded at COS and PTAS at the same time (no beam correction).

Second, the small OAP of the first PTAS telescope was milled from copper and therefore, as described before, likely introduces high spatial frequencies into the WF. This may be the reason why the FF at PTAS appears to be slightly blurred out compared to the one recorded at COS, which can be seen in the red-ish areas in Fig. 5.20.
Unfortunately, none of these effects can be quantified due to two reasons. First, no WFS was installed at COS, which could be used to identify differences in low order aberrations. Adding one would require

a complete upgrade of the NF imaging system, which was not justified for the sole purpose of this comparison. And second, no surface measurement of the small OAP in PTAS was available at the time of writing this thesis.
Still, the similarity shows that PTAS can be used as a sensor calibrated to a relative reference, while at the same time aiding the laser operation by providing an approximation of the FF at the target position.

5.2.6.2. Nearfield Imaging Quality

As discussed in section 2, the surface of the DM has to be imaged onto the WFS, where the imaging quality influences the linearity of the response matrix. Therefore, I compared the spatial resolution of the imaging to the one of the WFS. To do so, I recorded images of the NF using the NF camera in the camera box of PTAS, which is located in the same image plane as the WFS, only separated by a single beamsplitter. In Fig. 5.21, the NF of the CW alignment beam of PHELIX is shown on the left and the NF of the front-end beam on the right.
In order to evaluate the imaging quality, I placed an object right in front of the DM (left picture), being in physical contact with the membrane. The quality can be expressed in terms of spatial resolution, which can be measured on the edges of the object.
For this, I extracted local horizontal and vertical lineouts in the regions which are indicated by the white boxes in Fig. 5.21. I fitted a quadratic function to the areas that lie outside of the object and averaged the dark value over the shadow (see Fig. 5.22). From here, I could determine the spatial resolution in terms of the width where the lineout drops from 80 % of the beam value to 120 % of the average shadow value. The result stays below 2 mm both of the times.

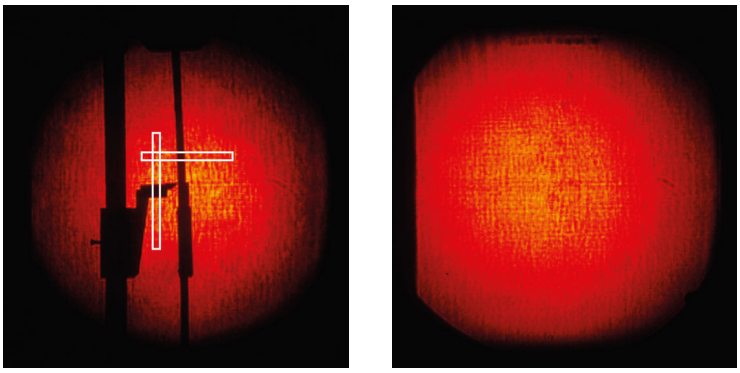

Figure 5.21.: The post compressor NF of PHELIX, recorded using PTAS. On the left, the CW alignment laser is used, while the frontend is used on the right. On the left, an object is placed directly in front of the DM (visible twice due to the reflection of the beam) in order to analyze the imaging quality. The white boxes are the regions that were used to do so.

This is now the spatial resolution of the imaging of the DM plane. To judge if this is sufficient, I invoked the spatial resolution of the WFS: the SHS at PTAS uses subapertures of 150 µm. The demagnification of the imaging system at PTAS adds up to a factor of 64 over all three telescopes, which means that one subaperture captures light from an area with a diameter of 9.6 mm in the NF, which is significantly larger than the spatial resolution determined above. Therefore, the WFS does sample the WF in a localized manner, leading to a linear response matrix. This means that the imaging quality is sufficient for this application.
After I verified this imaging quality, we can now use PTAS for further beam analysis: the NF image of the femtosecond front-end beam of PHELIX, shown in Fig. 5.21 on the right, shows that the beam profile of this source is much wider as of the CW alignment beam, filling out the aperture more homogeneously. Furthermore, features of the second compressor grating are clearly visible: on the left and right, the beam is clipped as the grating is not wide enough to take the full spatially dispersed pulse. On the top, one can see the shadow of a damage that has been imprinted on the grating by a diffraction artifact, which lead to locally raised intensity. After the damage was discovered on this picture, the grating was slightly elevated to move it out of the beam.

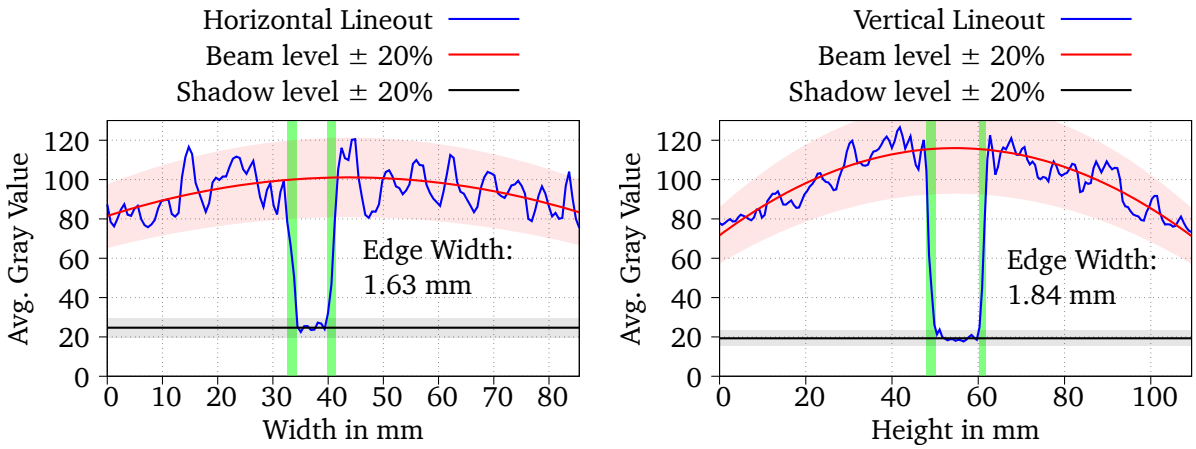

Figure 5.22.: The horizontal and vertical lineouts from the NF shown in Fig. 5.21 on the left, calculated from the marked areas. The green strips mark the transition zone where the average gray value drops from 80 % of the beam level to 120 % of the shadow level. This edge width gives the spatial resolution of the imaging.

5.2.6.3. Response Recording of the Deformable Mirror

As shown in the last section, the plane of the DM is well imaged to the surface of the SHS. At this point, the next step was to record a response of the DM at PTAS and to test the closed loop, even though no reference WF to the target chamber is known at this point.
The response is recorded using the same setup as for the final AO loop. A schematic of the components and the information flow between them is shown in Fig. 5.23 on page 65: the SHS camera records an image, which is sent to the PTAS computer via ethernet, using the GigE camera interface. Here, WOMBAT evaluates the image and calculates a WF, which is used either for response recording or the feedback loop with the mirror. In both cases, a set of control values, i.e. a digital representation of the desired stepper motor positions, is calculated and sent to a locally hosted server application via a TCP/IP connection. This server application, written in C++, hands over the commands (in this case: the stepper motor positions) from WOMBAT to the ISP libraries, which run the ethernet connection to the DM driver. The driver then directly controls the actuators of the DM, which influences the laser beam that propagates to PTAS. The loop is completed when the new WF is recorded by the SHS.
The response is recorded by performing a poke of each individual actuator of the DM at 40 % of its range and then averaging the WF over several seconds. Afterwards, the response is normalized to the full stroke of the actuator. The result is shown in Fig. 5.24 on page 66. Each response plot shows the influence of one actuator on the WF. In all cases, the influence function is smooth and stretches over wide parts of the aperture. This is beneficial for compensating low order aberrations. Combined with the circular distribution of the actuators over the aperture, which becomes apparent when comparing all influence functions, the mirror promises to accurately reproduce the lowest order Zernike polynomials. From the individual actuator responses, WOMBAT calculates the eigenmodes of the mirror, as discussed in section 2. The result is shown in Fig. 5.25 on page 67, where the mirror modes are sorted by their singular value in descending order. Obviously, the largest influence on the WF relative to the required actuator strokes are composed of low order WF deformations, approximating the Zernike polynomials. With rising eigenmode index, the spatial frequencies of the mode increase as well. At the same time, the Signal-to-Noise Ratio (SNR) gets worse up to the point where noise is the dominating component of the mode. This is approximately the case up from number 46. These modes have to be excluded from the AO loop in order to provide a stable WF control.

For localized DMs, the SNR of the response recording could be improved by moving from a simple

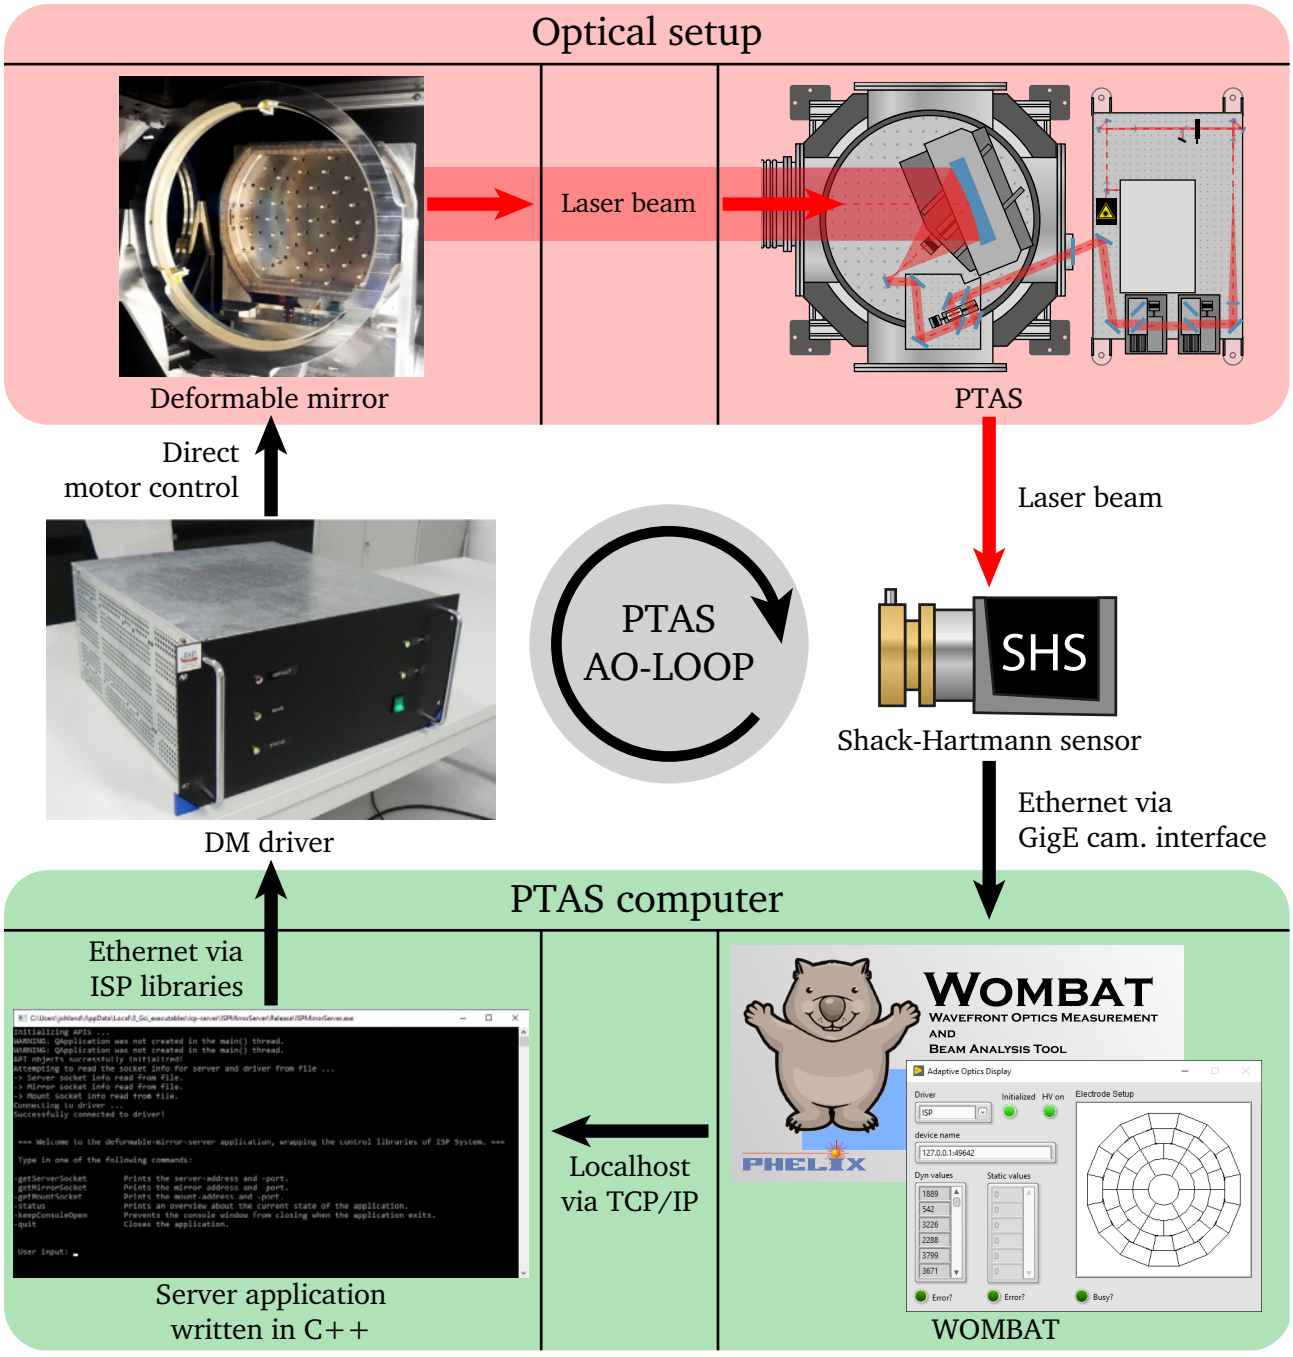

Figure 5.23.: Scheme of the information flow in the AO control loop using PTAS.

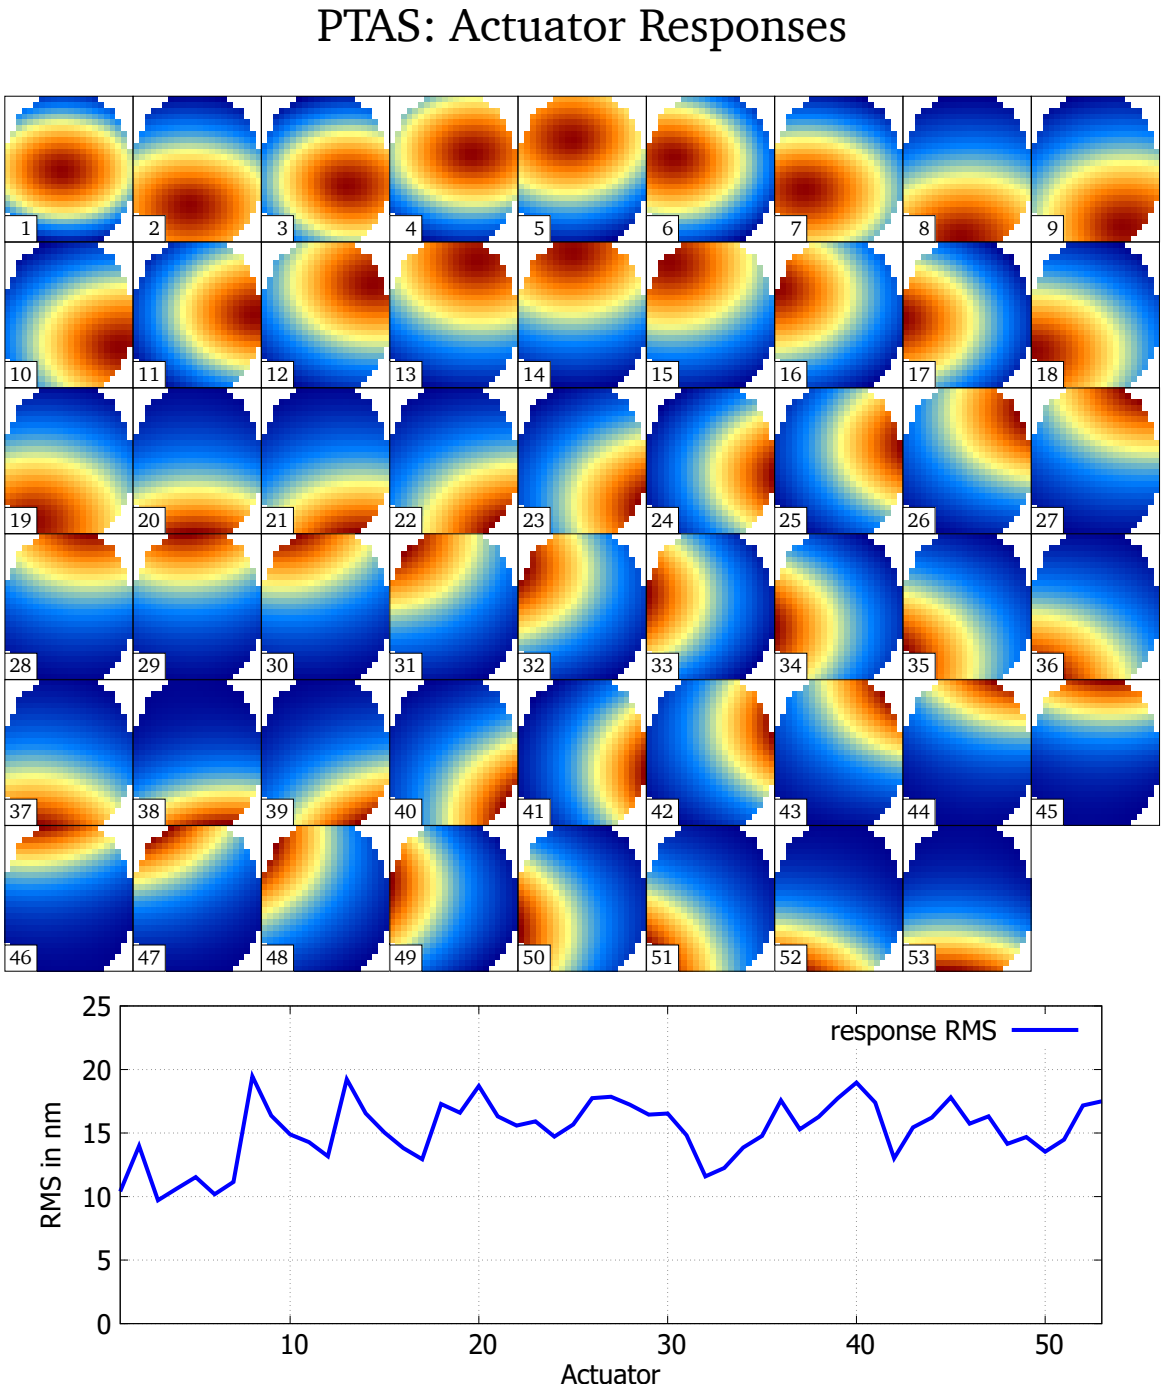

Figure 5.24.: The first response of the DM, recorded using PTAS. Each actuator performed a single poke and the response was normalized to the maximum possible stroke. The associated WFs are shown on the top, while the graph on the bottom shows the associated WF RMS. The beamsize on the mirror was 240 mm.

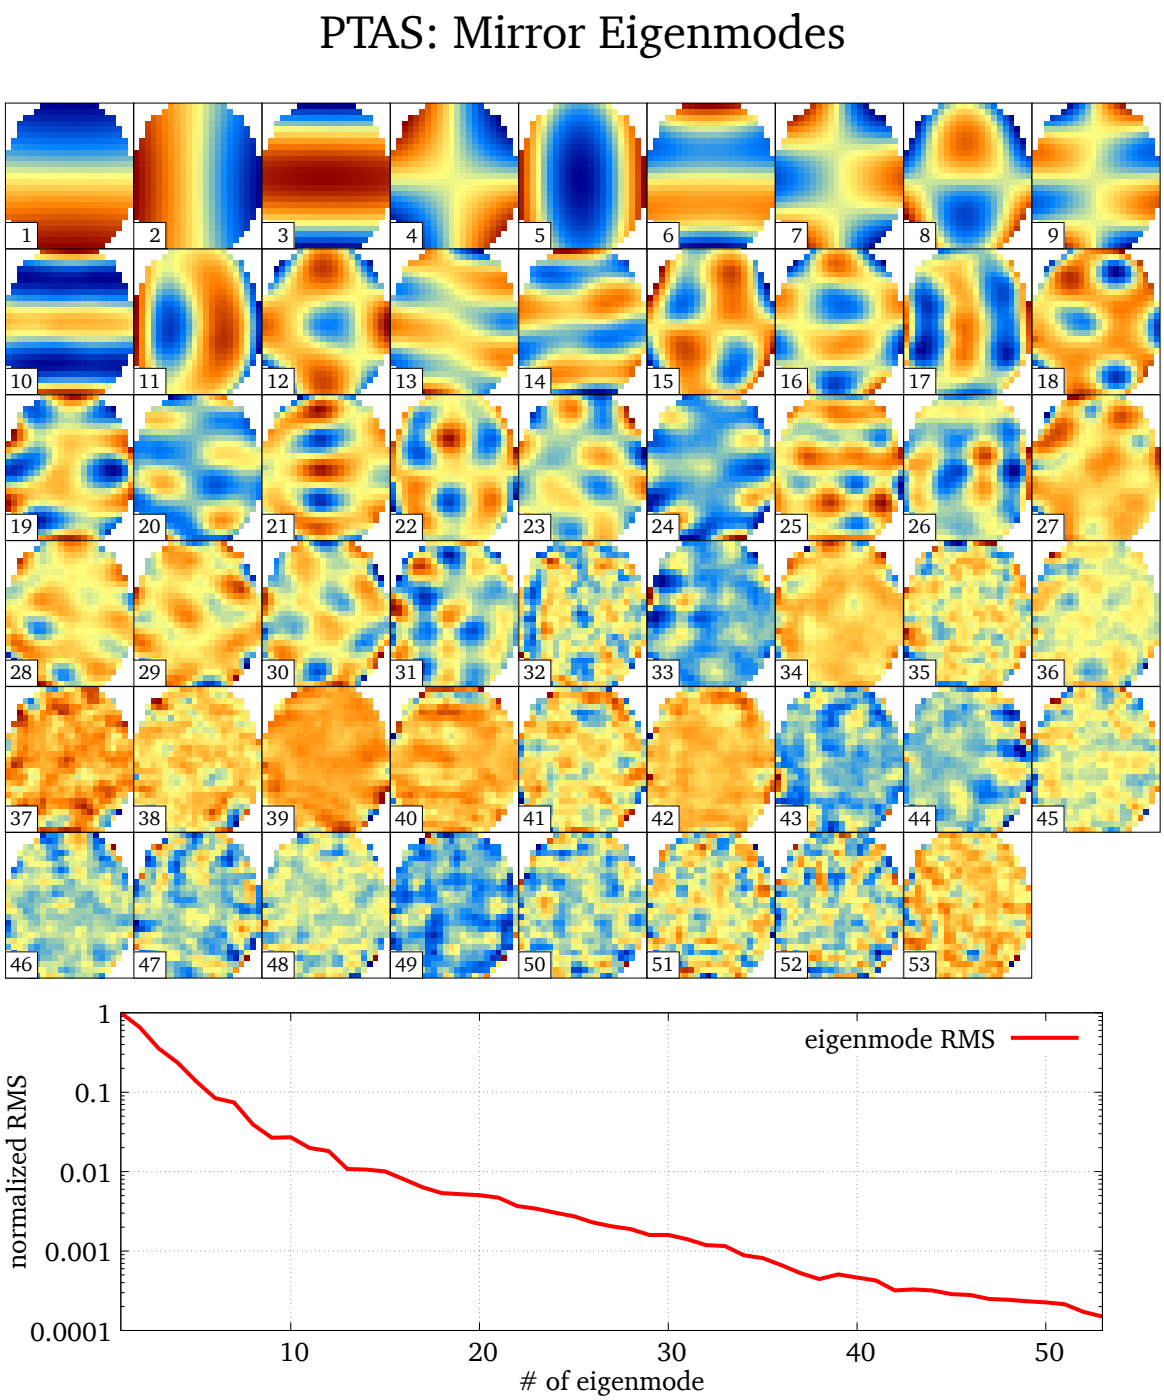

Figure 5.25.: The eigenmodes of the DM, calculated from the actuator responses shown in Fig. 5.24, sorted by singular values in descending order. The associated WFs are shown on the top, while the associated WF RMS, normalized to the eigenmode with the largest singular value, is shown in the graph below.
One can clearly see that the last eigenmodes are dominated by noise instead of mirror response.

actuator poke to a Haramard-encoded poke sequence, which changes the values of all actuators at once in each response step and thus distribute the signal over the whole aperture [33]. This greatly increases the sensitivity of the calibration and reduces the relative influence of noise. In a future work, it could be worthwhile to investigate if this also holds true for DMs with global influence function, such as the mirror used in this work, where the actuator influence spans the whole aperture (see Fig. 5.24).

5.2.6.4. Measurement of the Shot Aberrations

In order to demonstrate the capability of PTAS to measure pulse energies ranging from 400 µJ in alignment mode up to 200 J in full energy mode, I aimed to record the shot aberrations of the system. This equals the difference of the WF before the shot to the WF on shot, which spans the full energy range of PHELIX.
In order to have some reference for comparison, I recorded the shot aberrations at both the Main Amplifier Sensor (MAS) (see section 3.1.4) and PTAS. To do so, I followed the following steps within a timeframe of a few minutes:

1. I collected a total of 20 successive WF measurements at MAS and PTAS in the 10 Hz alignment mode in parallel and averaged them respectively.
2. We switched to the shot mode of PHELIX, but did not activate pre-compensation measures (moving lenses, bending mirrors) from acting out on the WF.
3. I changed the filter settings at MAS and PTAS to prepare for the full energy shot.
4. We charged the amplifier and triggered the shot, recording the WF at MAS and PTAS.
5. I subtracted the averaged WF measurements recorded in step 1 from the on-shot measurements.

The result is shown in Fig. 5.26 in the top row. One can see that the shot aberrations recorded by both devices are mainly composed of defocusing and astigmatism (a Zernike decomposition is shown in Fig. 5.27), where the order of magnitude is comparable between MAS and PTAS. For this reason, two movable lenses are installed at PHELIX which compensate for this amount of on-shot defocusing and a bending mirror compensates for the 0° astigmatism. Therefore, it makes sense to view the shot aberrations after these compensations are applied. In the bottom row of Fig. 5.26, the same measurements are shown, while both defocusing and astigmatism are removed.
Here, we see significant differences between the recordings at MAS and PTAS, where the residual aberrations in the first case appear to be composed of spherical and quadrafoil aberrations, while coma seems to be dominant in the second case. This is confirmed by the according Zernike decomposition and is valid regardless of the thermal state of the amplifiers, as can bee seen in Fig. 5.27.
In order to understand this, we need to consider that the pulse compressor at PHELIX is a Treacy compressor [45] where the beam passes both gratings only once. This means that the beam after compression is spatially dispersed, leading to local WF gradients that depend on the considered WL. In other words, aberrations entering the compressor are converted into STC. These WL-dependent WF slopes in turn generate sub-aperture focal spots on the SHS camera which are spectrally smeared out, too. As the evaluation routine of the SHS can only determine the central position of each spot, PTAS measures the spectrally weighted mean of the local WF gradient, resulting in an according erroneous WF reconstruction.
For this application, however, the knowledge about the spectrally resolved WF is not needed as the goal is only to remove the spectral mean WF distortion using the DM. Therefore, the DM can directly apply a WF offset corresponding to the inverted phase shown in Fig. 5.26, bottom right, in order to pre-compensate the residual mean errors caused by shot aberrations.

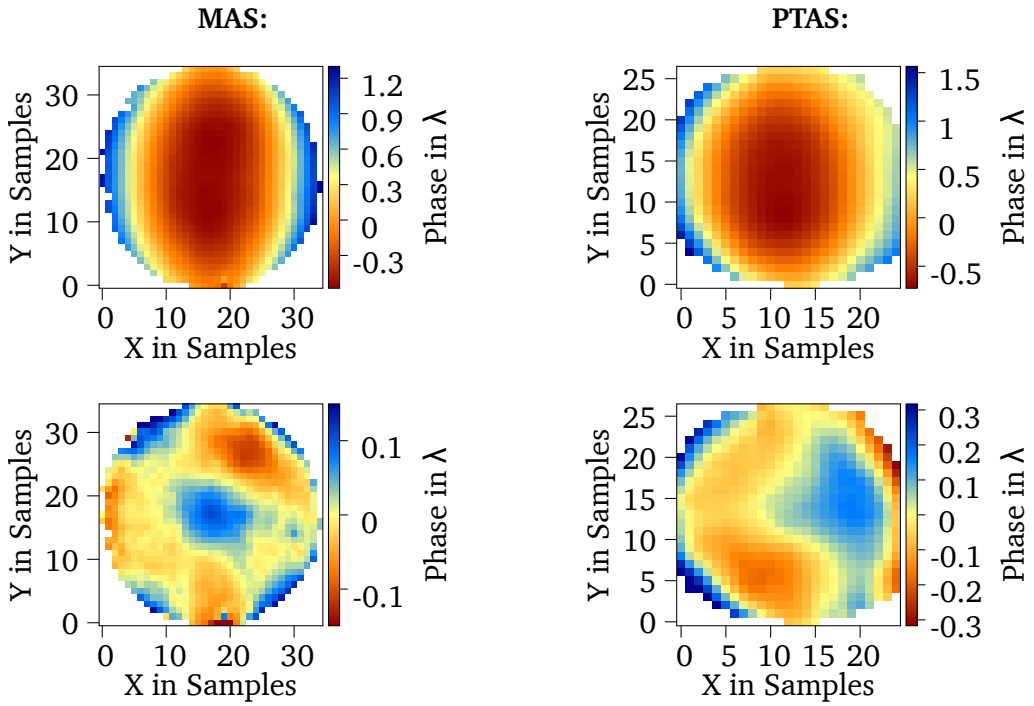

Figure 5.26.: The shot aberrations at PHELIX on the first shot of a day, recorded at MAS (left) and PTAS (right). The top row shows the full shot aberrations while the bottom row shows the same data with defocusing and astigmatism removed.

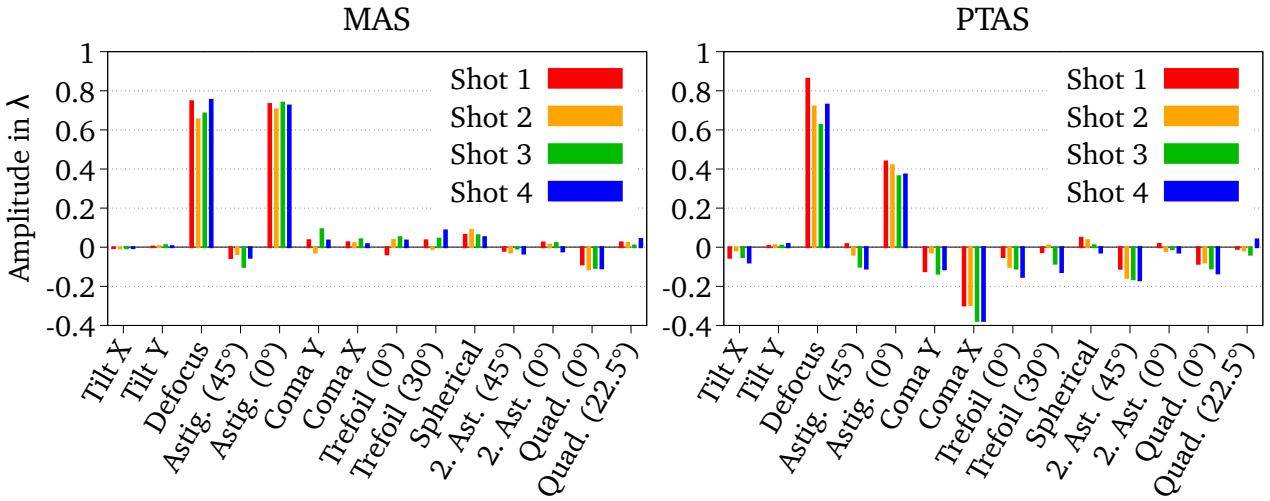

Figure 5.27.: A decomposition of the shot aberrations shown in Fig. 5.26 (”Shot 1”) and three more shots on the same day into the first 15 Zernike polynomials (piston is not shown).

5.3. Calibration Routine

As discussed in section 5.1.3, PTAS can only provide relative WF measurements. For the optimization of the focal spot in the target plane however, knowledge about the absolute WF in front of the target is necessary. This is why a suitable calibration method is an essential component of an AO system in scientific HI lasers (see also section 4.3.2).
This is not a trivial task as the beam is highly divergent and has to be measured in the image plane of the DM at the same time, which rules out using a stand-alone, off-the-shelf SHS. Therefore, I developed a self-calibrating device which serves this purpose.
In this section, I will first describe the calibration concept, which is inspired by preceding work done by L. Doyle in 2019 [74]. Afterwards, I will give a brief overview of the prototype device that I built and tested in this work as a proof-of-principle. I then give an overview of the test results I obtained and conclude this section by lining out the road map for future work.

5.3.1. Calibration Concept

The only way to generate a reference for any beam sensor relative to the WF incident on the target is to observe the beam after the final focusing optic. One approach is to perform the WF measurement directly in the target plane. Here, only the FF is observed, and as a consequence, the phase in the NF has to be determined. This process is generally called Phase Retrieval (PR) [75]. However, these approaches often assume a purely monochromatic beam and require measurements of not only the FF but several planes in front and behind the focal plane. As the PHELIX beam is spatially dispersed after pulse compression due to the employment of a Treacy compressor [45], these methods are ruled out entirely. And while other approaches using machine learning are documented [76, 77] and may be feasible, the investigation of suitable models and the acquisition of sufficiently large data sets to train these go beyond the scope of this work.
An alternative to classical PR and machine learning approaches is currently investigated in the frame of a supervised master's thesis [78]. Here, Orbital Angular Momentum (OAM) shall be added to the beam, as the focal spot is exceptionally sensitive to aberrations with angular variations [79], which can potentially provide precise insight into the NF WF. However, this approach is likely to fail for purely radial aberrations, which makes it more suitable for other applications like the alignment of OAPs and may therefore be used as a refinement of the alignment procedure for OAP telescopes described in section 5.2.4.
Logically, as measurements in the FF cannot be used at this point, the WF has to be measured in the NF behind or in front ot the focal plane. This holds some challenges as the F-numbers of the focusing optics at PHELIX are quite small, which leads to an extreme beam divergence. The image plane of the DM after the OAP of $f = 400\,\mathrm{mm}$ is located 35 mm behind the focus and therefore features a diameter of sightly above 24 mmm which is way larger than the camera chips used in standard SHSs. This means that an imaging system has to be employed, which itself will add aberrations to the beam.
The solution for this problem is to use an in-situ calibration routine for this WFS. Opposed to PTAS, this can be done as the beam passes the actual target plane before entering the calibration WFS.

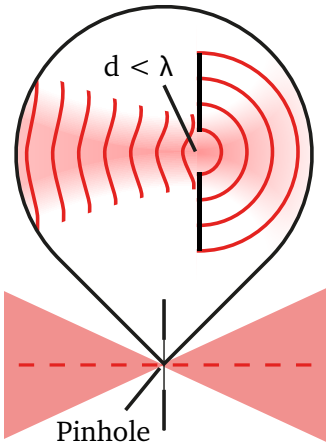

Figure 5.28.: Illustration of the generation of a spherical WF from a beam as used in our calibration concept. The beam enters from the left.

It has been well documented that extreme spatial filtering can generate close-to-perfect spherical WFs [80, 81, 82], which is equivalent to the aberration free case in the NF. To do so, one has to place a PH with a diameter of less than one WL in the focal plane (see Fig. 5.28). Having this perfect WF, one can now calibrate any WFS system after the focal plane, even if large amounts of aberrations are present. Therefore, one can use an imaging system to image the DM to a SHS with a suitable beam diameter.

The concept of this is shown in Fig. 5.29. Here, a microscope lens forms a telescope with the focusing OAP, where the intermediate focal spot corresponds to the target plane. Two more lenses form another telescope which images the DM onto the SHS. The PH can be placed in the intermediate focal spot in order to calibrate the SHS.

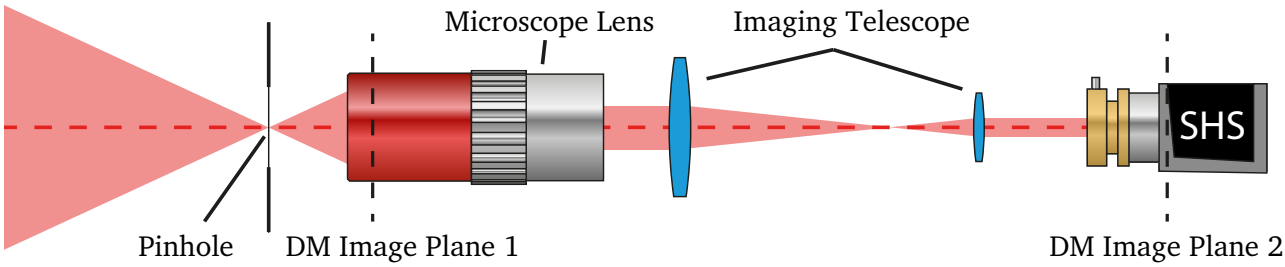

Figure 5.29.: Schematic of the calibration setup, where the beam enters from the left. The PH in the target plane may be removed according to the procedure described in the text.

Using this concept, the calibration of PTAS would work according to the following steps:

1. Place the calibration device behind the focal spot and insert the PH to the focal plane.
2. Record a reference at the SHS. This is the ideal WF.
3. Remove the PH. The difference to the reference corresponds to the aberrations in the beam.
4. Use the DM to optimize the WF with respect to the reference. The result is the best WF that can be reached.
5. Take a reference at PTAS and remove the device from the target chamber.

Now, PTAS has been calibrated. Whenever the WF corresponds to the reference at PTAS, the focal spot takes its best achievable from in the target plane.

5.3.2. Prototype Calibrator

The technical realization of the concept described in the last section holds some challenges. In this work, I built a prototype called Calibrator for Lasersystems with Absolute Wavefront Sensing (CLAWS) which addresses most of these. While the development is still ongoing, I will describe the current state of the implementation in this section and elaborate on the thoughts that went into it.
My goal was to build a device that, first, can provide a proof of concept for this calibration method and, second, could be used in pilot experiments, which I will discuss in chapter 6.
The latter dictates that CLAWS has to be compatible with the corresponding experimental setup in the target chamber. This is due to the fact that the weight of the experimental equipment may deform the breadboard of the target chamber and therefore influence the reference, which is why the calibration has to be done after the experiment has been installed.

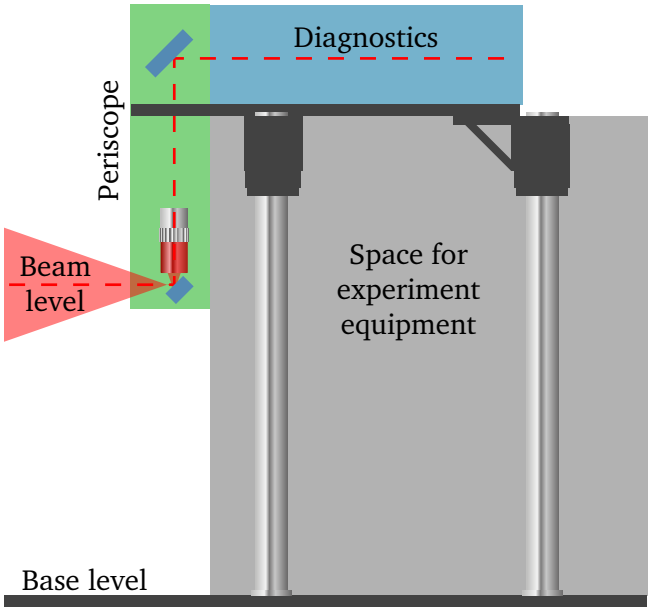

Figure 5.30.: Illustration of the coarse setup of CLAWS. The beam, entering from the left, is elevated above experiment equipment using a periscope.

For this prototype, I decided to realize CLAWS as a platform that can be placed above the experimental setup on four high legs (see Fig. 5.30). A periscope (green block) samples the beam directly behind the focal spot and transports the light onto the platform with the diagnostics (blue block), while imaging the DM onto the WFS. The gray block marks the space for the experimental setup.

This setup can, of course, not be used for any experiment as the experimental setup design has to take the position of the legs of CLAWS into account and allow for the removal of the device. Any future implementation should therefore work without this interference, e.g. by a motorized mount lowering the device from the ceiling of the target chamber.
In the next section, I will explain this prototype in detail on a technical level. Afterwards, I describe the procedure of aligning CLAWS in the target chamber in order to perform the calibration process.

5.3.2.1. Setup Description

A detailed sketch of CLAWS as built in this work is shown true to scale in Fig. 5.31 in different viewing angles. In the following, I will describe the setup and its components using this sketch.

Diagnostics In total, CLAWS features three cameras: the SHS, a FF camera with a wide viewing angle, labeled ”FF LR” in Fig. 5.31 and a 16-bit CMOS camera (pco.edge 4.2, Excelitas PCO GmbH, Germany), labeled ”FF HR”, with a narrow viewing angle of the FF and thus a good spatial resolution. The SHS measures the WF and the FF LR camera is used for the alignment of CLAWS, while the FF HR camera provides a high quality image of the focal spot to evaluate the intensity that can be reached with the current WF.

Platform The platform of CLAWS is made of a 30 cm x 45 cm breadboard which supports all optical components. The breadboard is supported by four 1.5” aluminum posts (not shown in Fig. 5.31, but in Fig. 5.30) with a top surface height of 565 mm. The baseplates of the posts are milled from Polyoxymethylene, which is a rigid plastic with a low friction coefficient. With these base plates, CLAWS slides easily on the breadboard of the chamber into its target position.
These pillar posts can be repositioned in order to comply with a wider range of experimental setups.

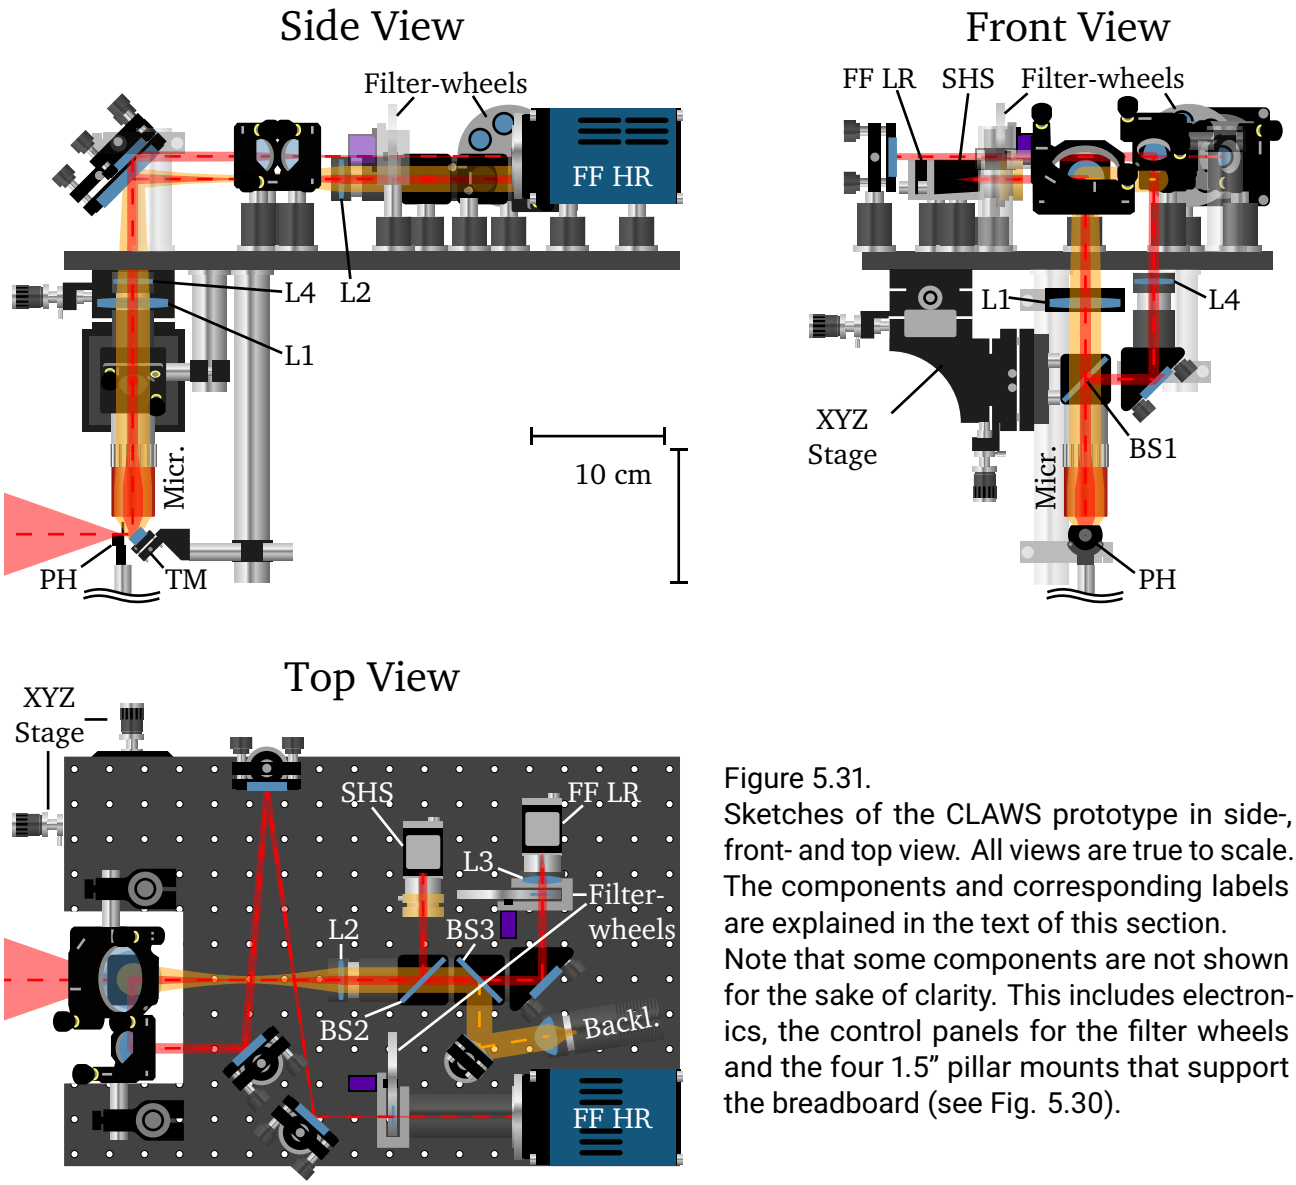

Figure 5.31.
Sketches of the CLAWS prototype in side-, front- and top view. All views are true to scale. The components and corresponding labels are explained in the text of this section. Note that some components are not shown for the sake of clarity. This includes electronics, the control panels for the filter wheels and the four 1.5” pillar mounts that support the breadboard (see Fig. 5.30).

Periscope As the beam height in the target chamber is 360 mm, a periscope elevates the beam level onto the platform of CLAWS. The lower mirror of the periscope is a 0.5” silver mirror, labeled ”TM”, which is directly mounted to the breadboard and deflects the beam upwards right after the target plane. The microscope lens, labeled ”Micr.”, is located above this mirror, collimating the light from the focal spot. The alignment is facilitated by a linear stage with three axes, labeled ”XYZ Stage”, which mounts

both the microscope lens and a beamsplitter (”BS1”), which splits the beam into the WFS beam path and the path to the FF HR camera, via a magnetic mount. For the coarse alignment, the microscope lens and BS1 can be replaced with two PHs.

WFS beam path Right after BS1, the transmitted part of the beam passes through the first lens (L1, f = +125 mm) of the imaging telescope towards the WFS. A downwards facing mirror then reflects the beam horizontally onto the platform, where the second lens (L2, f = +100 mm) of the telescope recollimates the beam and images the DM onto the SHS, which is located in the reflected beam arm of a second beamsplitter (BS2).

FF LR beam path The light transmitted trough BS2 passes another beamsplitter, which is used for the injection of a backlighter, which I describe below. After that, a mirror for alignment steers the beem through a filter wheel into a lens with +50 mm focal length (L3) that focuses the beam onto the FF LR camera.

Filter-wheels The filter-wheels in the setup, located in front of both FF cameras, are used to adapt the illumination of the cameras to different conditions, i.e. laser energies and defocusing, by inserting absorptive neutral density filters. For fast and repeatable exchange of the filters, the wheels are controlled by servo motors, which are wired to an ESP32 micro controller development board (not shown in Fig. 5.31). This board takes inputs of push buttons which are placed on three sides of the breadboard for easy access (also not shown in Fig. 5.31) and increment or decrement the position of the filter wheels.

PH Backlighter Aligning a 1 µm PH to a focal spot is a challenge by itself. In order to aid this process, I installed an infrared LED light source with a condenser lens (labeled ”Backl.” in Fig. 5.31). This light is injected backwards into the beam path via BS3, as indicated by the yellow beam area. As the LED generates spatially incoherent light, it illuminates an area around the focal spot. This helps finding the center of the PH, as described in the next section.

FF HR beam path The last component is the beam path to the FF HR camera, which starts at BS1 in the periscope. The reflected part of the beamsplitter is deflected upwards by a mirror into a f = +750 mm lens (L4). This lens is the objective of the FF HR camera and the beam is folded over the breadboard such that the camera chip sits in the focal plane of L4. Its long focal range and the absence of any further imaging optics provides a focal spot image with low aberrations, which can be used to estimate the on-shot intensity (see section 5.4) A filter-wheel is installed in front of the camera in order to adjust the illumination of the camera.

5.3.2.2. Alignment in the Target Chamber

As discussed before, the placement of CLAWS in the target chamber and the alignment of the PH pose a challenge themselves. Therefore I briefly describe a typical operation procedure in this section. This is done once the experimental setup is installed (taking into account the position of the pillar posts of CLAWS) as the calibration is sensitive to the weight on the breadboard of the target chamber. Note that the weight of CLAWS should always be compensated for as well when the device is not in place, e.g. by placing lead blocks in the chamber, in order to guarantee consistency between the steps.

1. Coarse alignment As the microscope lens and BS1 are fixed on the XYZ stage via a magnetic mount, these can be replaced with a dedicated iris diaphragm setup to aid the coarse positioning of CLAWS down to the millimeter scale. Here, one diaphragm sits in the desired target plane relative to CLAWS, marking the position where the focal spot must be located, while another one is centered in the former position of BS1. By observing the illumination in both of these planes, CLAWS can be shifted close to the ideal position.

2. Fine alignment After placing the microscope lens and BS1 back into CLAWS, one can observe the image of both FF cameras. Using the XYZ stage to move the microscope lens, the focal spot has to be centered in the FF LR camera image and focused in the FF HR camera. CLAWS is now aligned to the main beam.

3. PH alignment Now, one has to mark the position of the FF on both FF cameras before inserting the PH using the target manipulation stage of the PHELIX target chamber. By switching on the backlighter LED, the PH can be moved using the target manipulator in order to bring the back surface into focus find the PH opening (see Fig. 5.32). Here, the edges of the PH mount can act as guidance.
Once the PH is found, it can be aligned to the previously marked FF position. Now, light should be transmitted through the PH and the backlighter LED can be turned off.
In the last step, the position of the PH can be fine tuned until the amount of transmitted light is maximized, meaning that the PH now matches the position of the actual focal spot.

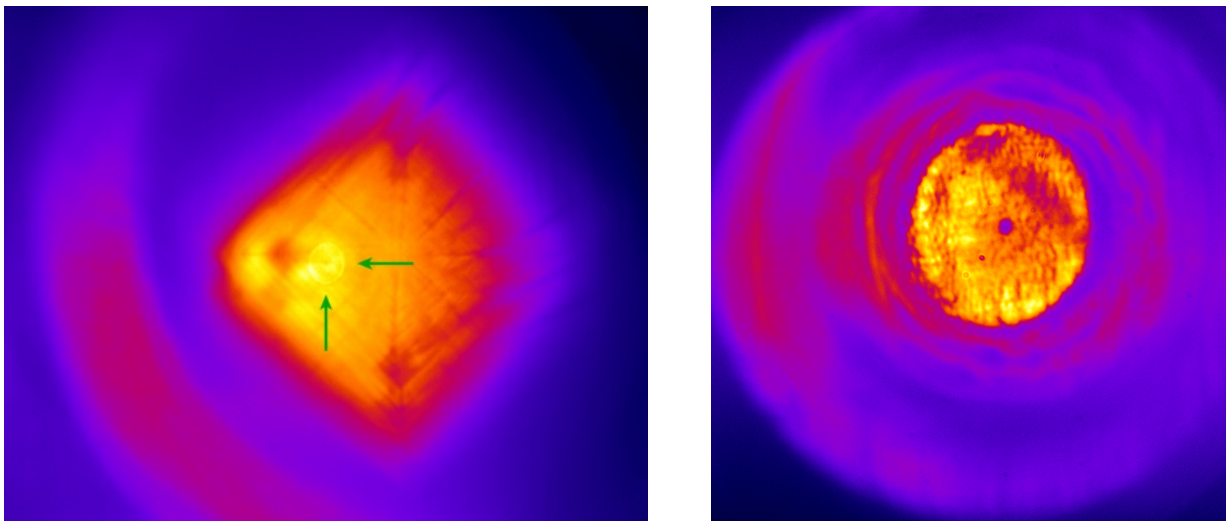

Figure 5.32.: Full images of the backlit PH, recorded using the FF LR camera (left) and the FF HR camera (right). In both cases, the actual hole is visible as a dark spot surrounded by a bright, circular area, which is an indent milled into the substrate.
On the FF LR image on the left, the LED chip of the backlighter is visible as a large diamond shaped area, which significantly lowers the contrast of the PH image (marked with green arrows). This is an artifact as BS3, which is used to inject the backlighter into the beam path, is a beamsplitter cube. The anti-reflective coating is insufficient as the transmitted part of the backlighter gets reflected on the rear side of the cube, passed the splitting surface again and gets partially reflected into the FF LR camera, which images the LED chip. The next version of CLAWS should include a replacement of BS3 with a plate beamsplitter or a cube with a suitable anti reflective coating.

4. Calibration procedure At this point, the intended calibration procedure can be performed, as described in section 5.3.1. Here, one encounters the problem of low illumination of the WFS, as the PH only transmits a small amount of the available light. For the recording of the reference at CLAWS, we generally use the CW alignment beam of PHELIX, coupled with dimmed room illumination and a long exposure time of the SHS. Once the PH is removed, however, one should switch back to the main beam, which matches the on-shot pulse more closely in terms of spectral distribution and illumination of the pupil.

5. Focal spot characterization Once the calibration routine is completed, an image of the focal spot should be recorded using the FF HR camera. This can be used to estimate the on-shot intensity, as I describe in section 5.4.

After that, CLAWS can be removed from the target chamber, which concludes the calibration routine.

5.3.3. In-Air Results

The current version of CLAWS can only be operated in air, mainly due to the cameras which cannot be cooled in vacuum. Therefore, the test results presented in this section have to be interpreted as a mere demonstration of the concept as pumping down the vacuum chamber and the compressor leads to deformations in the optical path and thus to loss of the reference taken at PTAS. One can expect, however, that a vacuum-compatible version of CLAWS will yield results with similar precision. Furthermore, I will show how to still make use of the in-air calibration in section 5.4.

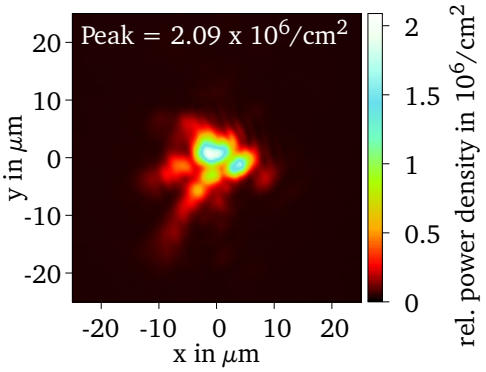

Figure 5.33.
The uncorrected focal spot in the PHELIX target chamber, using an f = 400 mm, 45° copper OAP.

In the following, I show the data of a corresponding calibration routine as a preparation for an experimental campaign. Here, a f = 400 mm, 45° copper OAP was used to generate the focal spot. As this kind of OAP is mechanically milled from copper, which is heavy and flexible itself, the aberrations introduced by this OAP are significant.

The uncorrected focal spot, recorded using the FF HR camera of CLAWS, is shown in Fig. 5.33. One can clearly see the effects of the aberrations. The relative power density was calculated following the procedure described in the following section 5.4.
From this point, the calibration routine described above was performed. The focal spot, corrected at CLAWS, is shown in Fig. 5.34 on the left. It is obviously much closer to the diffraction limit in this case. The peak of the relative power density multiplied by the peak laser power, is almost 3.4 times as high as in the uncorrected case.
The corresponding residual WF, also recorded using CLAWS, is shown on the right. It shows the deviations from the spherical reference wave obtained using the pinhole. Obviously, the low order aberrations are successfully removed and that the residuals are primarily high order aberrations which cannot be corrected with the DM.

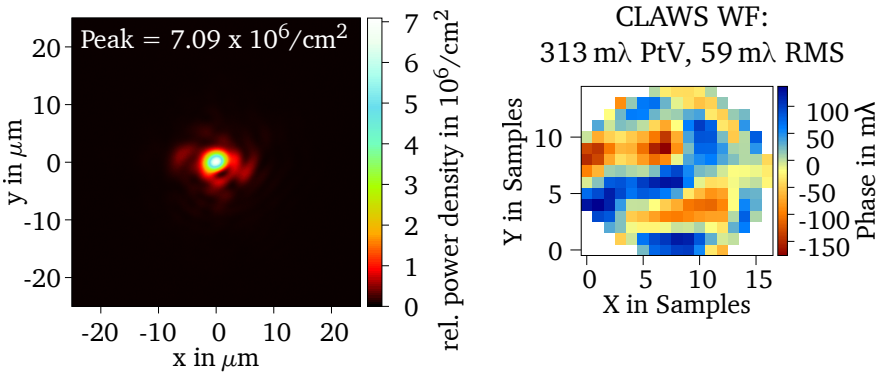

Figure 5.34.: The focal spot (left) and residual aberrations (right) at CLAWS after the SHS was calibrated and the closed AO loop performed.

Finally, the optimized WF was taken as a reference at PTAS. I set the DM actuators back to their factory positions and ran the closed loop using PTAS to see the accuracy of the calibration. The result is shown in Fig. 5.35, where the focal spot and WF at CLAWS are shown on the left and in the center respectively and the residual WF errors relative to the calibration reference at PTAS is shown on the right.

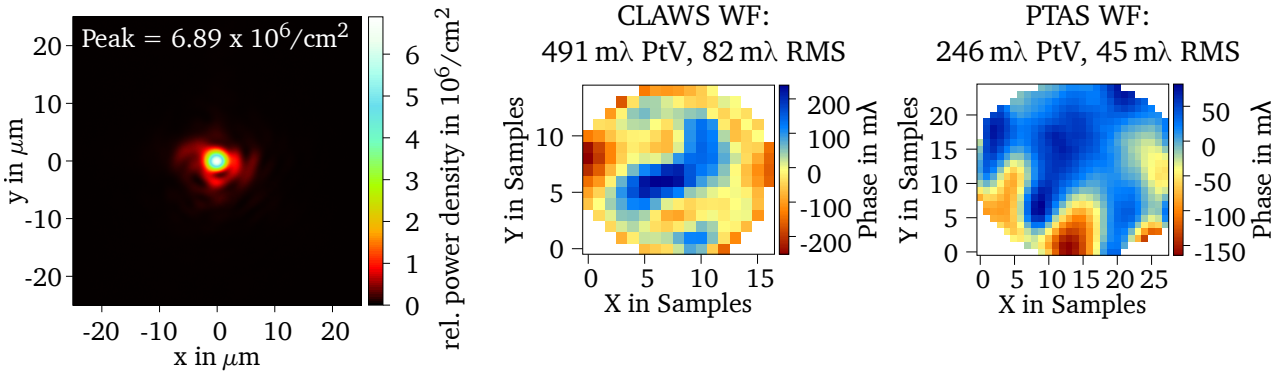

Figure 5.35.: The focal spot (left) and WF at CLAWS (center) and PTAS (right) after the closed loop re-optimized the WF using the WFS at PTAS. The WF residuals at CLAWS correspond to the deviations from the spherical reference, generated by the PH, while the residuals at PTAS reflect the difference to the calibration WF obtained using CLAWS.

Obviously, the WF could be approximately recreated. However, some residual errors of 45 mλ RMS were introduced by the closed loop, which slightly degraded the focal spot intensity. This is presumable due to differences in the recorded mirror responses between CLAWS and PTAS. The new intensity boost, compared to the uncorrected focal spot, is now a factor of 3.3, which is still a large improvement.
As shown above, the impact of the OAP on the beam is significant. While the AO loop is capable of absorbing the low order components of these aberrations, higher order artifacts cannot be compensated for, as can be seen in the residual WFs in Fig. 5.34 and Fig. 5.35.
In order to demonstrate the impact of the OAP imperfections on the optimization quality, I show the corrected focal spot using a high quality glass OAP of identical geometry in Fig. 5.36. This focal spot corresponds to the one shown in Fig. 5.34 from a procedure point of view, where the main difference is the OAP. We can clearly see that the reached peak intensity is 7.9 % higher, solely due to the higher quality of the focusing optic.

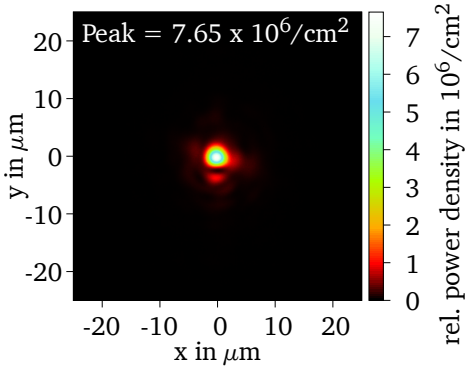

Figure 5.36.
The focal spot in the PHELIX target chamber after the calibration routine, using an f = 400 mm, 45° glass OAP

The residual deformations of the focal spot can still be attributed to high order aberrations and steep WF slopes at the edges of the aperture, which are hard to correct for. However, another possible source are chromatic aberrations in the beam, which cannot be corrected for with a DM and are yet to be characterized. This is especially threatening due to the Treacy compressor geometry [45] that is used at PHELIX.

Concluding this section, I point out that these results demonstrate the suitability of this calibration approach. While further work has to be done, which I will discuss in the next section, the calibration routine clearly works as intended and promises to boost the on-shot intensity of PHELIX by a factor of three or more.

5.3.4. Roadmap for Future Work

As stated above, the state of CLAWS presented in this work is only designed to deliver a proof of principle. Several things have to be done in order to make this device suitable for regular operation. I will present these issues alongside with possible solutions in the following.

Vacuum Compatibility As stated above, the presented version of CLAWS cannot be operated under vacuum, mainly due to a lack of cooling of the cameras. This is problematic as the atmospheric pressure does lead to a deformation of the beampath despite the efforts to decouple these effects. At PHELIX, this is especially problematic as the compressor vessel has to be under atmosphere as well if the full beam shall be measured in the vented target chamber, leading to further deformations and prolonged evacuation times.
In order to prevent the loss of calibration at PTAS due to deformations associated with the pump-down procedure, a vacuum compatible version of CLAWS has to be designed. Specifically, the following tasks have to be tackled:

- Usage of vacuum compatible cameras with suitable cooling
- Performing a survey on the vacuum-compatibility of the electronics and actuators associated with the filter-wheels
- Switching to a motorized alignment stage of the microscope lens to control the fine alignment of CLAWS from outside the chamber

Once these steps are completed, one can calibrate PTAS in vacuum with the full precision.

Reduced Interference with Experimental Setups As already mentioned, the calibration routine has to be performed with the experimental setup in place in order to account for deformations in the target chamber breadboard caused by the weight of the equipment and mechanical stresses induced by fastening the components.
Currently, the platform of CLAWS is supported by four 1.5” pillar posts to be mounted above the

experimental setup (see Fig. 5.30). While this works for experiments specifically designed to accept CLAWS, this is a huge disadvantage for the usage in regular experiments as the experimental setup not only has to offer the space for the pillar posts, but also needs to allow for the removal of CLAWS after the calibration is complete. Therefore, the need for a more detached solution arises.
A possible pathway to reach that would be to mount CLAWS onto a motorized stage attached to the ceiling of the chamber, e.g. in the fashion of an overhead crane. This would remove the necessities for pillar posts entirely, leaving only very little demands on the experiment to temporally leave some space in direct proximity to the target plane where the periscope of CLAWS has to ”dip in” from above.

PH Replacement The PH that I used to demonstrate this calibration concept is a mechanical hole, machined into a thin stainless steel foil, blackened on both sides. While this has the advantage that no bulk material is in the way of the laser light, imperfections at the rim of the steel can have drastic influences on the light transmission. Specifically, I observed a shift of the transmitted intensity maximum up to the point where the NF at the SHS was only partially illuminated, which obviously has a negative impact on the calibration routine. This shift in the NF is equivalent to a tilt of the WF in the focal plane, which itself must be caused by the PH geometry.
A way to prevent that would be to switch to a lithographically etched PH, where the PH is formed by a metallic coating on a transparent substrate. This means that the rim of the PH is extremely thin, presumably introducing less imperfections into the beam.
In this case, however, one would need to adapt the procedure to the bulk material of the substrate in the focal plane.

Improvement of the Backlighter Beampath As discussed above, the current implementation of the LED backlighter is less than ideal as a significant portion of the light is directly reflected to the FF LR camera by BS3 (see Fig. 5.31), significantly reducing the contrast of the actual PH image (see Fig. 5.32, left). Here, the solution would be to replace the beamsplitter cube with a regular plate beamsplitter.

These points would lead to better integration of the calibration routine into the experiment and higher precision. Still, the measurements were conclusive enough to allow first on-shot tests, where I will describe the estimation of the on-shot intensity in the next section.

5.4. Estimation of the On-Shot Intensity

After the calibration routine has been completed, the resulting focal spot has to be characterized. On one hand, this can be used as verification that the routine has been successful. On the other hand the obtained knowledge about the focal spot can greatly benefit the evaluation of the experimental data that are obtained in the following campaign (see chapter 6).
While the knowledge of the incoming field is in principle sufficient to simply calculate the electrical field, this task is far more challenging in this case for two main reasons: first, the lack of knowledge about the incoming field and second, the computational effort that is required to calculate the focal spot:
The knowledge is not sufficient as the WF is spatially chirped when entering the target chamber as PHELIX features a Treacy compressor [45], which means that the SHS at CLAWS only measures a spectrally weighted average of the local slopes, loosing the information about the STCs in the beam. The amount of assumptions to circumvent this issues is huge and the associated results do not accurately reflect reality.
And even if the full information was available, the computational effort to calculate the focal spot would

be huge. This is because the diffraction integral that is necessary to calculate the field is on one hand poly-chromatic, requiring the sum of a multitude of integrals, and on the other hand not valid in terms of the small-angle approximation due to the small F-number of the focusing OAP. This means that each monochromatic integral has to be expressed as sums over samples within the NF aperture and the often-used Fraunhofer-diffraction, described by a computationally cheap Fast Fourier Transform (FFT), does not apply. In practice, this means that a suitable calculation of a single focal spot would take about an hour on a consumer PC.
I give a detailed discussion of the associated diffraction theory in appendix B with a set of numerical considerations, which forms the basis of this argument.

For this reason, I equipped CLAWS with a 16-bit CMOS camera, receiving a high resolution image of the focal spot via a low aberration imaging system, as described in the last section. Using these images, I was able to establish a procedure to estimate the on-shot peak intensity. This process involves several steps:

1. Calculate the maximum relative power density from the image
2. Estimate the impact of the target chamber evacuation process on the focus
3. Estimate the impact of the corresponding on-shot aberrations
4. Multiply the relative power density with the laser peak power

I will now discuss each of these steps in detail.

5.4.1. Determination of the Maximum Relative Power Density

The relative power density is the measure on the portion of the total power that acts on a surface element and is thus given in units of inverse area. Following the conventions of the field, I will give the values in units of $[1/\mathrm{cm}^2]$. Multiplying it with the laser power directly yields the intensity. In order to obtain this property from the recorded focal spot image, one has to subtract the background, normalize the data to the energy in the image and apply a size calibration.
For the latter, I mounted a needle on the target manipulator, applied some backlighting and brought it into focus on the camera chip. I then moved the needle by a predefined amount and observed the movement on the camera chip. By taking into account the uncertainties of the motorized stages and the feature detection in the image, I obtained a spatial calibration value of (86 ± 1) nm per pixel, which leads to an error of approx. 2.5 % when calculating areas from the image. One should note that the spatial calibration value does not correspond to the optical resolution, which is of the order of microns.
The background subtraction is shown in Fig. 5.37 for two different focal spots. First, I determined the center of the focal spot by the means of a simple center of gravity. Afterwards, I cropped the image to an area of roughly 120 µm squared, centered around the focal spot (see Fig. 5.37, left column). In order to determine the background, I defined an annular area around the center with the outer diameter corresponding to the cropped image size and the inner diameter chosen such that no signal of the focal spot was included. In fact, I determined the inner diameter by iterating this process several times to verify this, visually taking small-scale features as reference (see Fig. 5.37 on the right).
Afterwards, I fitted a 2D-polynomial to the data in the annular region. For this, I chose the sum of the first 15 Zernike polynomials. The fitted background is shown in the central column of Fig. 5.37. Finally I subtracted this fit from the original data. It is easy to tell qualitatively if the fit worked well when viewing the result in logarithmic scale (see Fig. 5.37, right column).

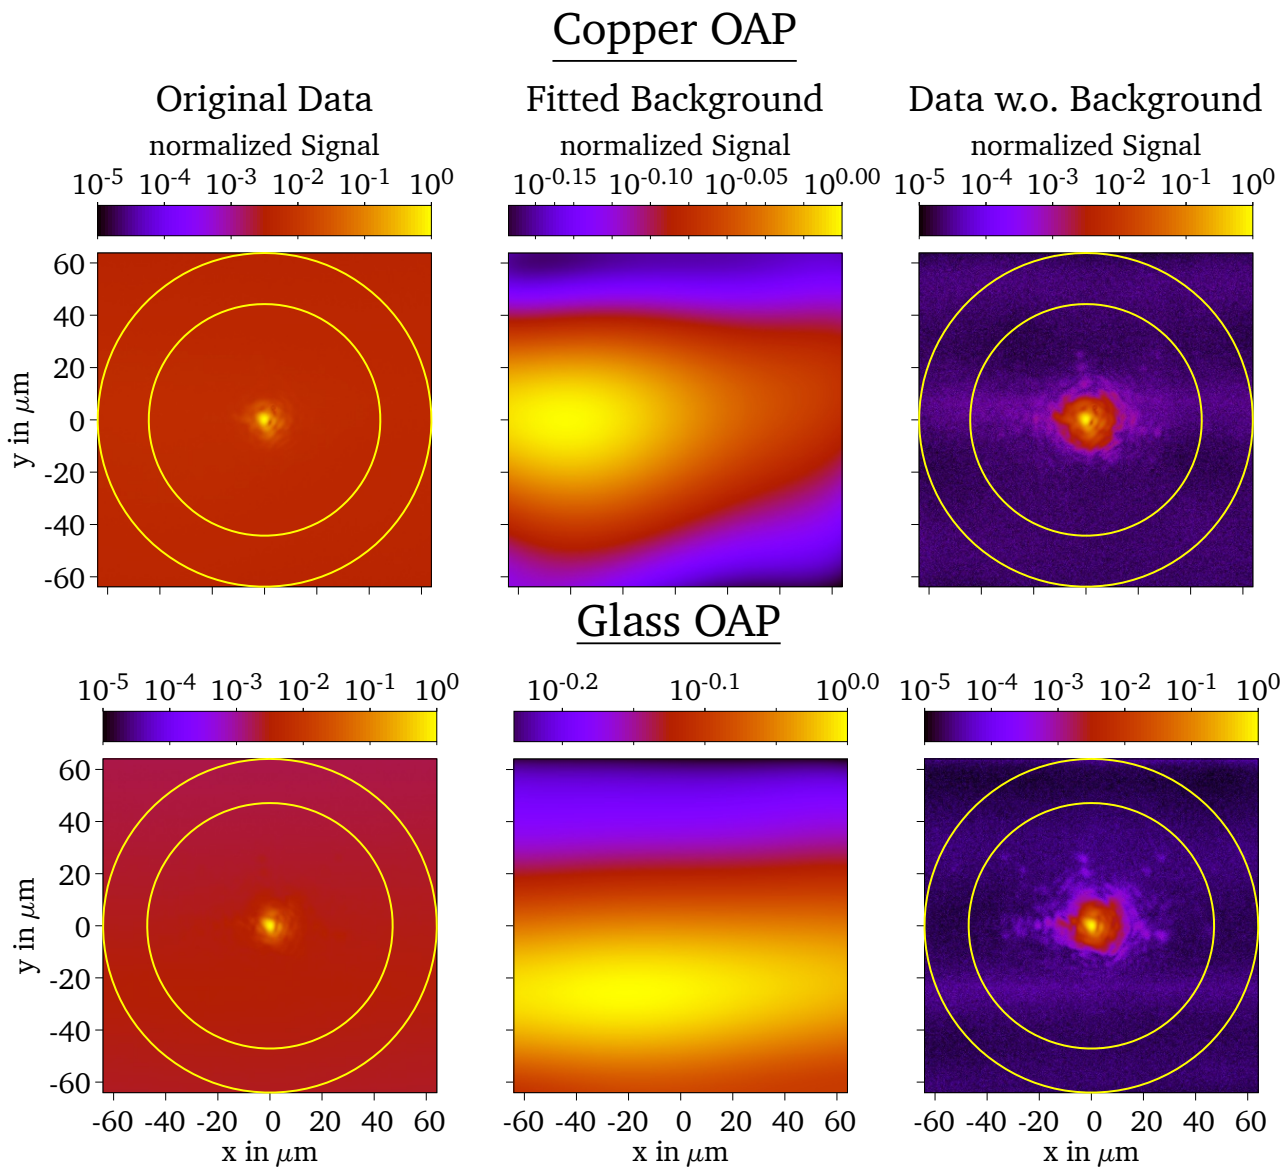

Figure 5.37.: Images of the background subtraction procedure of optimized PHELIX focal spot, recorded using CLAWS, for a copper OAP (top) and a glass OAP (bottom). The left column shows the raw image data, where the yellow circles indicate the annular region where the background is fitted using the first 15 Zernike polynomials. The result is shown in the central column, while the right column shows the image with the background subtracted. All data is color coded in logarithmic scale.

However, this should also be verified quantitatively. In order to do so, I calculated the encircled energy over the radius around the center of gravity of the focal spot. The result for the focuses which are shown in Fig. 5.37 is plotted in Fig. 5.38. In the case of a good background subtraction, the curves asymptotically approach 100% without overshooting or featuring residual slope in either direction. This is easily visible when the curve is shown using linear axes. A logarithmic scaling of the radius, however, is useful to reveal differences close to the center of the focus, where a steeper slope generally corresponds to a tighter focal spot.

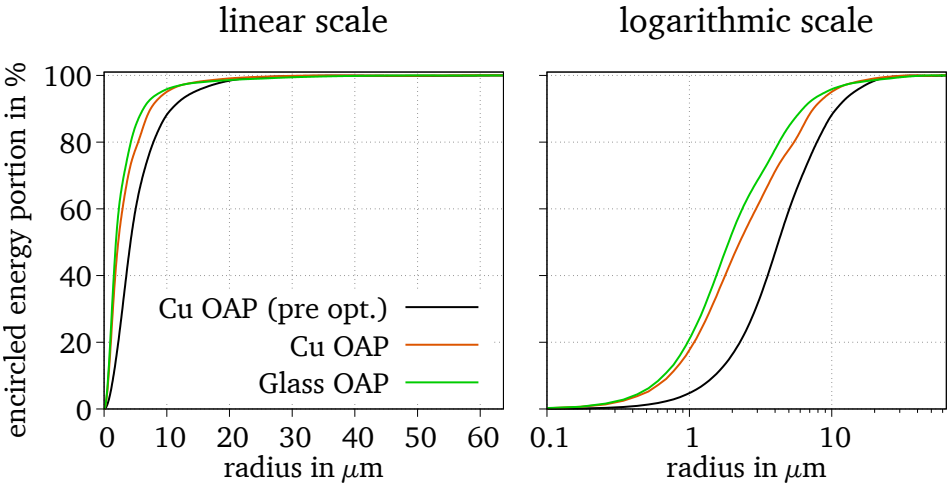

Figure 5.38.: Radially integrated energy portion of the focuses shown in Fig. 5.37 and the uncorrected focus of a copper OAP with linear x-scale (left) and logarithmic x-scale (right).

In the last step, I normalized the image to its integrated energy and take the maximum. The uncertainty introduced by different background fit parameters while checking the validity using the encircled energy approach, was 0.5 %. This concludes the determination of the maximum relative power density. For the two focuses shown in Fig. 5.37 I therefore obtained $(6.80 \pm 0.17)\ 1/\mathrm{cm}^2$ for the focus using the copper OAP and $(8.00 \pm 0.20)\ 1/\mathrm{cm}^2$ for the focus using the glass OAP.

5.4.2. Estimation of the Impact of the Evacuation Process

While an intensity can already be calculated at this point, the estimation only holds true while no other influences on the focal spot come into play. However, this is usually the case.
As I discussed earlier, the calibration routine currently only works while the target chamber, the compressor and the vacuum vessel of PTAS are vented. We are also aware that the atmospheric pressure does introduce deformations into the optical setup once the system is evacuated, leading to the loss of the calibration reference.
While we have no way to absolutely measure these differences, we know that the WF deformations caused by this are composed of low order aberrations only. By making the assumption that only defocusing and astigmatism are present here, one can compensate them very well using the degrees of freedom of the focusing OAP (comp. section 5.2.3).
However, the accuracy of this approach is limited as the only observable is the shape of the focal spot itself through the standard focal diagnostic and the optimum is determined subjectively. In practice, we found that we were able to find the optimal WF offset to compensate for the pump-down aberrations to a precision of 0.2 λ PtV of defocusing and 0.3 λ PtV of astigmatism in the Zernike sense. If both coefficients were at their maximum, this corresponds to a Strehl ratio of 0.81, given that the other aberrations were successfully eliminated by the calibration process. While this clearly indicates an error bar reaching from a Strehl ratio of 0.81 to 1.0, the expected value will be located somewhere in between.

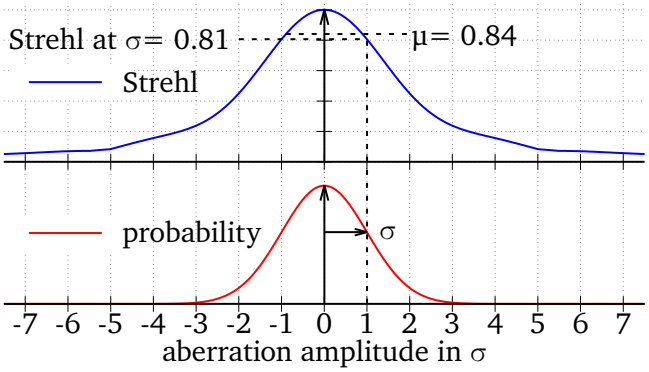

Figure 5.39.: Illustration of the impact of a standard distribution of an aberration amplitude (bottom) on the expected value of the Strehl ratio (top). One standard deviation corresponds to 0.2 λ PtV of defocusing and 0.3 λ PtV of astigmatism in the Zernike sense. The Strehl ratio was calculated numerically for a circular flat-top beam.

In order to calculate the expected value, one has to weight the Strehl ratio with the probability over the aberration amplitude (see Fig. 5.39) and integrate. For the Strehl ratio, I numerically calculated focal spots with defocusing and astigmatism in a constant ratio, given by the alignment precision mentioned above. I assumed the probability distribution to be Gaussian with a standard deviation equal to the point where the Strehl ratio hits 0.81 in order to match the observed precision.
Using this approach, I found that the expected Strehl ratio using this alignment scheme with its current limitations was $0.84^{+0.16}_{-0.03}$. Removing these uncertainties will therefore, on average, raise the intensity by $\frac{0.16}{0.84} = 19\,\%$.

5.4.3. Estimation of the Impact of On-Shot Aberrations

The next step, one has to consider the deviations of the WF from the calibration difference on shot. In this section, I will refer to these as ”on-shot aberrations”, while one is not to confuse these with the ”shot aberrations” as described in section 3.4. The latter one refers to the aberrations introduced due to the change of the refractive index in the amplifiers due to the pumping process, but these are mostly compensated for by adaptive optics. The on-shot aberrations on the other are composed of residuals of the AO loop and dynamic aberrations which cannot be predicted beforehand, e.g. caused by air turbulence.
Generally, these on-shot aberrations can be measured using PTAS, delivering an individual WF deviation from the optimum for each shot. Some examples are shown in Fig. 5.40.

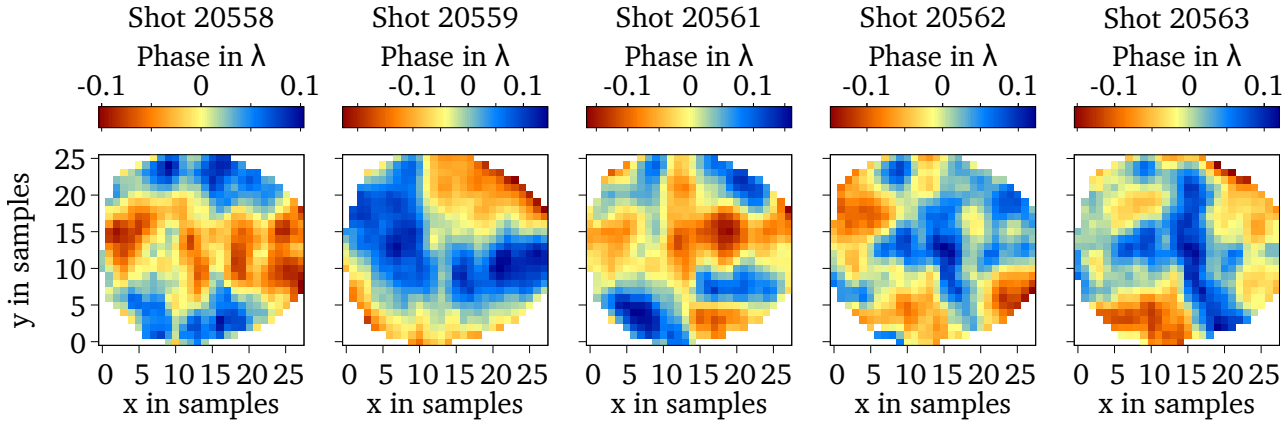

Figure 5.40.: Examples of on-shot aberrations, recorded at PTAS, over the course of one experiment day. The plots are labeled with the corresponding ID from the PHELIX shot data base.

For a perfectly calibrated WFS, these aberrations would directly translate into a Strehl ratio that one could apply to the maximum relative power density determined above. In reality however, some deviations are always present, including chromatic aberrations and STCs which require a significant amount of work to measure and an elaborate simulation routine to determine the intensity of the focal spot (see appendix B).
Therefore, an estimation is a viable trade off at this point. The idea is to calculate the Strehl ratio of the measured on-shot aberrations and multiply it with the other estimations. The quality of this estimation depends on the nature of the mostly unknown calibration errors and will, in general, get poorer the larger these errors are as the overall Strehl ratio will deviate from a quadratic dependency on the aberration stength with rising amplitude (comp. Fig. 5.39, top graph).
Once CLAWS can be operated in vacuum, the calibration quality should be sufficient such that this effect of aberration stacking is of no concern anymore. In this work, where the evacuation process leads to larger uncertainties, as discussed in the last section, this cannot be assumed in general. Even more, this issue cannot be fully quantified as the actual WF in the target chamber is unknown once the system is evacuated.
Therefore, the best way would be to perform a similar estimation as shown in the last section, where the uncertainties introduced by the evacuation process are varied in terms of defocusing and astigmatism and the expected value is calculated (comp. Fig. 5.39). In this case, one would not assume an aberration free WF as zero point, but the measured on-shot aberrations.
However, this leads to a complication: as the on-shot aberrations may include defocusing and astigmatism themselves, one cannot vary them using a unified amplitude multiplier, as done above. Instead, defocusing, oblique and normal astigmatism have to be varied separately, where every point in this three-dimensional space has to be weighted with a three-dimensional Gaussian likelihood. This tremendously increases the programming and computational efforts one has to do in order to estimate the on-shot intensity.
For this reason, I followed a statistical approach instead. For a number of shots, I added and subtracted one standard deviation of evacuation uncertainty (defocusing and astigmatism, as described above) to the on-shot aberrations and calculated the Strehl ratio. When comparing the results with the Strehl ratio of the original on-shot aberration, I observed that the average reduction of the peak intensity due to the uncertainty was 15 %. At the same time, most data points were included in a bidirectional error range of 15 %.
Here, one has to remember that this is based on statistical assumptions. These error bars include most of the data points, but do not reflect the specific properties of individual shots, which in part perform significantly better.

Conclusion

In this chapter, I described several major aspects of my work, building on the knowlege of the prior chapters. I started by giving a detailed description of the design considerations behind PTAS, which acts as the sensor of the AO loop. Furthermore, I described the technical realization in the laser hall of PHELIX and provided standalone test results. After that, I continued by discussing the calibration concept using the prototype CLAWS. Finally, I analyzed the optimized focal spots using the full AO loop and provided estimations for the on-shot intensity that can be expected using this setup. I will provide the experimental proof of an increased on-shot intensity in the next chapter.

6. Experimental Proof of the On-Shot Intensity Gain

In the last section, I discussed the design of the post-compressor Adaptive Optics (AO) loop that I built at Petawatt High-Energy Laser for heavy Ion eXperiments (PHELIX) in great detail. I also provided test results of both the sensor (section 5.2.6) and the calibration routine (section 5.3.3). However, I was only able to test the latter in air as the prototype was not designed for in-vacuum operation, which led to a large uncertainty of the estimated on-shot intensity gain (see section 5.4). This makes it clear that a full proof-of-principle of the AO loop must include an experimental verification of the on-shot intensity gain. In this chapter, I will describe an experimental campaign which I planned and conducted with our group in summer 2021, which aimed to demonstrate this intensity gain. However, the complexity of a potential absolute intensity verification, which I will briefly describe at the end of this chapter, made this method unsuitable for the time frame we had. This is the reason why we came up with a much simpler way to qualitatively demonstrate the on-shot intensity gain: measuring of the transmission of the laser pulses through foil targets ov various thicknesses.
In the following section, I will explain the concept of our experiment and describe the expected outcome. After that, I will briefly explain the setup in the target chamber and describe how we conducted the measurements. In section 6.3, I will show and discuss the results we obtained in this campaign.
Finally, I will conclude this chapter with a brief description of a potential method to provide an absolute proof of the reached peak intensity level based on field ionization of ambient gas, which can be investigated in a future work.

6.1. Concept and Expectations

In this experiment, we observed the transmission of the laser pulse through the target as we expected a higher transmission for higher intensities at the same energy and pulse duration. Before moving on to the description of the actual experiment though, I will first describe the inspiration for our approach and give some basic thoughts on the expected outcome. This will help to interpret the experimental data in order to draw a conclusion.

6.1.1. Concept

In the core, our experiment relies on the Hole Boring (HB) mechanism of a laser inside a plasma. Hornung et al. derived an expression for the temporal evolution of the HB velocity $v(t)$ of a laser pulse of the intensity $I(t)$ inside a plasma with an exponential density profile $n_e(x) \propto e^{L_c x}$ with the scale length L_c [83]:

$$v(t) = \frac{2L_c\sqrt{I(t)}}{\int_{-\infty}^{t}\sqrt{I(\tau)}d\tau + 2L_c\sqrt{\frac{c\,n_c M_i}{Z\cos(\theta)}}}. \tag{6.1}$$

Here, c is the speed of light in vacuum, n_c is the critical density at the laser WaveLength (WL), M_i the ion mass, Z the charge state and θ the angle of incidence of the laser pulse. This formula is very useful for laser-matter interaction as the pre-plasma, which forms due to laser light arriving prior to the main pulse, is approximately exponential due to the thermal distribution of particle velocities.

The authors coarsely divide the interaction into three qualitative phases: the pre-plasma expansion, where the pressure of the laser is not sufficient to act against the plasma pressure, the HB phase, where the light is intense enough to push into the plasma, and the plasma-pressure domination, where the dropping laser intensity is not sufficient to overcome the raising plasma density (see Fig. 6.1).
According to equation (6.1), more intense laser pulses will initially feature a higher HB velocity, which means that the laser penetrates deeper into the target. If the target thickness is limited, the critical surface will break through at some point. This is not included in the model, but has to be determined manually. From here, the rest of the laser light will be transmitted. Therefore, a higher intensity means faster breakthrough and thus more transmitted light. This is schematically illustrated in Fig. 6.2.

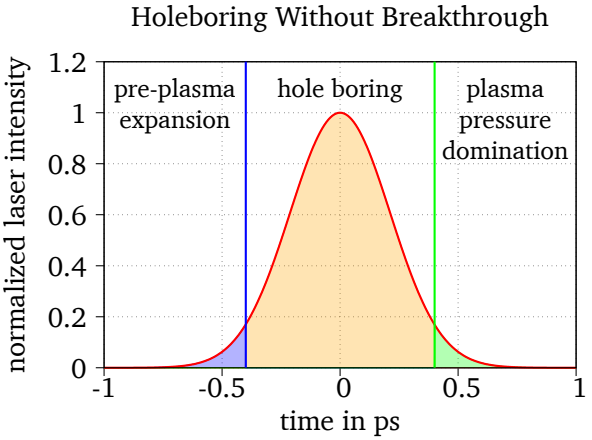

Figure 6.1.: Simplified laser penetration scheme according to Hornung et al. [83]. The graph shows the laser intensity over time, where the different interaction phases are indicated as shaded areas, labeled accordingly.

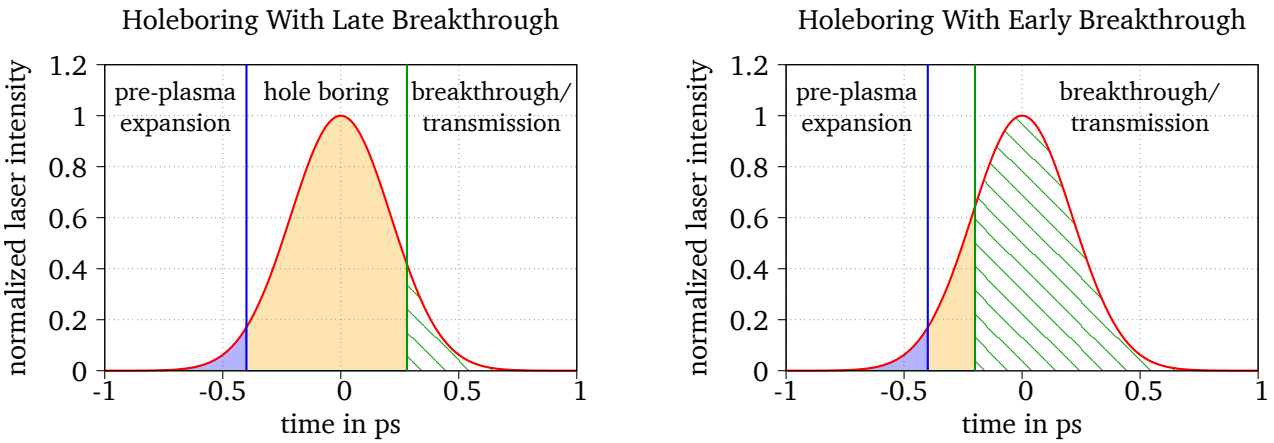

Figure 6.2.: Two scenarios where the laser breaks through the target before entering the plasma pressure domination phase as indicated in Fig. 6.1. The ruled area after the breakthrough indicates the portion of the pulse which is transmitted through the target.

This mechanism can be used to prove that a larger intensity is reached on the target. This will be purely qualitative as the interaction is too complex to derive a quantitative model for targets well beyond the skin depth of the laser.
Nonetheless, we used polystyrene foil targets and observed the transmitted light for shots with and without AO. Moreover, we scanned the target thickness to make sure to catch the relevant area, i.e. thicknesses above which the targets already have expanded to a fully undercritical plasma when the main pulse arrives and below thicknesses where the target stays completely opaque for the duration of the interaction. I will provide more details on this in section 6.2. Before that, however, I will introduce an expectation on the outcome in the next section.

6.1.2. Expectations

How exactly the transmission of laser light through foil targets works depends on a multitude of parameters. While the description is quite simple for very thin targets (comparable to the skin depth of the laser field) due to relativistic transparency in the volume [84], the interaction becomes increasingly complex for thicker targets and higher laser energies. For a truly realistic depiction of this process, a full three-dimensional consideration would be necessary, including self-focusing, filamentation and other mechanisms. The common way to do so would be to run a full simulation of the interaction, e.g. using a Magneto Hydro Dynamic (MHD) model or a Particle-In-Cell (PIC) method, which is a huge effort. And even here, the meaningfulness of the findings would be limited as the interaction is extremely sensitive to the plasma conditions prior to the shot and the distribution of the light on the target itself. Both of these properties are hard to estimate.
In this section, I will use equation 6.1 as a tool for some informed discussions instead. As mentioned above, this cannot serve to form realistic, quantitative expectations. The goal is rather to find a rough, qualitative prediction of the outcome of the experiment. One should keep this in mind over the course of the argumentation.
For this simplified model, I assume that the plasma of the pre-heated target features a fully exponential density distribution normal to the initial target surface (see red line in Fig. 6.3). This is realistic for the pre-plasma on the target, but will naturally fail at some point close to or inside the unheated target due to the limited thickness of the foil prior to the interaction. I will treat this with density- and thickness limits which depend on the initial target thickness and which I describe in more detail below.
Considering a fixed scenario, I assume a plasma scale-length L_c of one micrometer. This strongly depends on the intensity of the laser light that arrives prior to the main pulse and is hard to predict precisely, but is motivated from experience in the group [83] and the improved temporal contrast due to improvements in the fs FrontEnd (FE) of the laser [41]. Furthermore, I set the point where the plasma density equals the critical density n_c for the laser light, above which the laser cannot propagate anymore, to three micrometers in front of the initial target surface. Similarly to L_c, this fully depends on the temporal characteristics of the laser prior to the shot and is based on published MHD simulation results [85], which are based on the contrast of PHELIX. I used this plasma condition for all calculations in the following.

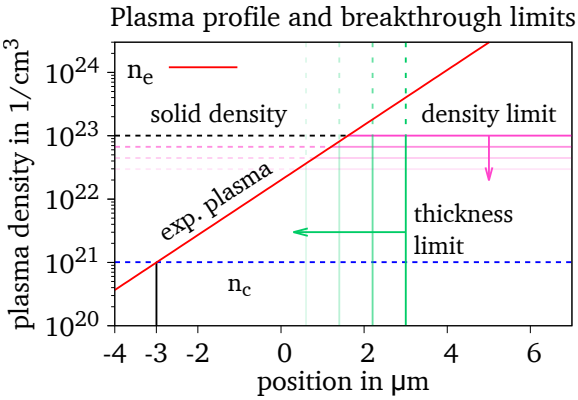

Figure 6.3.: Depiction of the plasma profile of this discussion. The pink and green lines show examples for the density and thickness limits, respectively, which serve as breakthrough thresholds.

As mentioned before, the exponential density distribution becomes unrealistic at some point inside the target. To account for this, I introduced two limits:
The first one is the foil thickness. Naturally, the plasma density will drop behind the rear side of the foil, which leads to the assumption that the laser breaks through the target once this point is reached. Some examples for this limit are shown as green, vertical lines in Fig. 6.3.
The second limit accounts for the maximum density of the target. Neglecting compression, the largest density is the solid density of the undisturbed target, which will be fixed for thick foils. Thinner foils, however, may expand at this location, which leads to a reduced peak density. Here again, I used the

findings from published MHD simulations [85] and assumed a clipped power model which matches the simulated target peak density well beyond the critical density (see Fig. 6.4). Here, the peak density $n_{\max}$ is given via the target thickness D in micrometers by

$$n_{max}(D) = \begin{cases} p(D) & \text{if} \quad p(D) < \text{solid density} \\ \text{solid density} & \text{if} \quad p(D) \geq \text{solid density} \end{cases}$$
$$\text{with} \qquad p(D) = 1.49 \times 10^{24} \cdot D^{3.16}.$$

This formula poses a second limit for the breakthrough, which I have depicted schematically using pink, horizontal lines in in Fig. 6.3. Again, if the surrounding density at the HB position exceeds this limit, I assume that the laser broke through the target at this instant.
Commonly, either the thickness limit or the density limit is reached first. In a realistic scenario, the breakthrough cannot be pinpointed to discrete events like these, but considering both cases will help to gain some understanding on what happens for thin and thick targets, respectively.
From this point, I performed a simple 2D simulation. I evaluated equation (6.1) for a grid of points over a Gaussian focal spot for a fully ionized CH target ($M_i = 13$, $Z = 7$) at an interaction angle of 20 degrees, reflecting the conditions in the real experiment.

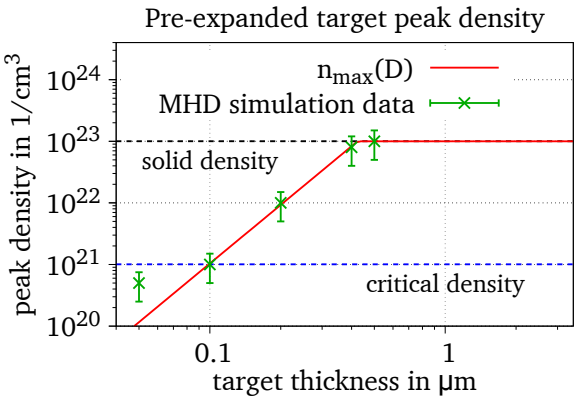

Figure 6.4.: Clipped power model of the maximum target density over the target thickness.

To do so, I assumed a laser pulse with a Gaussian temporal profil with a Full Width Half Maximum (FWHM) of 500 fs, where I varied the peak intensity to emulate the effect of AO. The integration within equation (6.1) starts at the point where the laser reaches 1 % of the peak intensity.
From the resulting set of HB velocity curves, I subtracted a speed of 3 nm/fs to account for the expansion velocity of the pre-plasma, which corresponds to a plasma temperature of about 7 keV (see Fig. 6.5).

Figure 6.5.: Examples of the HB velocity curves of a Gaussian laser pulse with a temporal FWHM of 500 fs and a peak intensity of 10^{21} W/cm^2. The thin curves show the HB velocities, depending on the radial position from the center of the Gaussian focal spot, where lower curves match regions of lower intensity.

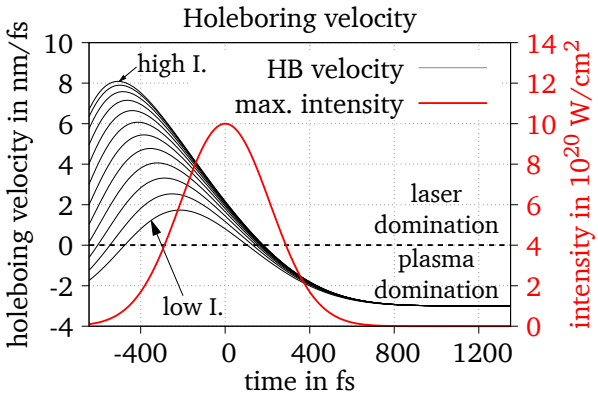

From here, a simple integration yields the hole-boring position over time for each sample of the 2D grid. The result is shown in Fig. 6.6. Obviously, the most intense regions of the laser pulse penetrate further into the plasma than the regions of lower intensity. One should note that in a real three-dimensional scenario, the gradient of the HB "wall" will affect the interaction, which is not considered in this simulation.

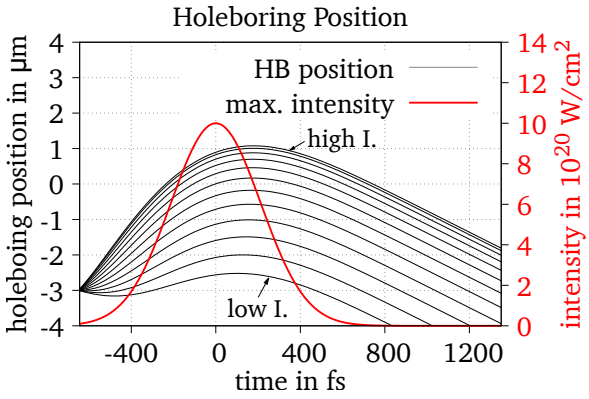

Figure 6.6.: Examples of the HB position curves for the 10^{21} W/cm^2 scenario. Again, the thin curves show the HB positions over the radius of the Gaussian focal spot, where lower curves match regions of lower intensity.

For each sample of the 2D grid, the density limit and the thickness limit are checked for different target thicknesses. If the sample crosses one of these limits, the transmitted energy portion is determined according to the time of this event. Integrating over all samples, the result is a curve that gives the total amount of laser energy that has been transmitted through the target.

In Fig. 6.7, the results for three realistic Intensities of the experiment are shown. The intensity is color-coded, while the transmission according to the density limit (pink horizontal lines in Fig. 6.3) is shown as dashed lines and the transmission according to the thickness limit (green vertical lines in Fig. 6.3) as dotted-dashed lines.

According to this model, the thinnest targets already lie below the critical density, which means that all energy is transmitted (compare Fig. 6.4). This is naturally not reflected by the thickness-limited model, as a significant part of the pulse duration is needed until the plasma is pushed to the rear side of the target. In this case, the density-limited model is therefore more realistic.

After that, the transmission drops continuously with increasing target thickness as the rays take longer and longer to reach the peak density up to the point where the maximum density is not reached anymore, which equals zero transmission. This scenario is, however, not realistic either, as the plasma density does not continue to grow exponentially with increasing depth in reality.

Applying the thickness-limit condition, where the pulse physically breaks through the rear side of the target, is an approximation which reflects this fact. According to the simulation, some transmission persists for thicker targets than for the density breakthrough, especially for higher intensities.

Of course, one does not expect a discrete transition from one case to the other, but rather a continuous fading from the density limited model towards the thickness-limited model with increasing target thickness. Furthermore, a full transmission, like predicted by the density-limited model, is not realistic as well, as even under critical plasma absorbs energy from the laser. A full transmission can therefore only happen if no target is present in the first place.

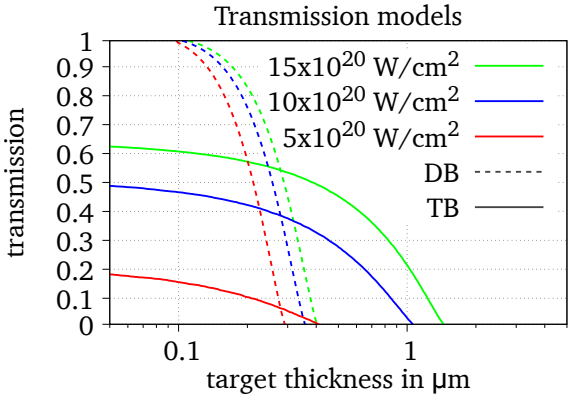

Figure 6.7.: The transmittance of polystyrene foil targets according to the simulation for three different intensities, applying the density breakthrough limit ("DB", dashed lines) and the thickness breakthrough limit ("TB", solid lines).

Despite all these deviations from reality, the simulation clearly indicates that a higher laser intensity leads to more transmission through the target and thus confirms the initial expectation. Furthermore, it predicts that the transmission features a maximum spread between the intensities at targets of a certain thickness, as thin targets simply transmit the pulse and thick targets fully absorb the energy.

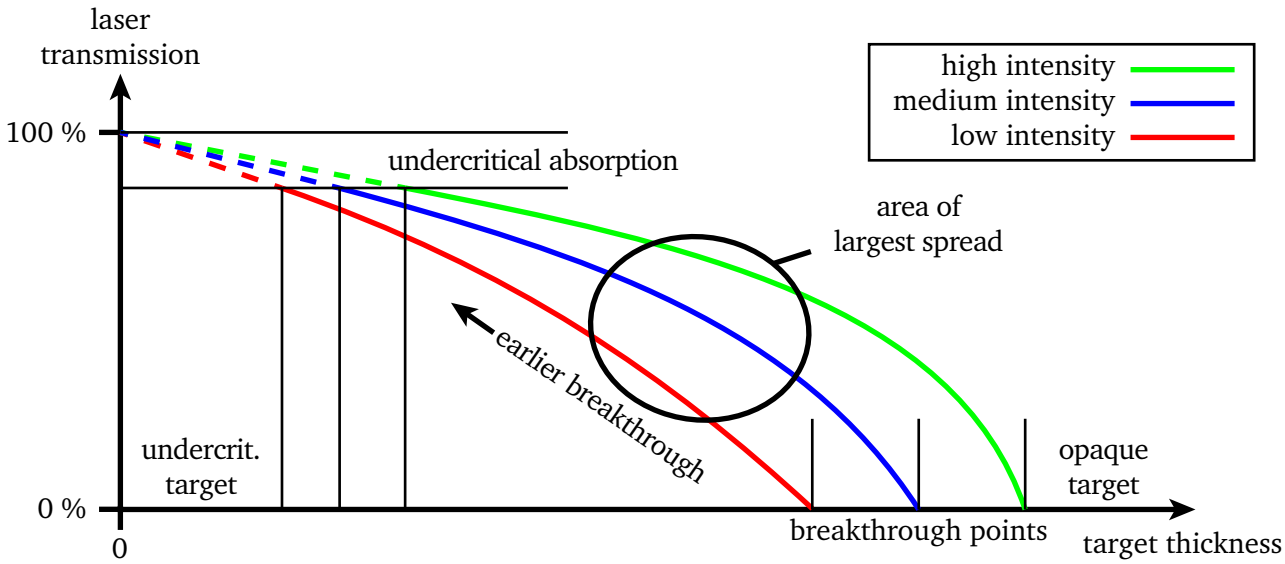

Figure 6.8.: Qualitative expectations on the experimental data, based on the two-dimensional HB simulation, shown for three different laser intensities. The figure is described in the text.

According to these investigations, a clear qualitative expectation towards the results can be formulated, which I have illustrated in Fig. 6.8. In the following I formulate the expectations from thick to thin targets:
In the limit of very thick targets, no light will be transmitted as the laser is not able to penetrate through the foil. If the target thickness is decreased, the pulse will break through at a certain point, from which the transmission rises quickly. This happens earlier for higher laser intensities.
While the transmission rises with decreasing thickness, there will be an area of largest spread between the transmissions for different intensities. Below this regime, the targets begin to expand over their entire volume before the interaction starts, lowering the maximum plasma density and thus leading to even earlier breakthroughs.
This continues, until the plasma density becomes fully undercritical. An unknown amount of the laser light gets absorbed in the undercritical plasma below this point, where the absorption approaches zero with vanishing target thickness.

6.2. Experimental Setup

We performed an experiment which we expected to produce data according to the established expectations (see Fig. 6.8). This meant shooting foil targets with a wide range of thicknesses while keeping the laser pulse energy constant, but varying the intensity using AO. Specifically, according to the goal of a qualitative proof of the intensity gain of the new AO loop, varying the intensity meant shooting with the uncorrected and corrected beam, respectively. While doing so, the transmitted energy had to be measured on each shot.

The experimental setup that we used to measure the laser transmission is shown schematically in Fig. 6.9.
We used two different f = 400 mm, 45° Off-Axis Parabolic mirrors (OAPs) for beam focusing, where we swapped between a well-polished glass OAP and a copper OAP, which features larger surface errors (see Fig. 6.10). The Angle Of Incidence (AOI) with the polystyrene foil targets was 20°.
We recorded the transmitted energy using an energy-meter, which received light that was reflected by a simple glass pane and collected by a Fresnel lens. We cross calibrated this sensor to another energy-meter (QE95 by Gentec) which is permanently located at Main Amplifier Sensor (MAS) to track the laser performance. This calibration was simply done by doing low-energy shots without target.
Furthermore, we measured the transmission through the reflector and the Fresnel lens by placing the head directly behind the target plane once, recording a shot, and relocating it to the indicated position afterwards.
Using these reference measurements, we were able to tell how much energy hit the target from the measurement at MAS and how much was transmitted from the measurement of the energy-meter in the target chamber.

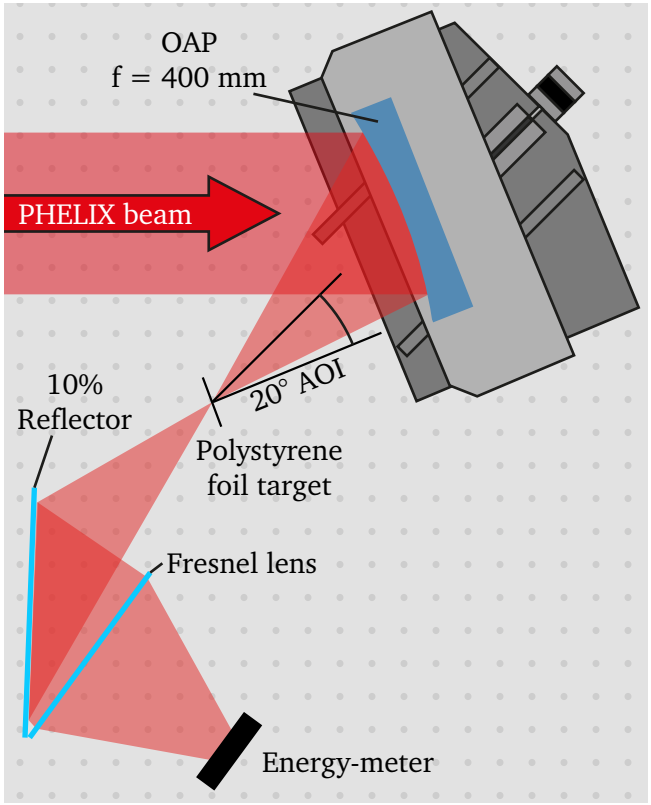

Figure 6.9.: The experimental setup for the target transmission measurements.

As we intended to scan the area of the gradually increasing transmission in Fig. 6.8, we chose a set of foil targets based on past experiences, spanning almost two orders of magnitude in thickness. The foil targets were produced at the target laboratory at the Technical University of Darmstadt using a spin-coating procedure. The produced foils were then fixed onto polycarbonate target frames which held them in place while insulating the foil from the target chamber. This acted as a countermeasure against the ElectroMagnetic Pulse (EMP) caused by the charge separation on shot, which could potentially interfere with measurements.
The details on the batches, also provided by the laboratory, are shown in Tab. 6.1. Each batch contained multiple targets, which we used to repeat shots with similar conditions.

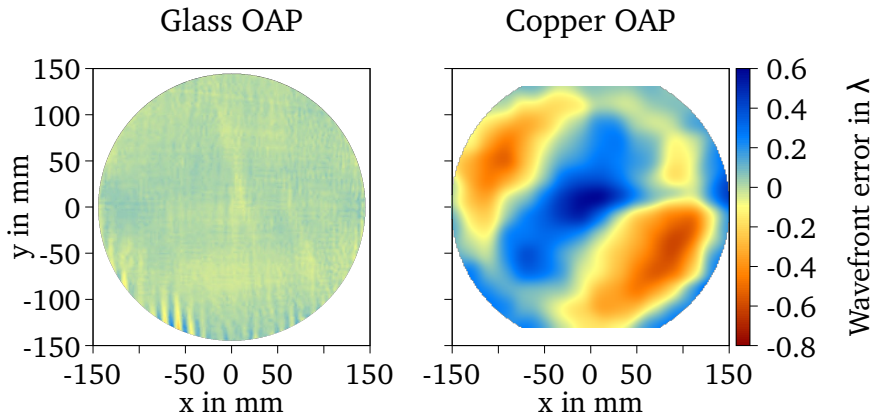

Figure 6.10.: The WaveFront (WF) errors introduced by the used glass OAP (left) and the copper OAP (right). Tip/tilt, defocusing and astigmatism were subtracted from the datasets. Even though the data of the copper OAP were given with a low spatial resolution, the large amplitude of the surface imperfections immediately becomes apparent. The datasets were delivered by the manufacturers.

Thickness in nm	Uncertainty in nm	Thickness measurement method	Nr. of targets
66	±1	Ellipsometry on wafer	16
130	±2	Ellipsometry on wafer	16
277	±1	Ellipsometry on wafer	16
718	±15	White light interferometry per target	12
765	±4	Ellipsometry on wafer	17
1.195	±6	Ellipsometry on wafer	10
1.735	±5	Ellipsometry on wafer	11
2.975 to 2.993	±15	White light interferometry per target	12
3.129 to 3.206	±15	White light interferometry per target	6

Table 6.1.: Batch details on the polystyrene foil targets used in this experiment. The targets and data were provided by the target laboratory at the Technical University of Darmstadt.

6.3. Results and Discussion

Within the experimental campaign, we performed 30 successful shots, with an average of three shots per day, without counting calibration shots. For each shot, we recorded the laser energy, using the cross-calibrated energy-meter at MAS, the laser transmission through the target and the on-shot WF at Petawatt Target Area Sensor (PTAS).
The shot energy on the target was 139.9 J ±9.0 J on average with an individual error of 7.1 %, where the maximum energy was 169 J ±12.0 J and the minimum energy was 119.7 J ±8.5 J. From here, I estimated the reached on-shot intensity for every shot using the procedure I described in section 5.4. This involved the specific energy of each shot, the on-shot WF and the focal spot of the calibration

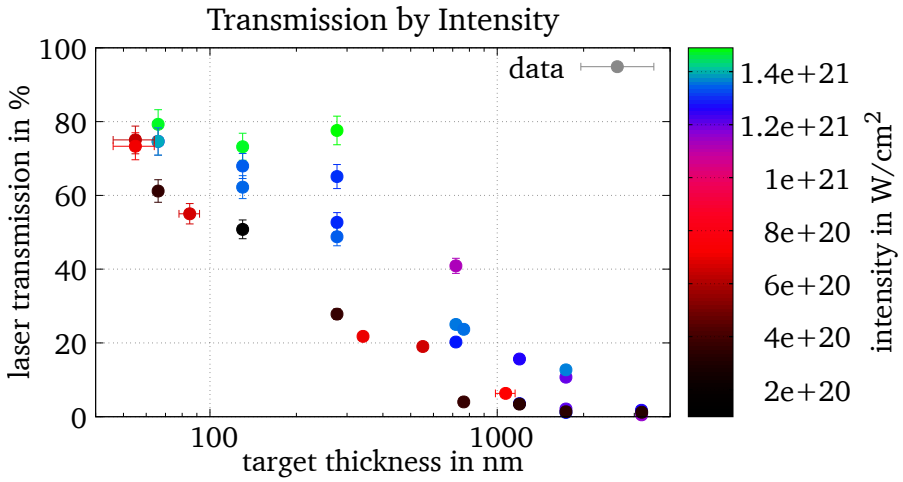

Figure 6.11.: The laser transmission over the thickness of the polystyrene foil targets, where the laser peak intensity is color coded. The same data is shown in Fig. 6.12, where the color code indicates the associated setup.

routine, which was recorded with the Far Field (FF) HR camera of Calibrator for Lasersystems with Absolute Wavefront Sensing (CLAWS). All errors accumulated to a bidirectional error of 23 % for shots where all data was ideally recorded (13 of 30 shots) and 34 % for the rest of the shots where the WF measurement was biased due to an accidental loss of reference by moving the focusing OAP in the target chamber. This was an error in the routine, which can be avoided in future experiments.
Due to the fluctuations in the pulse energy and the on-shot aberrations, the experiment did not provide a scan following a fixed intensity as shown in Fig. 6.8, but a scattered field of data points. In order to compare the findings to the expectations, I plotted the shots as points over the target thickness, where the y-axis corresponds to the laser transmission and the intensity is encoded in the color of the points. The result is shown in Fig. 6.11.
Obviously, the data points follow the qualitative expectations nicely. We did not observe any transmission for targets thicker than two micrometers, while reducing the thickness yielded higher transmission values. Clearly, the targets started transmitting earlier for higher intensities, as predicted by the model. We observed the largest spread of the data points at a target thickness of about 300 nm, ranging from 28 % up to 78 %. The spread is clearly reduced for both much thinner and thicker targets, matching the expectations as well.
For the thinnest target, we observed a strongly reduced spread in the transmission, where the shots with the smallest and largest intensity almost overlap. This may be interpreted as an effect of the absorption inside an undercritical target (as suggested by the MHD simulations [85]), which would indicate an absorption of about 20 %. However, this cannot be verified from the recorded data set.
At this point, it should be noted that the choice of displaying the intensity by color prohibits plotting the error of the intensity as well. As the error can be as large as 34 % for over half of the shots, this may be the largest source for the scattering of the data points in Fig. 6.11.

As the experiment matches the expectations, i.e. the transmission scales with the laser intensity, we were now able to verify that the AO loop was accountable for the increase of the on-shot intensity. To illustrate this, I plotted the same data set again in Fig. 6.12. This time, however I used the color code to indicate the setup used for each individual shot, rather than the intensity. The first setup used the

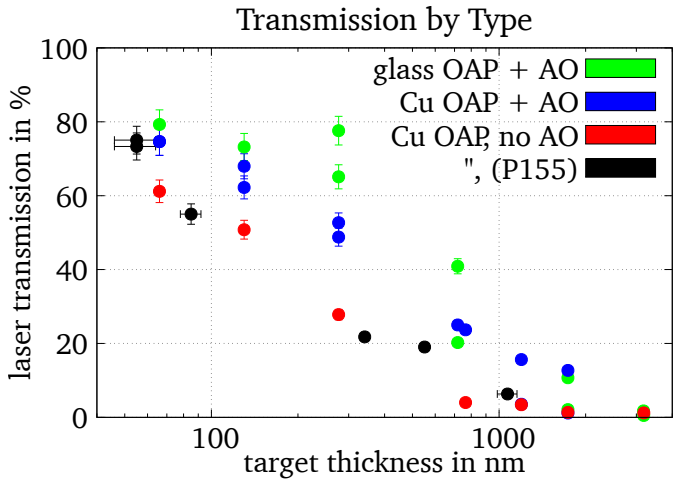

Figure 6.12.: The laser transmission over the thickness of the polystyrene foil targets, where the color of each data point indicates the corresponding setup. The black data set was recorded in another experimental campaign (P155), but used the same setup as for the red data points.

fine glass OAP (see Fig. 6.10) together with active beam-correction (green dots). The second one used a copper OAP with larger surface errors and beam-correction as well (blue), while the last one used the copper OAP without beam correction (red), which corresponds to the standard configuration at PHELIX and shows a good agreement with a previous conducted experiment (P155), given by the black dots.
The result is clear: the shots with beam correction using PTAS and the Deformable Mirror (DM) feature a consistently larger transmission than the shots without beam correction, which was the main goal of this experimental campaign.
Furthermore, the shots with the glass OAP perform better than the ones with the copper OAP, which means that the aberrations of the copper OAP cannot be fully corrected using this AO setup, underlining the advantages of more expensive glass OAPs.
As a last remark, it should be pointed out that the shots in the P155 experiment featured a consistently higher intensity than the ones recorded in this experiment (compare the red and black points in Fig. 6.11 and Fig. 6.12), but do not feature a significant spread over the transmission. This indicates that the intensity was rather the same in both campaigns and that at least one of the two experiments included a systematic error in the estimation of the intensity. The fact that PTAS did not exist in P155 makes the intensity estimations of this experiment more trustworthy as more data was available.
From this data, I conclude that the intensity on target is indeed increased by the adaptive optics loop and, despite only shown qualitatively, that a gain of factor three remains highly plausible.

6.4. Possible Quantitative Intensity Proof using Optical Tunnel Ionization

In our experiment, we were able to show the intensity increase due to the AO loop in a qualitative manner. However, the absolute values for the peak intensity remain estimations with a high level of uncertainty. In order to gain confidence, a qualitative experiment could be performed in the future. Such an experiment needs to rely on a physical process that is directly sensitive to the intensity in the regime of High Intensity (HI) lasers, namely intensities starting from $10^{20}\,\mathrm{W/cm^2}$.

One experiment that has been demonstrated before uses tunnel ionization of atoms [86]. In tunnel ionization, the laser field tilts the potential of the nucleus of an atom, basically turning one side of the potential wall into a potential hump. Electrons, which are trapped inside the potential of the nucleus, can now tunnel through this barrier. This ionization process gets more likely the smaller this hump is, eventually transitioning to field ionization, where the hump drops below the energy level of the electron. The height of the potential barrier in turn gets smaller the stronger the laser field is. This means that with rising laser intensity, electrons that are more tightly bound to the nucleus may be stripped from the atom. This process is highly nonlinear, where the likelihood of ionizing the atom up to a certain energy level abruptly rises once a corresponding laser field strength is reached. This makes this process suitable to prove that a given intensity is reached.

Instead of shooting at a solid target, the idea of the experiment is now to shoot into a thin ambient gas with pressures of 1E-6 mbar or less. This gas should be a noble gas as other gases may form molecules. The combination of low pressure and single-atom gas means that concessional ionization by free electrons can be neglected in the process.

Two electrodes around the focal spot now separate the ions from the electrons (see Fig. 6.13). The ions are accelerated towards the anode and gain a given energy from the electric field, depending on their ionization state. After they leave the electric field through a wire mesh, the ions need to travel a certain distance. The Time Of Flight (TOF) up to a detector therefore can be directly linked to the ionization state of the according ions.

The signal on the detector will consist of a series of peaks, where each peak scales with the abundance of an ion species. With a suitable theory, this spectrum can be linked back to the conditions in the focal spot of the laser. By comparing the height of the peaks relative to each other, this setup can even provide a continuous measurement of the peak intensity.

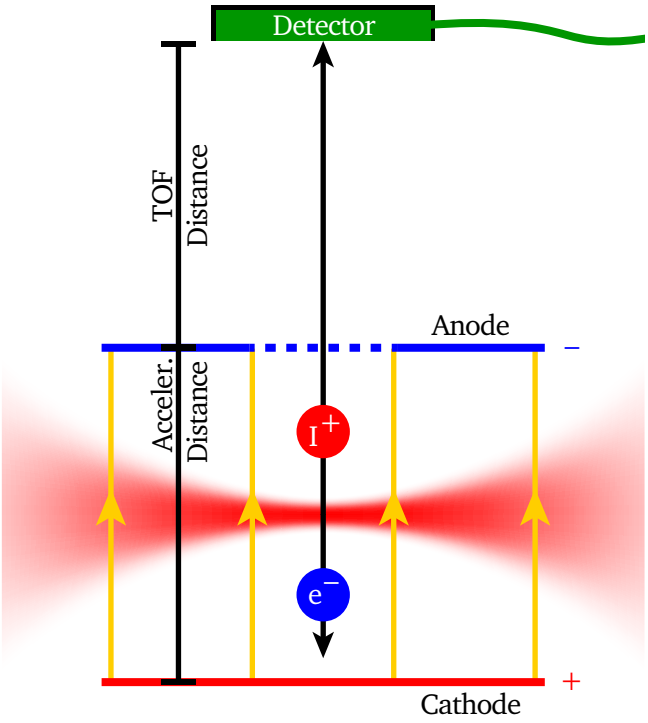

Figure 6.13.: The basic experimental setup for the quantitative intensity measurement using the tunnel ionization of the ambient gas by the laser field.

The evaluation of such an experiment is challenging, as the theory not only relies on the ionization process, but also on the three-dimensional geometry of the experiment. Research on this topic is still ongoing in the field.

If one succeeds, however, a quantitative measurement of the intensity is possible. For lasers with a high repetition rate, this setup can even serve as an alignment aid for the compression of the laser pulse or as additional feedback for AO.

7. Summary, Outlook and Acknowledgments

The motivation for this work was to obtain a better control of the on-shot focal spot at the Petawatt High-Energy Laser for heavy Ion eXperiments (PHELIX) laser. This would greatly increase the scientific value of experimental results and make higher peak intensities accessible. Such a beam control can be achieved by using Adaptive Optics (AO), which is capable of compensating WaveFront (WF) aberrations in the beam.
In the first chapters of this thesis, I have discussed that any focal spot optimization in Chirped Pulse Amplification (CPA) laser systems has to be done after the pulse compression stage in order to avoid transforming the WF aberrations into Spatio Temporal Couplings (STCs) in the compressor. Therefore, this work dealt with the design, implementation and testing of a post-compressor AO loop at PHELIX.

A major effort went into the construction of the WaveFront Sensor (WFS) system of the loop, capturing the leakage light of the last turning mirror prior to the target chamber. This system, named Petawatt Target Area Sensor (PTAS), basically consists of a drastic beam size reduction by a factor of 64, imaging the 28 cm aperture of the mechanical Deformable Mirror (DM) at the compressor output onto a Shack-Hartmann Sensor (SHS). Furthermore, a multi-stage energy attenuation system was added to provide measurements in the low-energy alignment mode as well as in shot mode.
The main challenge of this was the tight space constraint, as PTAS needed to fit within a footprint of only 2×1 meters, while still capturing the full aperture of the beam. I overcame this by implementing a compact 8:1 Off-Axis Parabolic mirror (OAP) telescope with an F-number of 1.43, which reduced the large beam to a manageable size. The small F number indicates the sensitivity of the setup to misalignment, which is why I developed a dedicated alignment scheme, which is based on overlapping the focal spots of both OAPs using an intermediate focal diagnostic. I performed a comprehensive analytical investigation on this scheme and showed that it yields sufficient precision in the context of High Intensity (HI) laser beam transport in general. The only exception is the defocusing of the telescope, which has to be fine-tuned using an external reference.
The tests of PTAS showed that the imaging featured a spatial resolution of less than 2 mm on the DM, which is five times better than the resolution of the SHS. With this setup, I was able to record the on-shot aberrations of PHELIX and control the mechanical DM in a closed loop. The characterization of the mirror matched the specifications of the manufacturer.

The next step was to implement a calibration routine. This was necessary as the WF at PTAS does correlate but not correspond to the WF in front of the target due to aberrations in the beam transport of PTAS and the focusing OAP in the chamber.
To do so, I designed and built a prototype of a mobile WFS called Calibrator for Lasersystems with Absolute Wavefront Sensing (CLAWS). This device is placed behind the focal spot in the target chamber and thus can measure the relevant WF. CLAWS is self-referencing by generating a spherical wave, centered on the target, by extreme spatial filtering, using a 1 µm pinhole. After recording the reference, the WF is optimized using CLAWS and the DM, which yields the best achievable WF, acting as reference for PTAS. The corresponding focal spot is recorded by a 16-bit camera on CLAWS.

The resulting intensity was over three times higher than the one prior to the optimization, delivering up to $1.4{\cdot}10^{21}\,\mathrm{W/cm^2}$ for a 100 J pulse with a duration of 500 fs. The main source of uncertainty in this estimation is the fact that CLAWS could only operate under atmosphere, while the subsequent evacuation process of the vacuum system adds aberrations to the beam, which could not be measured precisely. Expressed in a Strehl ratio, this manifests in a factor of $0.84^{+0.16}_{-0.03}$. A vacuum-compatible version of CLAWS could therefore additionally increase the intensity by up to $\frac{0.16}{0.84} = 19\,\%$.

In order to prove the on-shot intensity gain, we performed an experimental campaign. Here, we observed the transmission of uncorrected and corrected shots through polystyrene foil targets of different thicknesses. The expectation was that higher intensities would yield a higher transmission due to a faster Hole Boring (HB) velocity. The expectation was confirmed with significance, which delivered the qualitative proof of the increased intensity.
Overall, I was able to successfully optimize the focal spot, systematically reaching intensities over $10^{21}\,\mathrm{W/cm^2}$ for shots of 100 J and a duration of 500 fs. This is a factor of three more than the average shot prior to the implementation of the loop, and provides more stability over the course of an experiment day.

Outlook

As the implementation of the AO loop, which I presented in this work, is only a proof of principle, a lot of work has to be done until the focal spot optimization can be available for community experiments. This is both due to bugs in the current state of the setup and the operation of the system, which currently requires a high level of expertise. For the future of PTAS and CLAWS, three milestones can be set, which I will briefly describe in the following.

1. Integration of PTAS into the control system: The PHELIX laser is operated via the PHELIX Control System (PCS), which is an elaborated software network, spanning the computers of the whole facility. The PCS provides a Graphical User Interface (GUI) for the operating crew, where sensor outputs are displayed and axes of motion can be controlled. A subsystem of PHELIX is only applicable for daily operation if it is well integrated in the PCS with a simple GUI that requires little knowledge of the particularities of the hardware.
Naturally, PTAS needs to be integrated here as well in order to be accessible to the operating crew and the experimentalists alike. From the current state of PTAS, there are still several hurdles to take:

1. **Hardware debugging:** As this is the first version of PTAS, there are several things that need to be made more reliable. This includes for example the optical endswitches of the Mach–Zehnder Interferometer (MZI) shutter stage, which need to be replaced with more reliable, mechanical ones, and the implementation of a permanent back-propagating alignment beam.

2. **Optical debugging:** The optical path of PTAS has to be optimized too. As an example, the continuous attenuation is currently not homogeneous over the Near Field (NF) on-shot, which points out a certain level of depolarization.

3. **Development of an operational concept:** Currently, the alignment of the OAP telescope, despite being conceptually simple, requires handling nine axes of motion at once. For daily operation, a layer of abstraction and a suitable GUI has to be implemented.

4. **Software integration:** Some additional hurdles come with the integration of individual components into the PCS. For example, no interface for the vacuum camera has been programmed by now. Also, the proximity of the WFS to the target chamber leads to outliers on the SHS due to X-ray induced hotpixels. A robust outlier removal algorithm has to be implemented.

A full list of open tasks has been provided to the PHELIX crew.

2. Focal spot optimization for community experiments: A more challenging process will be the preparation of the WF control and focal spot optimization for community experiments, as the calibration routine is rather complex. The following points need to be fulfilled:

1. **Implementation of a vacuum-compatible version of CLAWS:** The evacuation process introduces aberrations into the system, as some components deform under atmospheric pressure. As a consequence, the calibration using CLAWS should be done in vacuum. Therefore, the axes of motion on CLAWS need to be motorized for remote control and the cameras need to be vacuum compatible.

2. **Debugging of CLAWS:** The latest prototype of CLAWS still features some design flaws. As an example, the backlighter LED for the pinhole gets partially reflected on the Far Field (FF) LR camera due to the use of a beamsplitter cube, which makes finding the pinhole challenging.

3. **Adapt CLAWS to be experiment friendly:** The prototype is currently supported by four large pillar posts, which interfere with the experimental setup. A different mounting scheme, e.g. hanging from the chamber ceiling, needs to be implemented.

4. **Increase DM stability:** The software library of the DM manufacturer was specifically developed for PHELIX and is unstable. Crashes occur on a regular basis. For daily operation, this issue needs to be resolved.

5. **Partial automation of the calibration procedure:** The complexity of the calibration procedure poses a challenge to the operating crew. Some effort has to be taken in order to simplify it. For example, the operation of CLAWS could be added to the PCS as well and the recording of a reference WF at PTAS could be done via a one-click function.

Again, a comprehensive list of the open tasks and suggestions has been provided to the PHELIX crew.

3. Re-implementation of PTAS prior to TB2: The current location of PTAS is not ideal as some experimental setups require placing optical components in the turning boxes adjacent to the target chamber, which may interfere with the WF measurement. With sufficient budget, however, a large vacuum tank could be added to the compressor vessel and PTAS could be relocated there. This would decouple the sensor entirely from the experiment and make the AO loop a standard option for all experiments.

Quantitative intensity proof Aside from this stepwise integration of the AO loop into the system, a quantitative proof of the on-shot intensity still has to be delivered. As discussed before, an option for this is to investigate the tunnel ionization of a thin noble gas by the laser field at the focal spot. While the necessary expertise does not exist yet at PHELIX, such an experiment has been done in the group of J. Schreiber at the Ludwig-Maximilians University in Munich. A collaboration would certainly help to perform such a measurement.

Implementation of an integrated AO concept Currently, multiple components at PHELIX work together to keep aberrations under control. This includes two AO loops, two moving lenses and a bending mirror for the compensation of astigmatism. Therefore, it may be useful to implement an integrated control scheme which incorporates all these elements into a single, large control logic. This could potentially make the shot preparation a lot faster, while simplifying the integration into the PCS at the same time.

Acknowledgments

Writing a thesis like this does only work with the right support, be it technical guidance, organizational crisis management, mental healthcare or just someone to have a coffee or beer with. Luckily, I had a great bunch of superheroes at my side that could provide all the essentials and more! Even though it is impossible for me to list them all, I would like to give some credit to the most important ones.

Victory-Vincent Besides being my doctoral advisor, Vincent earned his place in this list with a wide range of things. Having the big picture, he showed me the directions for my work, was a consultant in any affair and always had some time to listen to my problems or help me out with the structure of my publications. Besides that, he greatly contributed to my health (running in the woods) and, from time to time, did the opposite by organizing excellent wine tastings. Thank you for the awesome time!

Ultra-Udo Udo had been my technical supervisor since my master's thesis and gave me the tools to do what had to be done pretty early on. Throughout this work, he stayed at my disposal and I greatly enjoyed doing sessions of brainstorming with him. This is especially true for everything software-related - maybe due to the fact that he is the core-developer of Wavefront Optics Measurement and Beam Analysis Tool (WOMBAT), our beloved laboratory all-rounder and adaptive optics control software. This project fertilized nerdy conversations like nothing else.

Buster-Bernhard Moving from software to hardware: a huge thank-you to Bernhard, who is accountable for the target area at PHELIX. He gave me an incredible amount of support with the implementations of PTAS. This consumed at least a month of his time - not talking about all my issues with motion, hardware, vacuum and the fact that he purchased the DM! I highly appreciate his help and will remember our "heavy metal" sessions with a smile on my face, cramps in my forearms and oil on my disposable gloves.

Dashing-Diana She is the one that provided crisis management whenever I was overwhelmed by the organizational complexity of GSI. This happened more frequent than I would like to admit (organizing business travels is arcane arts to me), so I had to bother her a lot. But even after more than three years, she remained patient and kept her sense of humor. This is among the most remarkable things I have seen over the course of my work.

Yeti-Yannik Also known as "The Expert" and a multitude of other names, we shared an office and he acted as a most welcome antidote to intellectual thoughts. We also shared a growing aversion to using words for communication, enjoyed wrecking the nerves of the department with a grotesque collection of music and still survived two experimental campaigns together. Thanks for sticking with me!

Jedi-Johorn As a lone ranger in a world full of laser engineers, Johannes stayed in touch with the mysterious reign of plasma physics. He was an excellent match for Yannik and me and did stay until the end. Without him, I would have had a hard time to come up with chapter 6.

Spider-Simon He joined the chaos-crew at GSI for his master's thesis and made a glorious transition into his PhD. He blends in so effortless that I can barely remember a time without him. Only his

magnificent beard-growth and the amount of hand-crafted memes give me some sens of time here ... I am glad to have called you one of my colleges.

Destiny-Dustin My one and only graduate student ... seriously, there are no others and there never will be. He survived my idealism up to this day and keeps simulating thousands of donuts an hour. Still, he found time to do some proofreading of this thesis.

The rest of the PHELIX squad Astro-Albert, Crocodile-Christian, Destruction-Dennis, Doomsday-Dirk, Laser-Lila, Slayer-Stefan, Super-Sabine and Zapping-Zsuzsi: you are incredible! One cannot wish to meet a crew that is organized, dynamic, helpful and funny at the same time, but I still received all of that while working with you. Without this team, neither the setup of PTAS, nor the experimental campaign would have been possible. If you ever have a free position, don't forget about me. ;-)

My social environment I am incapable of coming up with enough superhero-names for my friends and family out there, and likewise, it is impossible to express how thankful I am to have so many intact and healthy relationships. A special shout-out goes to **Blaster-Biggi**, which I am happy to call my girlfriend. With her love, humor and bright rational thinking, she easily made up for the bad days of stress and doubts and helped me navigate the emotional waves of doing a doctorate. I owe you so much!

8. List of Abbreviations and Bibliography

Abbreviations

AO	Adaptive Optics
AOI	Angle Of Incidence
AOM	Acousto-Optic Modulator
CACAO	Compute And Control for Adaptive Optics
CLAWS	Calibrator for Lasersystems with Absolute Wavefront Sensing
CU	Control Unit
CW	Continuous Wave
COS	COmpressor Sensor
CPA	Chirped Pulse Amplification
CPU	Central Processing Unit
DM	Deformable Mirror
EM	ElectroMagnetic
EMP	ElectroMagnetic Pulse
FE	FrontEnd
FF	Far Field
FFT	Fast Fourier Transform
FPGA	Field Programmable Gate Array
FWHM	Full Width Half Maximum
GPGPU	General Purpose Computation on Graphics Processing Unit
GSI	GSI Helmholtzzentrum für Schwerionenforschung GmbH
GUI	Graphical User Interface
HB	Hole Boring
HI	High Intensity
HR	High Reflectivity
LCA	Liquid Crystal Array
MAS	Main Amplifier Sensor
MEMS	MicroElectroMechanical System
MHD	Magneto Hydro Dynamic
MLA	Micro Lens Array
MOPA	Master Oscillator Power Amplifier
MWLSI	Multi-Wave Lateral Shearing Interferometer
MZI	Mach–Zehnder Interferometer
NA	Numerical Aperture
Nd:Glass	Neodymium-doped Glass
Nd:YAG	Neodymium-doped Yttrium Aluminum Garnet
NF	Near Field
OAM	Orbital Angular Momentum

OAP	Off-Axis Parabolic mirror
OPD	Optical Path Difference
PA	PreAmplifier
PCS	PHELIX Control System
PH	PinHole
PHELIX	Petawatt High-Energy Laser for heavy Ion eXperiments
PIC	Particle-In-Cell
PR	Phase Retrieval
PTA	Petawatt Target Area
PTAS	Petawatt Target Area Sensor
PtV	Peak-to-Valley
PWFS	Pyramid WaveFront Sensor
QWLSI	Quadri-Wave Lateral Shearing Interferometer
RMS	Root Mean Square
SHS	Shack-Hartmann Sensor
SLM	Spatial Light Modulator
SNR	Signal-to-Noise Ratio
STC	Spatio Temporal Coupling
SVD	Singular Value Decomposition
Ti:Sapphire	Titanium-doped Sapphire
TOF	Time Of Flight
UNILAC	Universal Linear Accelerator
uOPA	ultrafast Optical Parametric Amplifier
WF	WaveFront
WFS	WaveFront Sensor
WL	WaveLength
WOMBAT	Wavefront Optics Measurement and Beam Analysis Tool
Yb:YAG	Ytterbium-doped Yttrium Aluminum Garnet
YAG	Yttrium Aluminum Garnet

Bibliography

[1] J. B. Ohland *et al.*, Alignment procedure for off-axis-parabolic telescopes in the context of high-intensity laser beam transport, Opt. Express, **29**(21):34378–34393 (2021), doi:10.1364/OE.439658.

[2] J. B. Ohland *et al.*, Ultra-compact post-compressor on-shot wavefront measurement for beam correction at PHELIX, High Power Laser Science and Engineering, **10**:e18 (2022), doi:10.1017/hpl.2022.9.

[3] T. MAIMAN, Stimulated Optical Radiation in Ruby, Nature, **187**(4736):493–494 (1960).

[4] C. N. Danson *et al.*, Petawatt and exawatt class lasers worldwide, High Power Laser Science and Engineering, **7**:e54 (2019), doi:10.1017/hpl.2019.36.

[5] D. Kraus *et al.*, Nanosecond formation of diamond and lonsdaleite by shock compression of graphite, Nature communications, **7**(1):1–6 (2016).

[6] J. Hornung *et al.*, Enhancement of the laser-driven proton source at PHELIX, High Power Laser Science and Engineering, **8**:e24 (2020), doi:10.1017/hpl.2020.23.

[7] J. M. Cole *et al.*, High-resolution μCT of a mouse embryo using a compact laser-driven X-ray betatron source, Proceedings of the National Academy of Sciences, **115**(25):6335–6340 (2018), doi:10.1073/pnas.1802314115.

[8] N. Guler *et al.*, Neutron imaging with the short-pulse laser driven neutron source at the Trident laser facility, Journal of Applied Physics, **120**(15):154901 (2016), doi:10.1063/1.4964248.

[9] A. B. Zylstra *et al.*, Burning plasma achieved in inertial fusion, Nature, **601**(7894):542–548 (2022).

[10] J.-P. Zou, B. Wattellier and R. Tyson, Adaptive optics for high-peak-power lasers–an optical adaptive closed-loop used for high-energy short-pulse laser facilities: laser wave-front correction and focal-spot shaping, in "Topics in adaptive optics", chapter 5, 95–116, IntechOpen (2012).

[11] D. Strickland and G. Mourou, Compression of amplified chirped optical pulses, Optics Communications, **55**(6):447–449 (1985), doi:10.1016/0030-4018(85)90151-8.

[12] S. Zhang *et al.*, Analysis of wavefront effects for large-aperture tiled-grating compressor, Laser and Particle Beams, **36**:84–91 (2018), doi:10.1017/S0263034617000878.

[13] J. W. Yoon *et al.*, Realization of laser intensity over 10^{23} W/cm^2, Optica, **8**(5):630–635 (2021), doi:10.1364/OPTICA.420520.

[14] J. W. Hardy, Active optics: a new technology for the control of light., IEEE Proceedings, **66**:651–697 (1978).

[15] J. M. Beckers, Adaptive Optics for Astronomy: Principles, Performance, and Applications, Annual Review of Astronomy and Astrophysics, **31**(1):13–62 (1993), doi:10.1146/annurev.aa.31.090193.000305.

[16] H. W. Babcock, The possibility of compensating astronomical seeing, Publications of the Astronomical Society of the Pacific, **65**(386):229–236 (1953).

[17] A. Buffington *et al.*, Correction of atmospheric distortion with an image-sharpening telescope, J. Opt. Soc. Am., **67**(3):298–303 (1977), doi:10.1364/JOSA.67.000298.

[18] J. W. Hardy, Solar Imaging Experiment., Technical report, ITEK CORP LEXINGTON MASS OPTICAL SYSTEMS DIV (1980).

[19] E. Gendron *et al.*, Come-on-plus project: an upgrade of the come-on adaptive optics prototype system, in M. A. Ealey (editor), Active and Adaptive Optical Systems, volume 1542, 297 – 307, International Society for Optics and Photonics, SPIE (1991), doi:10.1117/12.48818.

[20] S. S. Olivier and C. E. Max, Laser Guide Star Adaptive Optics: Present and Future, Symposium - International Astronomical Union, **158**:283–292 (1994), doi:10.1017/S0074180900107739.

[21] O. Guyon, Extreme Adaptive Optics, Annual Review of Astronomy and Astrophysics, **56**(1):315–355 (2018), doi:10.1146/annurev-astro-081817-052000.

[22] F. Rigaut and B. Neichel, Multiconjugate Adaptive Optics for Astronomy, Annual Review of Astronomy and Astrophysics, **56**(1):277–314 (2018), doi:10.1146/annurev-astro-091916-055320.

[23] R. Tyson and B. Frazier, Principles of Adaptive Optics, CRC Press (2022), ISBN 9781000531343.

[24] R. K. Tyson, How adaptive optics may have won the Cold War, in T. George, M. S. Islam and A. K. Dutta (editors), Micro- and Nanotechnology Sensors, Systems, and Applications V, volume 8725, 184 – 190, International Society for Optics and Photonics, SPIE (2013), doi:10.1117/12.2010801.

[25] P. S. Salter and M. J. Booth, Adaptive optics in laser processing, Light: Science & Applications, **8**(1):1–16 (2019).

[26] C. Rodríguez and N. Ji, Adaptive optical microscopy for neurobiology, Current Opinion in Neurobiology, **50**:83–91 (2018), doi:10.1016/j.conb.2018.01.011.

[27] M. Pircher and R. J. Zawadzki, Review of adaptive optics OCT (AO-OCT): principles and applications for retinal imaging [Invited], Biomed. Opt. Express, **8**(5):2536–2562 (2017), doi:10.1364/BOE.8.002536.

[28] R. Wang *et al.*, Demonstration of horizontal free-space laser communication with the effect of the bandwidth of adaptive optics system, Optics Communications, **431**:167–173 (2019), doi:10.1016/j.optcom.2018.09.038.

[29] M. Chen *et al.*, Performance verification of adaptive optics for satellite-to-ground coherent optical communications at large zenith angle, Opt. Express, **26**(4):4230–4242 (2018), doi:10.1364/OE.26.004230.

[30] C. Cai *et al.*, Experimental demonstration of an underwater wireless optical link employing orbital angular momentum (OAM) modes with fast auto-alignment system, in 2019 Optical Fiber Communications Conference and Exhibition (OFC), 1–3, IEEE (2019).

[31] A. E. Willner and C. Liu, Perspective on using multiple orbital-angular-momentum beams for enhanced capacity in free-space optical communication links, Nanophotonics, **10**(1):225–233 (2021), doi:10.1515/nanoph-2020-0435.

[32] C. Brabetz, Wellenfront kontrolle am phelix-lasersystem mit adaptiver optik, Ph.D. thesis, Master's thesis, Technische Universität Darmstadt, Darmstadt (2010).

[33] O. Guyon *et al.*, The compute and control for adaptive optics (CACAO) real-time control software package, in D. Schmidt, L. Schreiber and L. Close (editors), Adaptive Optics Systems VI, Proceedings of SPIE - The International Society for Optical Engineering, SPIE (2018), ISBN 9781510619593, doi:10.1117/12.2314315.

[34] G. Golub and W. Kahan, Calculating the singular values and pseudo-inverse of a matrix, Journal of the Society for Industrial and Applied Mathematics, Series B: Numerical Analysis, **2**(2):205–224 (1965).

[35] P. Piatrou and L. Gilles, Robustness study of the pseudo open-loop controller for multiconjugate adaptive optics, Appl. Opt., **44**(6):1003–1010 (2005), doi:10.1364/AO.44.001003.

[36] L. E. Hargrove, R. L. Fork and M. A. Pollack, Locking of He–Ne Laser Modes Induced by Synchronous Intracavity Modulation, Applied Physics Letters, **5**(1):4–5 (1964), doi:10.1063/1.1754025.

[37] T. Brabec *et al.*, Kerr lens mode locking, Opt. Lett., **17**(18):1292–1294 (1992), doi:10.1364/OL.17.001292.

[38] R. Paschotta, article on 'Master Oscillator Power Amplifier' in the RP Photonics Encyclopedia, URL `https://www.rp-photonics.com/master_oscillator_power_amplifier.html`, accessed on 2022-04-29.

[39] V. Bagnoud *et al.*, Commissioning and early experiments of the PHELIX facility, Applied Physics B, **100**:137–150 (2010), doi:10.1007/s00340-009-3855-7.

[40] F. Wagner *et al.*, Temporal contrast control at the PHELIX petawatt laser facility by means of tunable sub-picosecond optical parametric amplification, Applied Physics B, **116**(2):429–435 (2014).

[41] Y. Zobus *et al.*, Reaching the Millijoule-Regime via Ultrafast Optical Parametric Amplification–An Alternative to First Stage Regenerative Amplification Stages?, in 2021 Conference on Lasers and Electro-Optics (CLEO), 1–2, IEEE (2021).

[42] P. Tournois, Acousto-optic programmable dispersive filter for adaptive compensation of group delay time dispersion in laser systems, Optics Communications, **140**(4):245–249 (1997), doi:10.1016/S0030-4018(97)00153-3.

[43] W. W. Simmons and R. O. Godwin, Nova laser fusion facility—design, engineering, and assembly overview, Nuclear Technology-Fusion, **4**(1):8–24 (1983).

[44] Z. Major and S. Franz, Laser-driven x-ray diagnostics for heavy-ion heatod dense plasmas, GSI Darmstadt Reports (2021), doi:10.15120/GSI-2020-01321.

[45] E. Treacy, Optical pulse compression with diffraction gratings, IEEE Journal of Quantum Electronics, **5**(9):454–458 (1969), doi:10.1109/JQE.1969.1076303.

[46] M. Born and E. Wolf, Principles of optics: Electromagnetic theory of propagation, interference and diffraction of light, 7th (expanded) edition, Fourth printing, Cambridge U. Press, Cambridge, UK (2006).

[47] J. Ye *et al.*, Modal wavefront estimation from its slopes by numerical orthogonal transformation method over general shaped aperture, Opt. Express, **23**(20):26208–26220 (2015), doi:10.1364/OE.23.026208.

[48] J. C. Wyant and K. Creath, Basic wavefront aberration theory for optical metrology, Applied optics and optical engineering, **11**(part 2):28–39 (1992).

[49] V. N. Mahajan, Orthonormal aberration polynomials for anamorphic optical imaging systems with rectangular pupils, Appl. Opt., **49**(36):6924–6929 (2010), doi:10.1364/AO.49.006924.

[50] F. von Zernike, Beugungstheorie des Schneidenverfahrens und seiner verbesserten Form, der phasenkontrastmethode, Physica, **1**(7-12):689–704 (1934).

[51] R. J. Noll, Zernike polynomials and atmospheric turbulence, J. Opt. Soc. Am., **66**(3):207–211 (1976), doi:10.1364/JOSA.66.000207.

[52] R. V. Shack and K. Thompson, Influence of alignment errors of a telescope system on its aberration field, in Optical Alignment I, volume 251, 146–153, SPIE (1980).

[53] Z. Li *et al.*, Degradation of femtosecond petawatt laser beams: Spatio-temporal/spectral coupling induced by wavefront errors of compression gratings, Applied Physics Express, **10**(10):102702 (2017), doi:10.7567/apex.10.102702.

[54] V. Leroux *et al.*, Wavefront degradation of a 200 TW laser from heat-induced deformation of in-vacuum compressor gratings, Opt. Express, **26**(10):13061–13071 (2018), doi:10.1364/OE.26.013061.

[55] D. Papadopoulos *et al.*, First commissioning results of the Apollon laser on the 1 PW beam line, in 2019 Conference on Lasers and Electro-Optics (CLEO), 1–2, IEEE (2019).

[56] A. Jeandet *et al.*, Survey of spatio-temporal couplings throughout high-power ultrashort lasers, Opt. Express, **30**(3):3262–3288 (2022), doi:10.1364/OE.444564.

[57] H.-M. Heuck *et al.*, Chromatic aberration in petawatt-class lasers, Applied Physics B, **84**:421–428 (2006), doi:10.1007/s00340-006-2230-1.

[58] Z. Li and J. Kawanaka, Complex spatiotemporal coupling distortion pre-compensation with double-compressors for an ultra-intense femtosecond laser, Opt. Express, **27**(18):25172–25186 (2019), doi:10.1364/OE.27.025172.

[59] R. Heathcote *et al.*, Collimation effects on large CPA compressors, Applied Physics B, **116**:805–809 (2014), doi:10.1007/s00340-014-5765-6.

[60] G. Figueira *et al.*, Pulse front tilt control using non-collimated beams in a single pass grating compressor, Opt. Express, **28**(5):7678–7690 (2020), doi:10.1364/OE.28.007678.

[61] J. W. Yoon *et al.*, Realization of laser intensity over 10 23 W/cm 2, Optica, **8**(5):630–635 (2021).

[62] A. L. Rukosuev *et al.*, Adaptive optics system for real-time wavefront correction, Atmospheric and Oceanic Optics, **28**:381–386 (2015), doi:10.1134/S1024856015040119.

[63] M. Negro *et al.*, Fast stabilization of a high-energy ultrafast OPA with adaptive lenses, Scientific Reports, **8**:14317 (2018), doi:10.1038/s41598-018-32182-y.

[64] C. Dorrer *et al.*, Investigation of an apodized imaged Hartmann wavefront sensor, Appl. Opt., **57**(25):7266–7275 (2018), doi:10.1364/AO.57.007266.

[65] J. M. Soures *et al.*, OMEGA: A short-wavelength laser for fusion experiments, Technical report, Rochester Univ., NY (USA). Lab. for Laser Energetics (1983).

[66] L. J. Waxer *et al.*, The OMEGA EP High-Energy, Short-Pulse Laser System, in Conference on Lasers and Electro-Optics/Quantum Electronics and Laser Science Conference and Photonic Applications Systems Technologies, JThB1, Optica Publishing Group (2008).

[67] S. Akturk *et al.*, Spatio-temporal couplings in ultrashort laser pulses, Journal of Optics, **12**(9):093001 (2010), doi:10.1088/2040-8978/12/9/093001.

[68] J. Ohland *et al.*, WOMBAT: A modular software for wavefront sensing and optimization at PHELIX, in 8th Conference of the International Committee on Ultrahigh Intensity Lasers, ICUIL, Lindau, Germany (2018).

[69] V. N. Mahajan, Strehl ratio for primary aberrations in terms of their aberration variance, J. Opt. Soc. Am., **73**(6):860–861 (1983), doi:10.1364/JOSA.73.000860.

[70] M. Tacca *et al.*, Tuning of a high magnification compact parabolic telescope for centimeter-scale laser beams, Applied Optics, **55**(6):1275–1283 (2016).

[71] A. S. Householder, UNITARY TRIANGULARIZATION OF A NONSYMMETRIC MATRIX, J. Assoc. for Computing Machinery, **5**(4) (1958), doi:10.1145/320941.320947.

[72] G. B. Airy, On the diffraction of an object-glass with circular aperture, Transactions of the Cambridge Philosophical Society, **5**:283 (1835).

[73] A. Maréchal, Etude des influences conjuguées des aberrations et de la diffraction sur l'image d'un point, Ph.D. thesis, Faculté des Sciences des Paris (1947).

[74] L. Doyle, Implementation of an Adaptive Optics System for Focus Optimization at the Centre for Advanced Laser Applications, Master's thesis, Ludwig Maximilian University of Munich (2019).

[75] C. Guo *et al.*, A review of iterative phase retrieval for measurement and encryption, Optics and Lasers in Engineering, **89**:2–12 (2017), doi:10.1016/j.optlaseng.2016.03.021.

[76] G. Ju *et al.*, Feature-based phase retrieval wavefront sensing approach using machine learning, Opt. Express, **26**(24):31767–31783 (2018), doi:10.1364/OE.26.031767.

[77] J. White *et al.*, Real-time phase-retrieval and wavefront sensing enabled by an artificial neural network, Opt. Express, **29**(6):9283–9293 (2021), doi:10.1364/OE.419105.

[78] D. Posor, Voruntersuchung zur Verwendung von Strahlen mit Bahndrehimpuls als Wellenfrontsensorik, Project proposal, Technical University of Darmstadt (2021).

[79] J. B. Ohland *et al.*, A study on the effects and visibility of low-order aberrations on laser beams with orbital angular momentum, Applied Physics B, **125**:202 (2019), doi:10.1007/s00340-019-7317-6.

[80] D. Wang *et al.*, Analysis of diffraction wavefront in visible-light point-diffraction interferometer, Appl. Opt., **52**(31):7602–7608 (2013), doi:10.1364/AO.52.007602.

[81] G. Ding *et al.*, Effect of illumination parameters on the quality of wavefront diffracted by pinhole, in Y. Zhang and W. Gao (editors), 7th International Symposium on Advanced Optical Manufacturing and Testing Technologies: Optical Test and Measurement Technology and Equipment, volume 9282, 436 – 442, International Society for Optics and Photonics, SPIE (2014).

[82] F. Gao, Z. d. Jiang and B. Li, Diffraction wavefront analysis of point diffraction interferometer for measurement of aspherical surface, in Y. Zhang *et al.* (editors), 5th International Symposium on Advanced Optical Manufacturing and Testing Technologies: Optical Test and Measurement Technology and Equipment, volume 7656, 1412 – 1419, International Society for Optics and Photonics, SPIE (2010).

[83] J. Hornung *et al.*, Time-resolved study of holeboring in realistic experimental conditions, Nature Communications, **12**:6999 (2021), doi:10.1038/s41467-021-27363-9.

[84] D. Jung *et al.*, Scaling of ion energies in the relativistic-induced transparency regime, Laser and Particle Beams, **33**(4):695–703 (2015), doi:10.1017/S0263034615000828.

[85] V. Bagnoud *et al.*, Studying the Dynamics of Relativistic Laser-Plasma Interaction on Thin Foils by Means of Fourier-Transform Spectral Interferometry, Phys. Rev. Lett., **118**:255003 (2017), doi:10.1103/PhysRevLett.118.255003.

[86] Y. Akahane *et al.*, Characterization of wave-front corrected 100 TW, 10 Hz laser pulses with peak intensities greater than 10^{20} W/cm^2, Review of Scientific Instruments, **77**(2):023102 (2006), doi:10.1063/1.2166669.

[87] J. Hartmann, Bermerkungen über den Bau und die Justierung von Spektrographen, Zeitschrift für Instrumentenkunde, **20**:47 (1900).

[88] B. C. Platt and R. Shack, History and principles of Shack-Hartmann wavefront sensing, Journal of Refractive Surgery, **17**(5):S573–S577 (2001).

[89] W. H. Southwell, Wave-front estimation from wave-front slope measurements, JOSA, **70**(8):998–1006 (1980).

[90] M. Harker and P. O'Leary, Least squares surface reconstruction from measured gradient fields, in 2008 IEEE Conference on Computer Vision and Pattern Recognition, 1–7 (2008), doi:10.1109/CVPR.2008.4587414.

[91] M. E. Riley and M. A. Gusinow, Laser beam divergence utilizing a lateral shearing interferometer, Appl. Opt., **16**(10):2753–2756 (1977), doi:10.1364/AO.16.002753.

[92] J. Primot, Three-wave lateral shearing interferometer, Appl. Opt., **32**(31):6242–6249 (1993), doi:10.1364/AO.32.006242.

[93] S. Velghe *et al.*, Wave-front reconstruction from multidirectional phase derivatives generated by multilateral shearing interferometers, Opt. Lett., **30**(3):245–247 (2005), doi:10.1364/OL.30.000245.

[94] J. Primot and L. Sogno, Achromatic three-wave (or more) lateral shearing interferometer, J. Opt. Soc. Am. A, **12**(12):2679–2685 (1995), doi:10.1364/JOSAA.12.002679.

[95] J.-C. F. Chanteloup and M. Cohen, Compact high resolution four wave lateral shearing interferometer, in R. Geyl, D. Rimmer and L. Wang (editors), Optical Fabrication, Testing, and Metrology, volume 5252, 282 – 292, International Society for Optics and Photonics, SPIE (2004), doi:10.1117/12.513739.

[96] A. J. I. Toepler, Beobachtungen nach einer neuen optischen Methode: Ein Beitrag zur Experimentalphysik, 157, W. Engelmann (1906).

[97] R. Ragazzoni, Pupil plane wavefront sensing with an oscillating prism, Journal of Modern Optics, **43**(2):289–293 (1996), doi:10.1080/09500349608232742.

[98] R. A. Frazin, Efficient, nonlinear phase estimation with the nonmodulated pyramid wavefront sensor, J. Opt. Soc. Am. A, **35**(4):594–607 (2018), doi:10.1364/JOSAA.35.000594.

[99] A. Burvall *et al.*, Linearity of the pyramid wavefront sensor, Opt. Express, **14**(25):11925–11934 (2006), doi:10.1364/OE.14.011925.

[100] I. Shatokhina, V. Hutterer and R. Ramlau, Review on methods for wavefront reconstruction from pyramid wavefront sensor data, Journal of Astronomical Telescopes, Instruments, and Systems, **6**(1):1 – 39 (2020), doi:10.1117/1.JATIS.6.1.010901.

[101] L. H. Schatz *et al.*, Design of the MagAO-X Pyramid Wavefront Sensor (2018), doi:10.48550/ARXIV.1807.04316, URL `https://arxiv.org/abs/1807.04316`.

[102] J. Lozi *et al.*, Visible and Near-infrared Laboratory Demonstration of a Simplified Pyramid Wavefront Sensor, Publications of the Astronomical Society of the Pacific, **131**(998):044503 (2019), doi:10.1088/1538-3873/ab046a.

[103] R. W. Gerchberg, A practical algorithm for the determination of phase from image and diffraction plane pictures, Optik, **35**:237–246 (1972).

[104] D. Meschede, Gerthsen Physik, Springer, 24th edition (2010), ISBN 978-3-642-12893-6, doi:10.1007/978-3-642-12894-3.

[105] O. Ikeda and T. Sato, Comparison of deformability between multilayered deformable mirrors with a monomorph or a bimorph actuator, Appl. Opt., **25**(24):4591–4597 (1986), doi:10.1364/AO.25.004591.

[106] P.-Y. Madec, Overview of deformable mirror technologies for adaptive optics and astronomy, in B. L. Ellerbroek, E. Marchetti and J.-P. Véran (editors), Adaptive Optics Systems III, volume 8447, 22 – 39, International Society for Optics and Photonics, SPIE (2012), doi:10.1117/12.924892.

[107] J.-C. Sinquin, J.-M. Lurçon and C. Guillemard, Deformable mirror technologies for astronomy at CILAS, in N. Hubin, C. E. Max and P. L. Wizinowich (editors), Adaptive Optics Systems, volume 7015, 151 – 162, International Society for Optics and Photonics, SPIE (2008), doi:10.1117/12.787400.

[108] J. E. Pearson and S. Hansen, Experimental studies of a deformable-mirror adaptive optical system, J. Opt. Soc. Am., **67**(3):325–333 (1977), doi:10.1364/JOSA.67.000325.

[109] A. Lazzarini, G. H. Ames and E. K. Conklin, Methods of hierarchical control for a segmented active mirror, in J. V. Walker and E. E. M. IV (editors), Laser Power Beaming, volume 2121, 147 – 158, International Society for Optics and Photonics, SPIE (1994), doi:10.1117/12.174179.

[110] R. G. Dekany *et al.*, 1600 actuator tweeter mirror upgrade for the Palomar Adaptive Optics system (PALAO), in P. L. Wizinowich (editor), Adaptive Optical Systems Technology, volume 4007, 175 – 179, International Society for Optics and Photonics, SPIE (2000), doi:10.1117/12.390293.

[111] W. D. Henderson and S. V. Gunn, System Performance Of A Large Deformable Mirror Using Differential Ball Screw Actuators, in S. Holly (editor), Adaptive Optical Components II, volume 0179, 51 – 60, International Society for Optics and Photonics, SPIE (1979), doi:10.1117/12.957293.

[112] F. Deneuville *et al.*, Large aperture adaptive optics for intense lasers, in J. Hein (editor), High-Power, High-Energy, and High-Intensity Laser Technology II, volume 9513, 89 – 96, International Society for Optics and Photonics, SPIE (2015), doi:10.1117/12.2178848.

[113] S. V. Gunn *et al.*, Novel deformable mirror design for possible wavefront correction in CO2 laser fusion system, in D. H. Liebenberg (editor), Los Alamos Conference on Optics 1979, volume 0190, 258 – 258, International Society for Optics and Photonics, SPIE (1980), doi:10.1117/12.957751.

[114] T. Bifano *et al.*, MEMS deformable mirrors for adaptive optics, in Solid-State Sensor and Actuator Workshop,(Hilton Head Island, SC, June 1998), 71–4 (1998).

[115] R. J. Zawadzki *et al.*, Performance of 97-elements ALPAO membrane magnetic deformable mirror in Adaptive Optics-Optical Coherence Tomography system for in vivo imaging of human retina, Photonics Letters of Poland, **3**(4):147–149 (2011).

[116] S. A. Cornelissen *et al.*, 4096-element continuous face-sheet MEMS deformable mirror for high-contrast imaging, Journal of Micro/Nanolithography, MEMS, and MOEMS, **8**(3):1 – 8 (2009), doi:10.1117/1.3158067.

[117] V. Bagnoud and J. D. Zuegel, Independent phase and amplitude control of a laser beam by use of a single-phase-only spatial light modulator, Opt. Lett., **29**(3):295–297 (2004), doi:10.1364/OL.29.000295.

[118] J. M. Andersen *et al.*, Characterizing vortex beams from a spatial light modulator with collinear phase-shifting holography, Appl. Opt., **58**(2):404–409 (2019), doi:10.1364/AO.58.000404.

[119] A. M. Weiner *et al.*, Programmable femtosecond pulse shaping by use of a multielement liquid-crystal phase modulator, Opt. Lett., **15**(6):326–328 (1990), doi:10.1364/OL.15.000326.

[120] A. Suda *et al.*, A spatial light modulator based on fused-silica plates for adaptive feedback control of intense femtosecond laser pulses, Opt. Express, **9**(1):2–6 (2001), doi:10.1364/OE.9.000002.

[121] V. A. Schanz, M. Roth and V. Bagnoud, Picosecond contrast degradation by surface imperfections in chirped-pulse-amplification stretchers, J. Opt. Soc. Am. A, **36**(10):1735–1742 (2019), doi:10.1364/JOSAA.36.001735.

[122] S. Roeder *et al.*, How the laser beam size conditions the temporal contrast in pulse stretchers of CPA lasers, High Power Laser Science and Engineering (2022).

[123] J. Mocci *et al.*, A Multiplatform CPU-Based Architecture for Cost-Effective Adaptive Optics Systems, IEEE Transactions on Industrial Informatics, **14**(10):4431–4439 (2018), doi:10.1109/TII.2018.2799874.

[124] N. Sherriff, Learn Qt 5: Build modern, responsive cross-platform desktop applications with Qt, C++, and QML, Packt Publishing Ltd (2018).

[125] R. Conan and C. Correia, Object-oriented Matlab adaptive optics toolbox, in E. Marchetti, L. M. Close and J.-P. Véran (editors), Adaptive Optics Systems IV, volume 9148, 2066 – 2082, International Society for Optics and Photonics, SPIE (2014), doi:10.1117/12.2054470.

[126] M. J. Townson *et al.*, AOtools: a Python package for adaptive optics modelling and analysis, Opt. Express, **27**(22):31316–31329 (2019), doi:10.1364/OE.27.031316.

[127] T. J. Maceina and G. Manduchi, Assessment of General Purpose GPU Systems in Real-Time Control, IEEE Transactions on Nuclear Science, **64**(6):1455–1460 (2017), doi:10.1109/TNS.2017.2691061.

[128] T. N. Truong *et al.*, Real-time wavefront control for the PALM-3000 high order adaptive optics system, in N. Hubin, C. E. Max and P. L. Wizinowich (editors), Adaptive Optics Systems, volume 7015, 995 – 1002, International Society for Optics and Photonics, SPIE (2008), doi:10.1117/12.790457.

[129] J. Hagelberg *et al.*, KalAO the swift adaptive optics imager on 1.2m Euler Swiss telescope in La Silla, Chile (2020), doi:10.48550/ARXIV.2012.09065, URL https://arxiv.org/abs/2012.09065.

[130] A. Surendran *et al.*, Scalable platform for adaptive optics real-time control, part 1: concept, architecture, and validation, Journal of Astronomical Telescopes, Instruments, and Systems, **4**(3):1 – 11 (2018), doi:10.1117/1.JATIS.4.3.039001.

[131] A. Surendran *et al.*, Scalable platform for adaptive optics real-time control, part 2: field programmable gate array implementation and performance, Journal of Astronomical Telescopes, Instruments, and Systems, **4**(3):1 – 8 (2018), doi:10.1117/1.JATIS.4.3.039002.

[132] A. L. Rukosuev *et al.*, 1.5 kHz adaptive optical system for free-space communication tasks, in H. Hemmati and D. M. Boroson (editors), Free-Space Laser Communications XXXII, volume 11272, 331 – 340, International Society for Optics and Photonics, SPIE (2020), doi:10.1117/12.2548337.

[133] D. Petrascheck, Elektrodynamik, SpringerLink : Bücher, Springer Spektrum, Berlin, Heidelberg (2015), ISBN 978-366-24345-7-4.

[134] G. Kirchhoff, Zur Theorie der Lichtstrahlen, Annalen der Physik, **254**(4):663–695 (1883), doi:10.1002/andp.18832540409.

[135] J. A. Stratton and L. J. Chu, Diffraction Theory of Electromagnetic Waves, Phys. Rev., **56**:99–107 (1939), doi:10.1103/PhysRev.56.99.

[136] J. Kim, Y. Wang and X. Zhang, Calculation of vectorial diffraction in optical systems, J. Opt. Soc. Am. A, **35**(4):526–535 (2018), doi:10.1364/JOSAA.35.000526.

[137] E. Wolf, Electromagnetic diffraction in optical systems - I. An integral representation of the image field, Proc. R. Soc. Lond. A, **253**:349–357 (1959), doi:10.1098/rspa.1959.0199.

[138] B. Richards and E. Wolf, Electromagnetic diffraction in optical systems, II. Structure of the image field in an aplanatic system, Proc. R. Soc. Lond. A, **253**:358–379 (1959), doi:10.1098/rspa.1959.0200.

[139] E. Wolf and Y. Li, Conditions for the validity of the Debye integral representation of focused fields, Optics Communications, **39**(4):205–210 (1981), doi:10.1016/0030-4018(81)90107-3.

A. Available Technology

Since Adaptive Optics (AO) has been invented, a lot of different approaches on how to implement control loops have been developed. In this chapter, I will briefly describe associated technology that is well established and of relevance for the field of High Intensity (HI) lasers. Furthermore, I will list strengths and weaknesses, which helps to determine which technology should be used in a certain case. In section A.1, I will start with different approaches to WaveFront (WF) measurement which are either established or of interest for HI laser systems. I will continue with Deformable Mirror (DM) technology in section A.2 and close this chapter with the most common control loop schemes in section A.3.

A.1. Wavefront Metrology

The measurement of the WF of a beam is an integral part of beam quality assurance as the phase determines the beam propagation. As discussed before, uncontrolled phase variations can transform into undesired intensity maxima in the Near Field (NF) or into Spatio Temporal Coupling (STC) when passing the pulse compressor, or ultimately diminish the peak intensity in the focal spot (see section 4.1). Therefore, some of the vast variety of WF measurement techniques that arose over the decades found their application in HI laser systems.
In this section, I will discuss a selection of WF metrology approaches that are either established in this field or find potential application in the future. However, I will only discuss techniques that first, do only measure the purely spatial properties of the beam and second, do not rely on an additional reference beam, as other methods are either out of the scope of this work or cannot be applied in the usual context of beam control.

A.1.1. Shack-Hartmann Sensor

In 1900, Johannes Hartmann measured the collimation of a telescope by using used the property of light to propagate normal to its local WF. To do so, he mounted a plate with holes in front of the aperture of the telescope and traced the individual rays from each hole through the optical system. From the displacement of each ray, he was able to determine the local WF gradient over each hole [87]. While this technique remains relevant today (see next section), two disadvantages make it unfeasible for the majority of applications: first, the precision of this method is limited by the size of the holes. Large ones result in larger uncertainties, while small holes introduce a significant amount of diffraction. Second, the plate obscures the main beam path, which is inconvenient if one intends to use the beam for an application while simultaneously observing the WF.
For this reason, Roland Shack adapted this approach for small beams by replacing the Hartmann plate with a Micro Lens Array (MLA) [88], which mitigates the diffraction issues. The position of each focal spot relative to the axis of the micro lens then corresponds to the mean WF gradient over the lens aperture. Today, this type of sensor is known as a Shack-Hartmann Sensor (SHS), where the positions of the focal spots is recorded using a camera (see Fig. A.1, left). These sensors are easy to manufacture and commercially available.

The measurement of a WF using an SHS works by measuring the mean gradient over each MLA aperture. The resulting gradient field can then be used to reconstruct the WF using zonal (e.g. the Southwell algorithm [89]), modal (i.e. by fitting an orthogonal basis over the aperture to the slopes [47]) or direct approaches (e.g. the Harker-O'Leary approach, solving a linear equation system [90]).

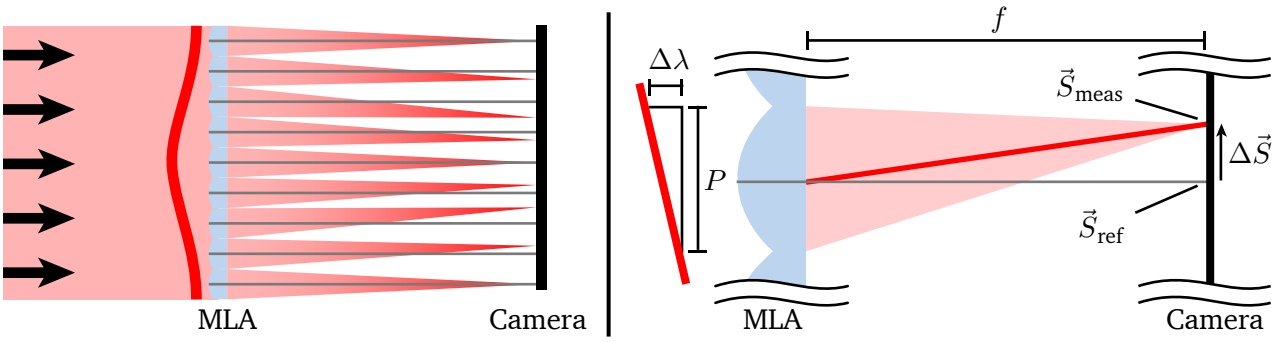

Figure A.1.: Left: Schematic of a SHS. Right: The geometry under a single microlens.

The gradient of a WF W over a subaperture i can be calculated by simple trigonometry in two steps (see Fig. A.1, right, for the geometry): first, one measures the focal spot position $\vec{S}_{\text{meas}}$ and calculates the gradient by taking the difference to a reference position $\vec{S}_{\text{ref}}$ and applying the small angle approximation:

$$\left.\frac{dW}{d\vec{r}}\right|_i = \frac{\vec{S}_{\text{meas},i} - \vec{S}_{\text{ref},i}}{f} = \frac{\Delta\vec{S}_i}{f},$$

where f is the distance of the MLA to the camera chip. In the second step, the gradient per unit length is converted to a gradient per microlens diameter $\delta\vec{\lambda}$, given in WaveLengths (WLs), which is much more convenient for WF reconstruction:

$$\left.\delta\vec{\lambda}\right|_i = \frac{P}{\lambda}\cdot\left.\frac{dW}{d\vec{r}}\right|_i = \frac{P}{\lambda f}\Delta\vec{S}_i,$$

where P is the pitch of the MLA, i.e. the diameter of the lens.
Aside from the simple manufacturing process, the SHS design has several advantages that led to the widespread use of this sensor type, most of which are linked to the robustness of the measurement. First, the SHS provides usable data no matter which part of the sensor is illuminated. Second, even strong WF deformations with large amounts of tilt yield usable data, as long as a corresponding reference is taken (note that extreme values will lead to a loss of the small angle approximation). These two points enable using an SHS on-the-fly while assembling and aligning an optical setup, just by manually placing it into the beam path. Third, the SHS is achromatic. While a changing WF changes the focal length of the MLA, the distance f to the camera chip remains unchanged. Therefore, the spot size on the camera chip changes, but not the position of its centroid. Therefore, one can use an SHS for measuring pulses with broad spectral ranges.
Another advantage lies in the flexibility of the concept itself. By using different focal lengths or MLA pitches, a broad spectrum of applications can be covered, ranging from high spatial resolutions (small pitch, small f) over ultrasensitive sensors (large pitch, large f) to large frame rates (low number of subapertures, low camera resolution).
The main disadvantage of the SHS concept is the limited spatial resolution compared to other WF diagnostics. Even models with high resolutions are restricted by the necessity of having several pixels per microlens, which is why Multi-Wave Lateral Shearing Interferometer (MWLSI) WaveFront Sensors (WFSs) (see section A.1.3) or Pyramid WaveFront Sensors (PWFSs) (see section A.1.4) generally perform better in this area.

A.1.2. Full Aperture Hartmann Mask

Even though the SHS provides a robust and flexible way of measuring the WF of a beam, its size is limited to the size of the camera chip. As discussed in section 4.2.3, this introduces the necessity to shrink the main beam of a laser system down to a suitable size before measuring it, where the imaging system will itself introduce aberrations. While careful alignment can minimize these, a calibration routine for the large beam diameter makes taking meaningful measurements a lot easier.
This is the point where the approach of Hartmann comes into play again. In 2018, Dorrer et al. [64] proposed to insert an apodized Hartmann mask (an opaque plate with a grid of holes with softened edges to mitigate diffraction artifacts) into the large beam, recording the positions of the holes using a plain NF camera in a following beam diagnostic (see Fig. A.2, left). The Hartmann mask is then moved along the beam propagation direction by a distance of a meter or more, where larger beam diameters require quadratically larger displacements for equal sensitivity.
Any aberrations in the beam will now shift the points on the NF camera independently from the aberrations introduced by the demagnification system (see Fig. A.2, right), granting direct insight into the WF of the large beam. The evaluation is analogous to the one of the SHS, which I discussed in the last section. Running a closed loop on this type of diagnostic can directly optimize the large beam. Once this is done, the mask can be removed and a reference WF recorded at the WFS.
Up to this point, the method has not been demonstrated on large scale beams and while it may be feasible, it requires a lot of space as the mask has to be moved on meter scales along the propagation direction, which may not be possible in all laser systems.

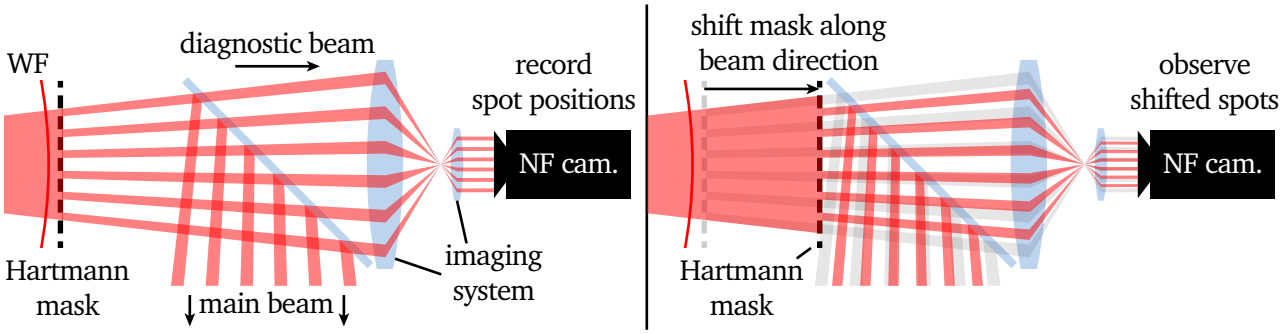

Figure A.2.: Scheme of the WF measurement method in the large beam proposed by Dorrer et al. [64], shown for a divergent beam.
Left: the Hartmann mask divides the main beam in subapertures. The beam positions are recorded by a regular NF camera within a beam diagnostic.
Right: after the mask has been moved along the beam propagation direction, the observed positions of the subapertures change according to aberrations in the main beam independently from the quality of the imaging system in front of the NF camera.

A.1.3. Multi Wave Lateral Shearing Interferometer

A different way to assess the local WF gradient is using so-called shearing interferometry [91]. Here, a beamsplitter generates two replicas of a beam, which then overlap with a spatial offset and tilt. This generates a fringe pattern that features straight lines for an unaberrated beam and tilted or curved lines if defocusing or other aberrations are present, respectively (see Fig. A.3). This has the advantage that the beam acts as its own reference, where the local WF gradient is translated to a larger or lower Optical Path Difference (OPD) due to the spatial offset of the replica with respect to each other.

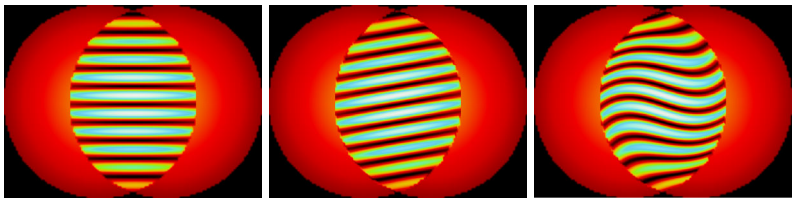

Figure A.3.: Simulated examples of shearing interferograms with horizontal shearing and vertical tilt for an unaberrated beam (left), a defocused beam (center) and a beam with spherical aberration (right).

While a shearing interferometer can be as simple as a wedged glass plate where the reflexes on the front and back side overlap on a screen, which is still used for manual optic alignment, this technique only grants insight into the derivative of the WF along the shearing direction. This makes the interferograms ambiguous. As an example, defocusing cannot be told apart easily from orthogonal astigmatism (relative to the shearing axis) as both feature a linear one-dimensional derivative in the shearing direction.
To counteract this, J. Primot proposed to overlap three or more replicas instead of two in 1993 [92], each displaced and tilted in a unique direction by the same amount, which he achieved using a series of beamsplitters. The resulting interferograms (examples for four overlapping waves, which is named Quadri-Wave Lateral Shearing Interferometer (QWLSI), are shown in Fig. A.4) grant insight into the WF derivatives along the primary shearing axes and, additionally, in all other directions in which beam replica overlap, e.g. the diagonals for QWLSI (compare Fig. A.5).

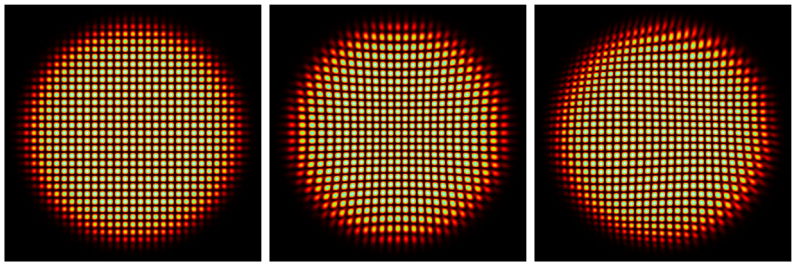

Figure A.4.: Simulated examples of QWLSI interferograms for an unaberrated beam (left), a beam with spherical aberration (center) and a beam with irregular aberrations (right).

The extraction of the approximate WF gradients along these axes is possible using a Fourier transform of the interferogram [93], which I have illustrated in Fig. A.5 for QWLSI. To understand this, it is useful to recall how the interferogram is formed: starting from the NF of the beam (pictures on the left), four displaced and tilted replicas overlap (which is shown next to the NF image, where the colored circles indicate the apertures of the replicas and the colored arrows the displacement $\vec{S}_i$). The tilt of the replica relative to each other now generate a regular pattern of constructive and destructive interference, where a larger tilt would increase the spatial frequency of this modulation. The shearing on the other hand leads to interference of different parts of replica, making the local frequency of the pattern sensitive to the curvature of the WF.
This information can be accessed in the Fourier domain of the image (blu-ish image on the right of the interferogram in Fig. A.5). Here, the central peak holds the information on the intensity distribution in the NF of the beam, while the surrounding peaks are formed by the high spatial frequency in the interferogram, and therefore contain the information about the phase of the pattern.

By first applying a window F_i (green boxes) and performing an inverse Fourier transform afterwards, one obtains a complex representation of the NF, where the argument depends on the phase difference between the replica at each position. For small amounts of shearing, this can be converted into to the approximate local WF gradient in the selected direction. In order to do so, the argument has to be unwrapped first as it can only lie in between -π and π, and the tilt T_i between the replica has to be subtracted. The approximate WF gradient in the direction i, given in WLs can then be obtained by dividing this function by 2π and the absolute value of the displacement between the interfering replica. Wrapped up, the extraction of the approximate gradient works as follows:

$$\frac{\partial W}{\partial x_i} \approx \frac{W\left(\vec{x}+\frac{\vec{S}_i}{2}\right) - W\left(\vec{x}-\frac{\vec{S}_i}{2}\right)}{\left|\vec{S}_i\right|} = \frac{\mathcal{U}_i(K)}{2\pi\left|\vec{S}_i\right|}$$

with the operation $\mathcal{U}_i(K) := \text{Unwrap}(\arg(\text{iFFT}(F_i(K)))) - T_i,$

where
- W is the WF in WLs,
- $\vec{S}_i$ the displacement between the overlapping replicas in direction i,
- K the Fourier transform of the interferogram,
- T_i the tilting phase between the overlapping replicas in direction i,
- F_i the windowing function in the Fourier-plane in direction i,
- iFFT the inverse Fourier transform
- arg the argument function on a complex number,

and
- Unwrap the unwrapping function which eliminates 2π phase jumps.

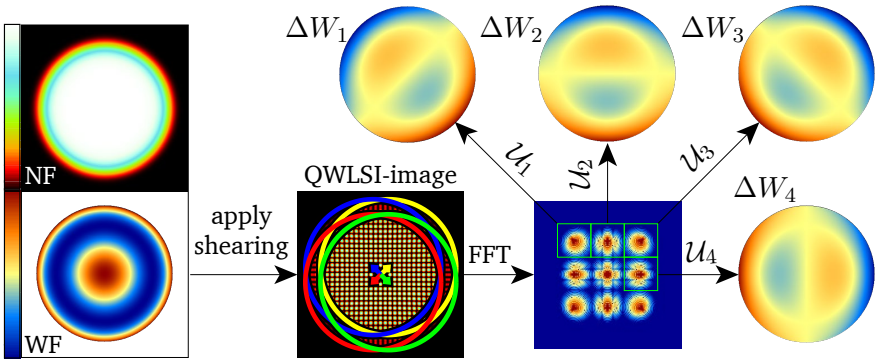

Figure A.5.: Schematic of the evaluation of a QWLSI WFS, as described in the text. In this example, the incoming wave features some spherical aberration.

While Primot used beam splitters to obtain the overlapping waves, one innovation made this approach much more convenient to use: by passing the beam through a grating and using the first diffraction orders as replica (other orders are removed using a spatial filter, see Fig. A.6), the design gets a lot more compact and thus portable while still being achromatic enough to measure broadband light-sources with acceptable systematic errors due to the small diffraction angle [94]. This has been demonstrated for three waves and QWLSI [94, 95].

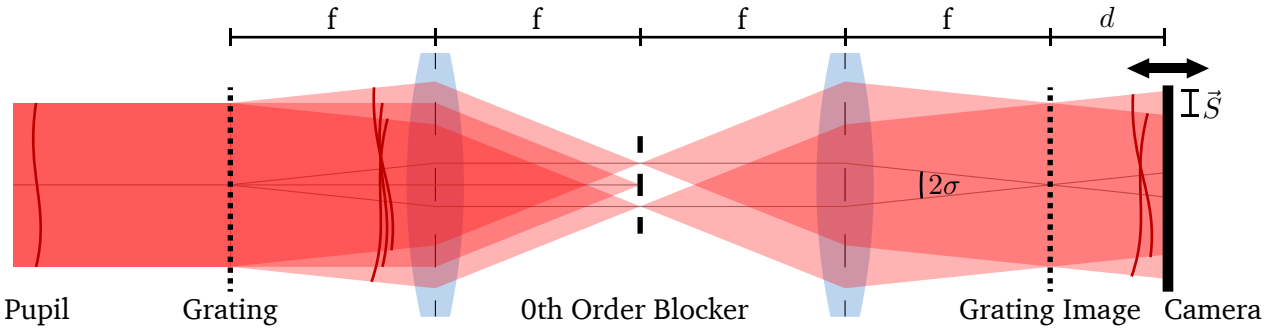

Figure A.6.: Schematic of a possible implementation of a MWLSI WFS. A grating generates beam replica. The zero-order beam has to be removed by a spatial filter, while the camera is located with an offset d from the image plane of the grating.

MWLSI is an established WF measurement strategy in HI laser systems as well, as this device is as compact and robust as the SHS and also commercially available. Compared to the SHS, MWLSI has additional advantages: it does feature a high spatial resolution, which is in principle only limited by the shearing distance ($\vec{S}$ in Fig. A.6) or the camera resolution. At the same time, it offers an adjustable dynamic range without changing any of the optics, as the shearing distance $\vec{S}$ can simply be changed by changing the distance d of the camera to the replication plane (see Fig. A.6). Here, a larger distance increases the sensitivity of the sensor at the cost of spatial resolution and vice versa.
The disadvantages of this technique are linked to the more complex setup compared to the SHS (which may introduce sensor-internal aberrations) and the evaluation of the sensor, which is computationally more demanding. The first issue can be mitigated by a suitable calibration, while the second one is primarily the price for the higher spatial resolution. This is why the MWLSI is not used for high-speed applications but in cases where the high resolution is necessary.

A.1.4. Pyramid Wavefront Sensor

A third way to measure the WF of a beam in the NF is the Pyramid WaveFront Sensor (PWFS). Even though this technique is not established in the HI laser community by now, it found widespread use in astronomy due to its desirable properties concerning the spatial resolution and computational costs when feeding the data to an AO loop in real-time applications. If fast dynamic aberrations shall be compensated in the laser community in the future, this technology may therefore be of interest here as well.
The concept is a modification of the schlieren imaging technique, developed by A. Toepler in 1864 [96]. The original approach uses a knife edge to block one half of the light in the Fourier plane of an aperture, which, geometrically speaking, removes all rays from the beam that feature a gradient towards the knife and let all rays pass that point in the other direction. This allows to visualize the WF gradient in the NF after the Fourier plane as dark and bright areas.
The PWFS, proposed by R. Ragazzoni in 1996 [97], takes this approach a step further by not blocking but deflecting parts of the beam by inserting a vibrating pyramidal prism into the Fourier plane (see Fig. A.7). For the static version of this sensor, the pyramid divides the beam into four parts exactly in the focal spot. Naturally, this generates four replica of the NF, each of which contain information about the beam gradient analogous to the schlieren image.

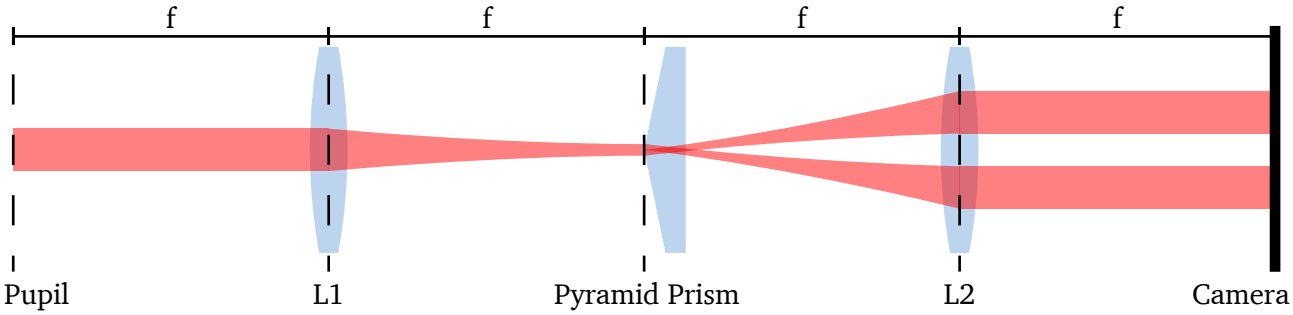

Figure A.7.: Schematic of a static PWFS, using a 4f-imaging of the pupil plane. The tip of the pyramidal prism divides the beam in the Far Field (FF) into four parts.

The pupil image of such a static PWFS is shown in Fig. A.8 on the top right. This image is extremely sensitive to low order aberrations, as even a slight tilting of the WF will move the focal spot off the pyramid tip, transferring all the light to only two or even one of the four pupil images. This also shows that such a sensor is highly nonlinear [98].
That is the reason why Ragazzoni proposed to modulate the prism with a rectangular or circular translation [97], making the tip of the pyramid circle around the center of the unaberrated focal spot (or alternatively: make the focal spot circle around the pyramid tip, see Fig. A.8, bottom central image) and integrate the light on the camera plane over an integer number of revolutions. Doing this, every ray of the pupil spends time on every facet of the pyramid, while the ratio between these times depends on the aberrations, i.e. in which direction the ray points in the NF. As an example, a ray pointing to the right in the NF will naturally spend more time on the right facets in the FF as a larger translation is required to move it onto the left ones.
The amplitude of the motion of the pyramid/focal-spot now determines the fraction of the time it takes for the focal spot to sweep over a single one of the pyramid edges compared to the time it takes to travel from one edge to the next. This means that a larger motion amplitude reduces the impact of the aberrations on the changes in the image plane of the PWFS, which effectively changes the dynamic range of the sensor.

Doing this, the linearity of the sensor drastically increases at the cost of sensitivity [99]. An assortment of example images is shown in Fig. A.9. Ragazzoni has shown that the horizontal and vertical WF gradients can be estimated from the pupil images P_i (see Fig. A.8 for the indicees):

$$\frac{\partial W}{\partial x} = \frac{(P_a + P_b) - (P_c + P_d)}{P_a + P_b + P_b + P_d} \times \frac{F}{\delta V}, \tag{A.1}$$

$$\frac{\partial W}{\partial y} = \frac{(P_a + P_c) - (P_b + P_d)}{P_a + P_b + P_b + P_d} \times \frac{F}{\delta V}. \tag{A.2}$$

Here, W represents the WF over the pupil, F the F-number of the lens that focuses the pupil onto the pyramid and δV the modulation amplitude relative to the focal spot diameter. This is, however, only an estimation and will not deliver precise results. In fact, the reconstruction of the WF from different types of PWFSs remains a challenging problem up to this day [100].

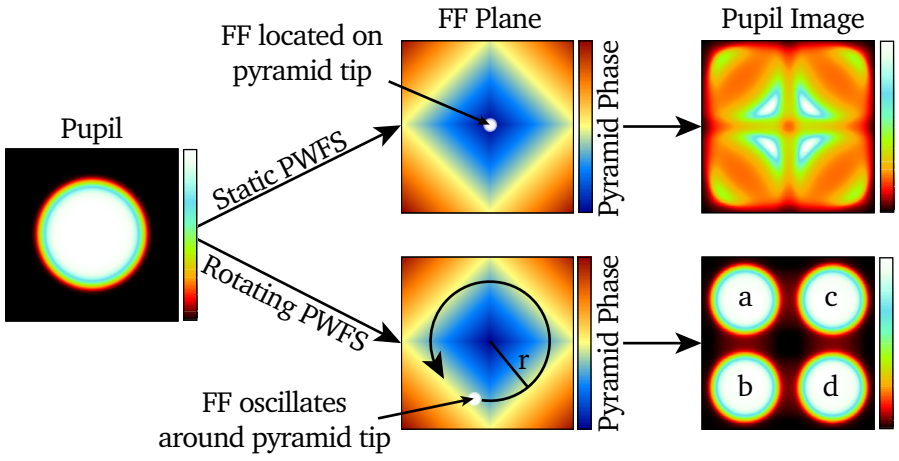

Figure A.8.: Comparison between the focal and image planes of a static (top) and rotating (bottom) PWFS.
Left: the intensity distribution of the NF. Center: the phase map of the pyramid with a white circle indicating the position of the FF. Right: the image plane of the PWFS for the aberration free case. The pupil-images labels of the rotating PWFS are explained in the text.

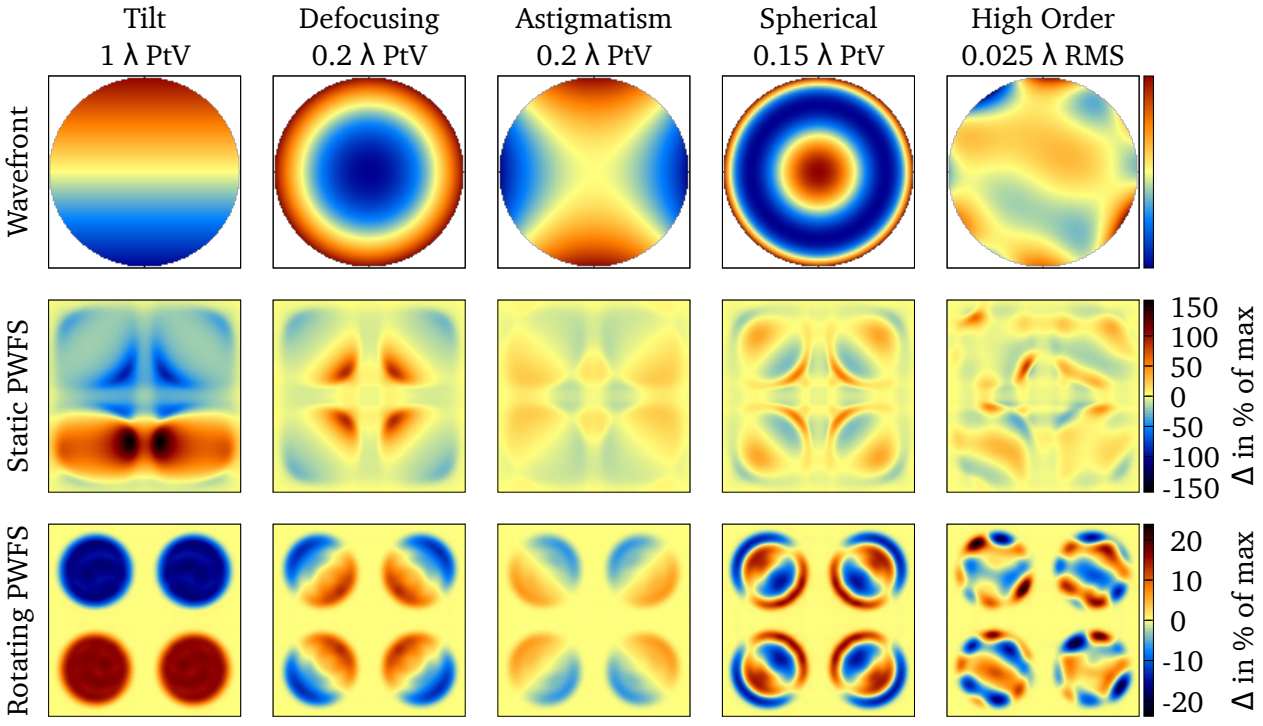

Figure A.9.: A simulated comparison of the pupil images of a static PWFS (central row) and a rotating PWFS (bottom row) for different input WFs (top row). For all images, the relative deviation in percent from the aberration-free case (see Fig. A.8, right) are shown, as this is used for AO control and shows the difference in sensitivity more clearly.

Speed The reason why this type of sensor has nonetheless such a widespread use in astronomy is a different one: as equations (A.1) and (A.2) indicates, the local intensity difference of the normalized pupil images from the aberration-free case is proportional to the local WF gradient. This means that one only has to normalize the full PWFS image, subtract a reference and can then use the resulting data to directly feed an AO loop. These operations are computationally very cheap and allow for WF control with low latency and repetition rates of several kHz, while still having a good spatial resolution. This is especially of interest for so-called *extreme AO loops*, featuring thousands of actuators that have to be controlled in real-time [33, 101].
For HI laser systems on the other hand, where the demand on actuator numbers will likely stay behind astronomy, the PWFS may still be applied for rapid beam stabilization if a Continuous Wave (CW) pilot beam is available for constant feedback.

Nonlinearity We have seen before that the PWFS is extremely sensitive and quickly becomes nonlinear (see Fig. A.9). Even the modulated PWFS runs a lot faster into limitations than the schemes I discussed before. In astronomy, this is usually less of an issue as the WFS can be placed close to the final imaging camera where the aberration free case corresponds to a flat WF. In HI laser systems however, as discussed in sections 4.2.3 and 4.3.2, the WFS has to be located after an imaging system, where the ideal case does not necessarily correspond to a flat WF anymore. This means that the PWFS will likely have to operate in a non-ideal regime, which does not allow for great freedom concerning calibration routines. A possible scheme to overcome this is to use successive or nested AO loops which guarantee that the PWFS operates within its linear regime while stabilizing the beam, where the remaining static aberrations are eliminated elsewhere.

After I have discussed the strengths, weaknesses and applications of the modulated PWFS, it is worth noting that a significant drawback of this scheme is the complexity of the setup, where moving parts have to be involved. Furthermore, the manufacturing of a pyramidal prism with sharp edges where all four planes intersect at the same point is a significant challenge, increasing the costs for these optics well beyond the prices of components for other sensor types. While this issue can be circumvented by using two roof prisms which are almost in contact with each other [102], this poses a further complication in the construction of a PWFS.

A.1.5. Phase Retrieval

The last method of WF measurement that I will discuss in this section does follow a completely different approach. Here, the WF in the NF shall be reconstructed using data on the fluence distribution in the FF, i.e. images of the focal spot, which is extremely useful for the cross calibration of NF WFSs at other locations to the focal spot in the target chamber.
Solving inverse optical problems of this kind is commonly referred to as Phase Retrieval (PR). A vast variety of applications rely on PR, e.g. phase- and X-ray microscopy, image encryption and, of course, the WF reconstruction as well [75]. The problems can be categorized into two kinds: the ones where the fluence distribution is known in the object- and frequency domain (e.g. the NF and FF) and the ones where only one side is known.
In the application of WF retrieval in HI laser systems, we usually deal with the case where the fluence distribution is known on both sides. Even if no NF measurement is available, the shape and size of the aperture is known and can be used to synthesize the data accordingly.
For this scenario, the Gerchberg-Saxton algorithm [103] is a typical approach to do exactly that by using

one image in the NF and one in the FF. Here, one propagates a numeric representation of the beam back and forth between the two data supports by applying Fourier transforms. After each transform, the magnitude of the complex electric field is replaced by the magnitude obtained from the measured fluence distributions. After several iterations, the phase of the beam approaches a stable solution to this problem.
One major drawback here is that the problem is ”ill-posed” when phrased like this, meaning that there is too little information to find a unique solution. In other words, a multitude of solutions exist for one intensity distribution. A countermeasure for this is to add more constraints to the problem, which is often done by measuring the fluence distribution not only in the FF but in additional intermediate planes in front and/or behind the focal plane (see Fig. A.10), and feed these to an adapted version of the Gerchberg-Saxton algorithm.

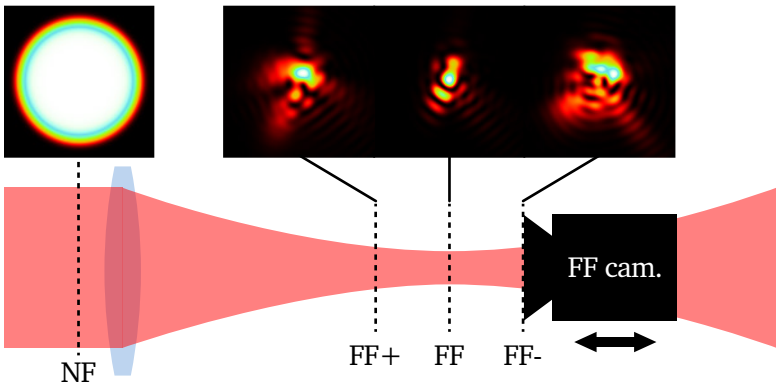

Figure A.10.: Schematic of the data acquisition process for the classical Gerchberg-Saxton PR approach.

An alternative way to record the same data set relies on the displacement theorem, which states that the three-dimensional intensity distribution around the focal plane of an optic can be shifted by applying moderate amounts of tip/tilt for the transversal and defocusing for the longitudinal direction [46, sec. 9.1]. Therefore, the focal plane can be shifted back and forth by moving a lens or using a DM to apply small amounts of defocusing. The applied WF delta between the corresponding steps can then be saved and fed into the algorithm.
From this point, there are different ways to feed this enriched data set to the Gerchberg-Saxton algorithm [75]. One adaption is illustrated in Fig. A.11, where the algorithm swaps between the different datasets for each iteration by adding the recorded WF delta in the NF (the blue ”Apply Offset” boxes in Fig. A.11) and applying the corresponding magnitudes in the NF and in the Fourier domain (the red ”Apply Intensity” boxes in Fig. A.11), which can be considered a serial adaption of the Gerchberg-Saxton algorithm.
Other approaches directly propagate the beam from one FF measurement to the next before going back to the NF, or run the algorithm in parallel for each data set and average the WFs in each iteration.

The Gerchberg-Saxton approach and its modifications face a set of limitations, which challenge its usability. First of all, the usage of Fourier transforms for the propagation between the FF and the NF imposes the small-angle approximation and the assumption of a purely monochromatic pulse. This automatically means that STC cannot be included in this model, which may lead to the failure of the WF reconstruction in the worst scenario. At Petawatt High-Energy Laser for heavy Ion eXperiments (PHELIX),

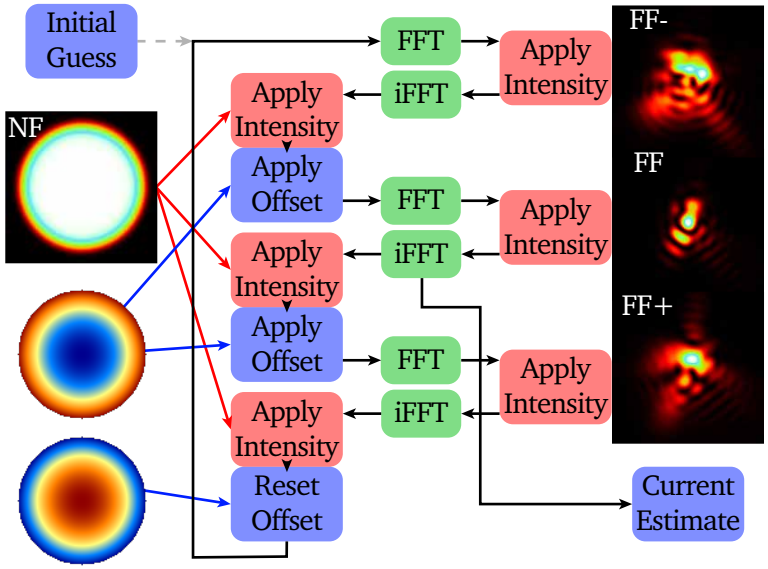

Figure A.11.: Flow chart of a modified, serial Gerchberg-Saxton PR algorithm with three FF measurements and two WF deltas.

this is exactly the case as the pulse is spatially chirped at the target chamber entrance due to the Treacy compressor geometry [45]. More elaborated models, using polychromatic vectorial diffraction integrals (see appendix B), are unfeasible due to the extreme numerical load even for a single iteration.
In the last decade, other approaches to PR arose which rely on machine learning. In these cases, the FF data are either fed to a machine learning model directly or decomposed into descriptors first, while the model has been trained on large data sets to estimate the WF that produced the set of measured FFs [76]. Other approaches go a step further and try to eliminate the necessity of several FF images, e.g. by enriching the available data with diffraction masks that are placed into the focal plane [77]. These methods may still be feasible in the presence of STC and provide reasonably accurate WF estimations. However, the development of a suitable machine learning model requires a substantial amount of background knowledge and the task of generating a sufficiently large data set is challenging (and potentially extremely time consuming if the data cannot be simulated easily). For these two reasons, corresponding investigations were consciously excluded from this work, but remain of interest for future projects.
As an alternative, I developed a calibration device for the NF WFS, which relies on extreme spatial filtering, which I presented in section 5.3. Furthermore, a promising measurement routine of non-radial aberrations in the FF, based on beams with Orbital Angular Momentum (OAM), is currently under investigation in the frame of a supervised master's thesis [78].

This concludes the first section of the chapter about AO technology, dealing with WF metrology in HI laser systems. However, the knowledge about the phase in the NF is only the first step towards beam correction.

A.2. Deformable Mirrors

The second important part of an AO loop is the active component, spatially changing the OPD over the aperture of the beam. There are different concepts on how to do so and I will describe a selection of technologies that repeatedly appeared in HI lasers over the last decades. While the most widespread concept is using a DM, I will also touch on the topic of Spatial Light Modulators (SLMs).

A.2.1. Piezoelectric Deformable Mirrors

The first category of DMs that is well established in the field of HI lasers uses piezoelectric actuators to deform the mirror membrane. These actuators are based on the piezoelectric effect, where a material (either a crystal lattice or a ceramic) changes its volume when exposed to an electric field [104, sec. 7.2.5]. The deformation is approximately linear over the applied voltage.

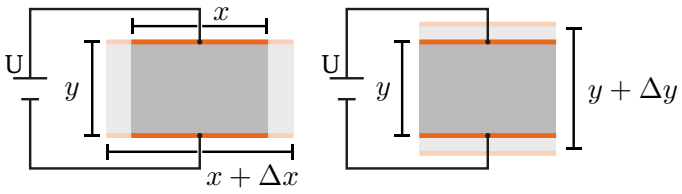

Figure A.12.: Schematic of a transversal (left) and a longitudinal piezo actuator (right). The piezoceramic is shown in gray, while the electrodes are shown in orange.

In general, piezo materials are anisotropic, which means that the direction of expansion depends on the orientation of the material relative to the applied electric field. This in turn means that the piezoelectric coefficient has to be written as a tensor.
In engineering, one is mainly interested in two special cases: transversal and longitudinal expansion (see Fig. A.12). If the electric field E is oriented such that $\vec{E} \parallel \vec{e}_y$, the change in length without the presence of any external force can be written as:

$$\Delta y = \delta_{\text{long}} \cdot E \cdot y = \delta_{\text{long}} \cdot U \qquad \text{in the longitudinal case and} \tag{A.3}$$

$$\Delta x = \delta_{\text{trans}} \cdot E \cdot x = \delta_{\text{trans}} \frac{U}{y} x \qquad \text{in the transversal case.} \tag{A.4}$$

Here, δ_{long} and δ_{trans} are the longitudinal and transversal piezoelectric coefficients, respectively, given in meters per volt, y is the size of the piezo element parallel to the electric field and x the transversal size, while U is the applied voltage. Obviously, the changes in length depend on the material properties and the electric field strength. In typical applications, the expansion is small compared to other actuator types, often measured in micrometers. This makes them particularly interesting for applications where a high spatial precision is needed, such as DMs.
In the following, I will describe three types of DMs that use piezo elements as actuators.

Monomorph Mirrors use a single layer of a transversal piezo element with electrodes on both sides, directly glued to the rear side of a mirror substrate with a thickness of a few millimeters [105] (see Fig. A.13). If voltage is applied to the electrodes, the substrate contracts or expands, bending the mirror in either direction. This amplifies the otherwise small changes in size of the piezo element, leading to strokes of several tens of micrometers while maintaining an intrinsic, smooth curvature.

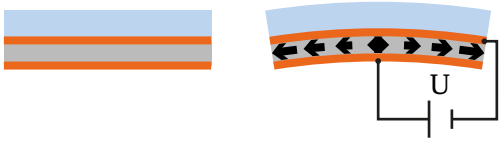

Figure A.13.: Scheme of a monomorph piezoelectric DM. Left: default state. Right: active state.

This type of DM has several advantages, compared to other types. First of all, the design is simple, making the production comparatively cheap [106]. This point also holds true for the voltage driver as the actuators only require manageable voltages of a few hundred volts at the maximum and feature a low capacity. The latter means that this type of mirror can be operated with kHz repetition rates.
Aside from the costs, this type of DM can be adapted easily to a lot of different applications. First, as the thickness of the mirror substrate is still in the range of millimeters, most available coatings can be applied to the surface, such as metallic, high-power- or chirped dielectric coatings. Second, as the electrodes on the rear side are freely accessible, the electrode pattern can be adapted to the application. Common geometries are rectangular or hexagonal tilings or concentric annular segments which are designed to compensate Zernike-polynomials as efficient as possible.
On the other hand, this technology has some downsides which permit the use of monomorph mirrors in some cases. One of the major drawbacks is the hysteresis that is inherent to piezoelectric elements: contrary to the simplified linear equations (A.3) and (A.4) for the change of length shown above, the length actually depends on the history of voltages applied before. This interferes with the linear model that is typically used for AO loops, making them converge slower than theoretically possible, and also includes a drift in length over time, which makes these mirrors unfeasible if long-term stability is needed. Another issue is that these mirrors cannot compensate tilt in the beam very well as the substrate needs to be side mounted. If one wants to do so, this requires complicated workarounds, like expanding the mirror diameter well above the beam aperture and thus emulate a tilt by tilting only the inner part of the mirror, or having special, piezo-driven mounts.
Last, the size range of this technology is limited in both directions: as the mirror substrate has to be relatively thin, large substrates will start to bend under their own weight and resonances in the membrane may start to notably interfere with the AO loop. On the other hand, the electrodes on the rear side of the mirror need to feature gaps in between, which limits how small the electrodes themselves can get without loosing efficiency. Therefore, commercial monomorph mirror diameters start at a diameter of a few centimeters and range up to a maximum of about 30 cm.

Bimorph Mirrors follow a very similar approach as monomorph ones. The main difference here is that not one layer of piezoceramic is glued to the substrate, but two. Electrodes are placed on the front and rear side of the piezoceramic layers, as well as in between [106] (see Fig. A.14). Two of these electrode layers are tiled and each tile can be controlled independently. In general, this approach features the same advantages and disadvantages as the monomorph design, aside from being slightly more complicated to manufacture.

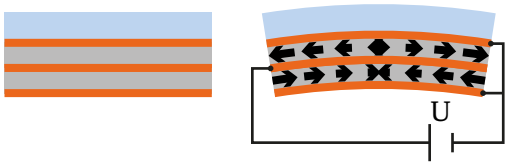

Figure A.14.: Scheme of a bimorph piezoelectric DM. Left: default state. Right: active state.

The main difference is that in this case, more elaborated actuator patterns are possible, such as having one layer of piezoceramic controlled by an electrode spanning the full diameter of the substrate, compensating only for defocusing and therefore reducing the load on the other actuators.

Piezo Stack Mirrors follow a different approach. Instead of transversal piezo actuators, longitudinal ones are used to push and pull on the membrane of the mirror. As the stroke of a single piezo actuator is too low, the actuators are stacked in a sandwich configuration with alternating polarity to achieve larger strokes [106] (see Fig. A.15).

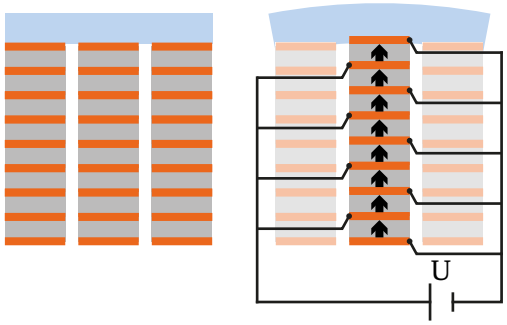

Figure A.15.: Scheme of a stacked piezoelectric DM. Left: default state. Right: active state.

As the strokes are larger than for the monomorph and bimorph designs (called curvature mirrors in the following), the actuators can be considerably smaller in diameter without loosing efficiency, enabling much higher actuator densities. While standard actuator pitches span from 5 mm to 10 mm, spacings down to 1 mm are commercially available for stacked piezo DMs [107].
At the same time, the DM is much more rigid than the curvature mirrors described earlier, as the membrane is supported by the piezo stacks over the full aperture. Therefore, the mirrors can be arbitrarily large [107], feature only little sag due to their own weight [106] and resonance frequencies are much higher, ceasing to interfere with the control loop even for multi-kHz control rates[108].
The downsides of this design are therefore different than for the curvature designs (except the hysteresis, which cannot be avoided). Contrary to the latter, piezo stack mirrors are complicated to manufacture and harder to drive as the actuators feature a much larger capacity. Naturally, this increases the price of these mirrors. At the same time, electrode failure is a lot more common due to the the more complicated design and the larger number of electrodes.
From an optical point of view, these mirrors can rarely be used as stand-alone DMs: the strokes are a lot lower than for the bending mirrors, only reaching single-digit micrometer amplitudes for off-the-shelf models. On top of that, a significant portion of this range is needed to flatten out the DM itself, which will never leave the factory completely flat. The last downside is due to the rigid grid of actuators, which leads to a ”print-through” effect, where the geometry of the actuators remains visible in the reflected WF even for an optimally corrected beam.
All these points indicate that a piezo stack DM is rather suitable for a ”tweeter-woofer” design, where a smooth, low resolution DM with a large stroke compensates for the low order aberrations, which usually feature the largest amplitude, and a piezo stack DM corrects for the high spatial frequencies afterwards [109]. While this approach is often used in astronomy [110], careful considerations have to be done if the additional correction is worth the costs.

A.2.2. Mechanical Deformable Mirrors

As mentioned above, piezoelectric actuators do have some issues with hysteresis and long-term stability. A technology that circumvents these problems is based on motorized linear actuators, which mechanically push and pull on the membrane of the DM [111]. Here, a gear converts the rotational motion of a motor, e.g. a stepper motor, into a linear motion. This motion, however, is usually too large to directly act on the mirror membrane. Therefore, the actuator is mounted to a spring-loaded mechanism that converts the linear motion into a force, which is then transferred to the membrane [112] (see Fig. A.16).

This approach has another advantage: if one actuator applies a force to the membrane and induces a deformation, the stroke is not completely blocked by the neighboring actuators and slowly decays with increasing distance. This means that each actuator has a smooth influence on the whole aperture, which makes this type of DM ideal to mimic low order polynomials with high precision.

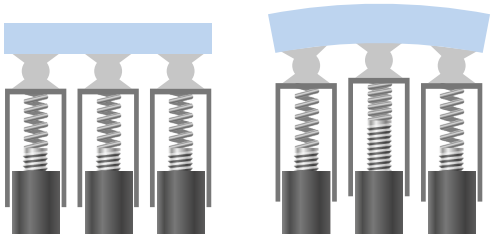

Figure A.16.: Scheme of a mechanical DM. Left: default state. Right: active state.

Mechanical DMs like these provide excellent long-term stability as the force on the membrane does not convert back into torque on the motor axis, which means that the motor drivers can simply stay inactive while the membrane remains in its current shape.
The main downside of this technology is the slow repetition rate. As each command to the DM means driving a motor to the desired position, the regulation speed is strongly limited, lying mostly below 10 Hz. This is far too low to compensate for fast evolving dynamic aberrations[1]. Furthermore, the motors generate heat, which, when running for too long, may negate the advantage of long-term stability as thermal expansion happens within the device.
Also, these DMs come at a huge cost compared to the relatively low actuator density which is achievable due to the diameter of the linear actuators.

A.2.3. MicroElectroMechanical System (MEMS) Deformable Mirrors

In the last two decades, MEMS DMs got available on the market [114]. These devices directly convert forces, carried by an ElectroMagnetic (EM) field, to the membrane of the mirror. This is different to piezoelectric DMs, which rely on the change of volume in the piezoceramic to change the shape of the mirror.
One example of MEMS DMs works with a flexible electrode, which is mounted on a grid of supports (one could imagine a tiny waffle made out of insulating material). In the gaps of the support, fixed electrodes can either attract or repel the flexible electrode, leading to a discontinuous deformation. This deformation is then transferred to the optical membrane via tiny posts, mounted on the front side of the electrode (see Fig. A.17).
Other types are available as well, using tiny magnets and coils to apply the force instead of electric fields [115].

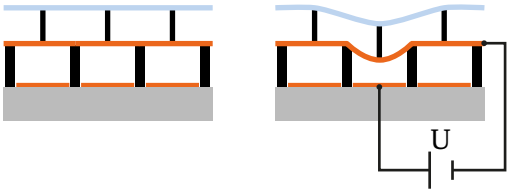

Figure A.17.: Scheme of an electrostatic MEMS DM. Left: default state. Right: active state.

The main advantage of this technology is the unprecedented density of actuators as the mirror is rather manufactured as an integrated device, mounted on a chip, than the individual assembly of actuators. DMs with actuator pitches of 300 µm have been realized, featuring more than 4000 actuators on mirrors smaller than 3 cm in diameter [116]. Furthermore, the influence function of the actuators is very localized, allowing for the compensation of high order aberrations, and can be driven with high speed. These traits make MEMS DMs an excellent choice to compensate for atmospheric turbulence.
The downsides of the MEMS technology are mainly linked to the damage threshold: as the influence of each actuator is highly localized, the membrane needs to be flexible in order to grant a feasible stroke.

[1] Efforts to circumvent the issue of slow reaction times in mechanical DMs include additional piezoelectric elements, which of course undermine the long-term stability of these mirrors [113].

This generally translates into using a thin membrane, which cannot be coated with arbitrary dielectric coatings. Therefore, these mirrors can only be used in low intensity environments, which limits their application in HI lasers to a few selected points in the system - mainly in the stretched beam that has been expanded right before entering the next amplifier. Furthermore, the high actuator density is only available for small mirrors and drastically drops for larger models. Their main advantage can therefore only be used very early in the system.
The last downside is the high pricing of the mirror, which is on one hand linked to the necessity of a specialized facility for manufacturing, and the challenges that come with controlling hundreds to thousands of actuators at the same time.

A.2.4. Spatial Light Modulators

In the last section on DM technology, I will touch on other options to modulate the phase of light instead of using DMs. In this context, I will use the more general term Spatial Light Modulator (SLM), which refers to any optic that is capable of modifying light over an aperture.
One common implementation of an SLM uses a Liquid Crystal Arrays (LCAs), which can both change the phase and the amplitude of a transmitted beam by changing the orientation of the particles in the liquid crystal by applying an electric field. Naturally, these devices have been used for the modulation of the amplitude and WF of a beam [117] and for the generation of different beam geometries [118].
Other applications of these devices include the modulation of the pulse shape by inserting an SLM into the Fourier plane of a pulse stretcher. This was proposed in the 1990s [119] and implemented many times since then using different SLM technologies, such as acousto-optic modulators or even mechanically tilting plates of glass [120].
However, even if SLMs offer great flexibility and potentially high spatial resolutions, these applications suffer from different problems: on one hand, the beam has to pass through material, which means that high intensities can quickly damage the device or suffer from dispersion and/or nonlinear effects. On the other hand, it has been shown that high frequency phase fluctuations in the Fourier plane of a pulse stretcher quickly generate a rising slope prior to the pulse, degenerating the temporal contrast on the picosecond timescale, which can interfere with the experiments in HI laser systems [121, 122].
For these reasons, transmitting SLMs are not used for the optimization of the focal spot and are not suitable for the use as active component in pre-compressor AO loops, optimizing for the compressor entrance (see section 4.1) as well. Therefore, this technology is of little interest in the context of this work, while still being extremely useful for other applications.

A.3. Control Units

The last building block of an AO loop is the Control Unit (CU). In this work, I refer to the CU as the sum of components in hard-, firm- and software that links the WFS with the DM driver, thus closing the loop, while simultaneously accepting user commands. Phrased in a little more detail, the CU receives data from the WFS, evaluates it, calculates the control signals for the DM and sends them to the DM hardware driver. At the same time, the CU often provides a user interface to change loop parameters and check on beam parameters, a data-logging mechanism or interfaces to third party software.
Up to this day, three technologies have been established to provide this functionality: Central Processing Units (CPUs), General Purpose Computation on Graphics Processing Units (GPGPUs) and Field Programmable Gate Arrays (FPGAs). In the following, I will briefly explain these terms, show how an associated AO loop architecture may look like and review the strengths and weaknesses of each one, which helps to determine suitable applications.

A.3.1. CPU

The first technology that naturally reached a widespread use is Central Processing Unit (CPU) based, which means the CU consists of a computer with a dedicated software. The big success of this technology is linked to the high availability of computers which offer large capabilities in memory and computational power, great flexibility and - last but not least - a very familiar working environment to the majority of users. Commercial hardware with plug-and-play drivers is readily available and software can be written and compiled with relatively low effort. This makes CPU based systems great testing environments while at the same time providing a feasible operating solution for a lot of applications.

Setup

The setup of a CPU based CU is straight forward (see Fig. A.18): the WFS connects to the computer either directly via established interfaces such as USB or GigE, or via a frame grabber, which can be inserted into a PCI slot of the computer.
Having the image data available, a dedicated software evaluates the raw image data and calculates corresponding control signals for the DM.
The DM is connected to the computer in a similar manner, where some examples use a USB or Ethernet interface. If high data throughput is expected, a PCI slot can be used as interface as well.
At the same time, the user can directly access the loop parameters and data streams on the computer as CU itself.

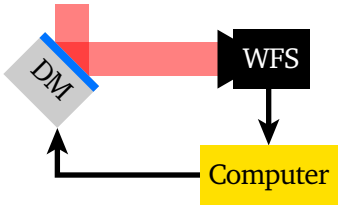

Figure A.18.
Schematic of a CPU based CU composition.

Examples

The distribution of AO software is easy due to the high-level interfacing with computers via the operating system. This is the reason why a broad variety of commercial software is available, which is often delivered together with purchased hardware. At the same time, open source solutions do exist, which offer the advantage of flexibility, enabling users to implement functionality according to their own needs.
One example for this is Wavefront Optics Measurement and Beam Analysis Tool (WOMBAT)[2], which the working group at the PHELIX laser continuously develops. I contributed to WOMBAT during this work myself [68]. WOMBAT is written in NI LabVIEW and the user can manually link individual modules for each evaluation step of the control loop via the graphical user interface. These modules run asynchronously, preventing blocking operations to slow down the rest of the evaluation chain. This offers great flexibility, proved to be extremely versatile in the laboratory environment and is very user friendly. However, WOMBAT does not provide synchronized real-time functionality and is not able to run faster than a few tens of frames per second.
Another example for an open source solution that aims for higher speed at the cost of flexibility is called "PhotonLoop"[123][3]. This project has been written in C++ using the Qt library [124] for the user interface. The code is platform-independent and traded the modular approach for higher throughput. Framerates of 500 Hz have been demonstrated on consumer level notebooks.
Besides that, a variety of open source simulation toolkits exist which aid theoretical investigations in AO in different scripting or programming environments such as MATLAB (e.g. Object-oriented Matlab adaptive optics toolbox [125]) or Python (e.g. AOTools [126]).

[2]The open source repository is hosted on an instance of GitLab on the servers of GSI Helmholtzzentrum für Schwerionenforschung GmbH (GSI):
`https://git.gsi.de/phelix/lv/wombat_ce`
[3]The source code for PhotonLoop can be downloaded on SourceForge:
`https://sourceforge.net/projects/photonloop/`

Strengths

The CPU based CU has, as mentioned before, the advantage of providing great flexibility in the way of implementation. This freedom manifests in the choice of the programming language, the way of implementation in terms of the choice of algorithms and the complexity of the application.
On top of that, one can use standard software development skills in a familiar environment, which makes CPU based CUs a great option for developing new AO schemes.

Weaknesses

The downsides of using a computer as CU are the limitations in speed and latency compared to the other technologies that I describe in this chapter. This is even true if the computer runs on a real-time operating system. The main reason for this is how the computer handles memory. Depending on the demands of the software, data may be transferred between different memory systems which are either optimized for quick access, larger bandwidth or capacity. Each of these copy operations takes time and applications can only partly influence when and where these are performed. This leads to larger latency and lower throughput, making current CPU based CUs unfeasible for loops that have to run faster than a couple of hundred Hertz.

A.3.2. GPGPU

Compared to pure CPU based CUs, one can increase the bandwidth using General Purpose Computation on Graphics Processing Unit (GPGPU). This term refers to using the processors of a graphics card that offer massive parallelization, which is often possible for processing images and large data grids. Naturally, this is the case in AO applications as the WFS delivers image data and the DM control matrix is a grid of numerical values as well.
Using GPGPU offers a similar degree of flexibility as the CPU based approach as the processing capacities of the graphics card are accessed from an application on the root complex of a computer, even though some specialized knowledge on GPGPU programming is required.
Nonetheless, the community around this technology is large and the interfaces very well documented. This means that one can greatly speed up the AO loop by designing a software which incorporates using GPGPU from the ground up.

Setup

Overall, the AO loop using a GPGPU based CU is constructed in a similar manner as the one using a CPU CU. In fact, the exact same layout is possible with the only additional component being a graphics card inserted into one of the PCI slots of the computer.
However, this brings some disadvantages, which are again linked to memory management. Here, the image data of the camera are streamed into the memory of the root complex of the computer, from where they have to be transferred to the graphics card memory. This takes time which may quickly become relevant for high-speed applications.
A possible way to circumvent this issue is to use a camera which directly streams the raw pixel data into the memory of the graphics card (see Fig. A.19). This can be done by using a camera with a PCIe interface and the GPUDirect RDMA interface by NVIDIA [127] or the DirectGMA interface by AMD, which are both designed to stream data between PCIe devices while avoiding unnecessary data copies.

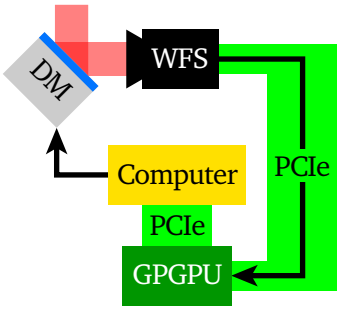

Figure A.19.
Schematic of a GPGPU based CU composition with a camera directly connected to the graphics card.

Examples

GPGPU has been used in AO for more than a decade now, where the first full-scale operational systems have been demonstrated in the field of astronomy [128]. Back then, graphics cards did not deliver the processing power that was necessary to run a loop with thousands of actuators with a multi-kHz frame rate and the systems were designed to run un multiple machines simultaneously.
Nowadays, even consumer grade graphics cards deliver the processing power to do so on a single machine, which made open-source software solutions viable. One example for this is the Compute And Control for Adaptive Optics (CACAO) real-time control software package, which offers full AO loops using NVIDIAs GPGPU interface ”CUDA” on a real-time linux machine [33].
This software offers high throughput: for a PWFS with a resolution of 120x120 and a DM of 2000 actuators, 3.5 kHz with a latency of less than 1.2 ms have been achieved on a computer with two consumer grade graphics cards (NVIDIA GTX 980Ti), which is described in the original application.
However, CACAO has also been used for an SHS, delivering a frame rate 1.8 kHz for an 11x11 SHS [129].

Strengths

The strength of GPGPU based CUs is that this technology spans a bridge between flexibility and high throughput. As the main application is still a piece of software running on the CPU of the computer, one has in principle the same freedom to implement functionality as for CPU based CUs. The main restrictions here are the choice of the programming language, which must not introduce large overheads after compiling in order to not slow down the execution, and the choice of the operating system.

Weaknesses

These restrictions also pose the main downside of the GPGPU approach. While the hardware supports high throughput, the software implementation can still hinder the required speed. Without a deep understanding about the memory management of the root complex and the graphics card, it is likely to accidentally trigger unnecessary memory-transfer- or copy operations which form a bottleneck for the data streams. This also means that the software architecture has to acknowledge the application needs from the ground up as using GPGPUs as an ”on-top” extension for another application will likely mean that the data flow is not optimized accordingly.
Aside from the implementation itself, it is beneficial to run the system on a real-time capable operating system, which greatly narrows the choice of the development environment. Together with the need of a dedicated software architecture, this reflects that building a GPGPU CU requires focusing much more on the structure of the implementation than for CPU based systems, which may require a trade-off in the usability in terms of system setup and user interface.

A.3.3. FPGA

The last established option for AO CUs is based on FPGAs. These are integrated circuits that feature a large array of logic gates that can be programmed by a computer, i.e. a pre-compiled logic can be uploaded to the chip. Using such an FPGA to directly receive the camera input from the WFS, evaluating it, calculating the control signals for the DM and sending them to the driver has the big advantage of avoiding the trouble of memory management on a computer as a whole, removing the primary bottleneck of CPU and GPGPU based CUs. Furthermore, FPGAs inherently offer a high amount of parallelization. These properties offer the capacity for extreme speed and ultra low latencies. Once set up, FPGA based systems deliver the highest performance for real-time AO.

Setup

The setup of an FPGA based AO loop differs from the ones discussed earlier as no computer is involved in the primary loop: in this scenario, the digital camera directly connects to the FPGA circuit, usually bypassing a frame grabber as the interface protocol (e.g. Camera Link HS) is directly implemented on the FPGA. This enables starting the calculations on the raw pixel data in real-time without having to wait for the full camera frame to be transferred. The FPGA then performs all necessary computations and feeds the control values to the DM driver. This step either happens directly ("Option 1" in Fig. A.20) or via a computer ("Option 2" in Fig. A.20), depending on the interface of the DM driver.
In such a setup, the computer does not take part in the control directly. Its purpose is limited to uploading the pre-compiled application onto the FPGA on startup (e.g. via a PCIe bus), communicating user commands, receiving data for logging- or visualization purposes, and perhaps interfacing with the DM driver if necessary.

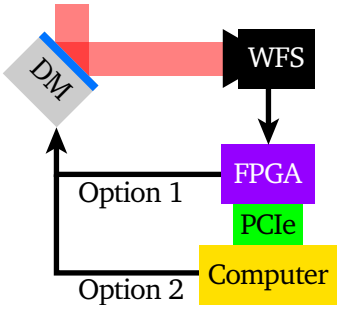

Figure A.20.
Schematic of an FPGA based CU composition. The DM is either fed via the FPGA directly or via the computer.

Examples

One approach from astronomy that demonstrates the capacities of the FPGA technology in AO was published in 2018, where the authors presented a scalable FPGA system for upcoming large telescopes [130, 131]. Here, the FPGA software was implemented in a modular way, allowing easy adaption for different AO scales. At the time of publication, the speed of their implementation was limited by the memory bandwidth that is necessary to transfer the large control matrices of extreme AO systems that can reach hundreds of megabytes. However, it is expected to reach an FPGA-internal latency (without camera and DM) of 100 µs for an SHS with 50x50 sub-apertures with upcoming memory technology in the near future. This shows that the FPGA technology has the potential to outrun the speed of cameras and DMs, allowing for multi-kHz operation even for high spatial resolution.
In 2020, A. Rukosuev et al. published details on another FPGA based AO system using an SHS and a bimorph piezoelectric DM that was intended for free-space optical communication [132]. They achieved a repetition rate of 1.5 kHz, which is well suitable to compensate for atmospheric turbulence. The choice of hardware is comparable to what one may find in a HI laser system, showing that this approach would be suitable in this field as well.

Strengths

The examples show quite impressively that the FPGA technology offers unprecedented speed and low latencies, which makes it especially suitable for demanding real-time AO applications. The possibility to operate an AO loop with a repetition rate of multiple kHz and latencies well below the millisecond regime offers the possibility to correct for aberrations introduced by strong turbulence or vibrations.

Weaknesses

The downside of FPGAs is that they are notoriously hard to program. This is due to two reasons: first, programming an FPGA requires very specialized knowledge as the way these devices work strongly differ from the logic one usually applies when programming computer applications in higher languages. Second, the compile times to translate the code into a binary that can be uploaded to the FPGA are extremely long and can span from minutes for even the most basic ”hello world” code to hours for advanced applications. This makes testing and debugging a time-consuming process.
These two reasons also mean that the code is very hard to maintain, especially by third party developers. If an AO system shall be set up in a facility that will be operated over several decades, this can be a game-breaker for FPGA.
Furthermore, the application code is often hardware dependent over large portions. Code that runs on one system can therefore not be adapted easily for another, which hinders open-source projects on this topic.
Last, the logic capacities of FPGAs are limited, i.e. only a certain amount of fixed computation steps can be programmed onto the circuit. This means that there is not much room for advanced AO schemes such as elaborated state-machines for response recording, predictive control or machine-learning. Implementing designs like these would mean to switch from one to several FPGAs, which again increases the complexity of the system and makes its maintenance difficult.

Conclusion

In this chapter, I reviewed different AO technologies in terms of WF metrology, DMs or CUs that are available today and are either established in the context HI laser systems or of interest in one or the other way. All these technologies feature unique characteristics, which determine the applications that can benefit from each of them.
With that knowledge, it is much easier to judge what technology may be appropriate for a certain application. Likewise, it helps to understand the design considerations for the AO loop I implemented in this work, where the application clearly determines what technology to use (see chapter 5).

B. Diffraction Theory

As I mentioned in the main text, the focal spot of Petawatt High-Energy Laser for heavy Ion eXperiments (PHELIX) cannot be easily simulated from the measured Near Field (NF) data. This is due to the incompleteness of the measurements and a high computational effort, which makes running a large set of simulations, e.g. for every shot of an experiment, unfeasible. In order to understand this, some knowledge on the associated diffraction theory is needed.
In this appendix, I summarized the underlying theory that is needed for a realistic calculation for the focal spot of a High Intensity (HI) laser system. At the end of this chapter, it will be obvious that first, a full frequency-resolved measurement of the WaveFront (WF) would be necessary for the calculation, second, that the calculation cannot be done with the computationally cheap Fraunhofer diffraction, and third, that a the numerical load is extreme if realistic results shall be obtained.
I will start this chapter with a brief summary on waves in the ElectroMagnetic (EM) field in section B.1 and will use the wave equation to outline the derivation of scalar diffraction formulae in section B.2, which will yield the Fraunhofer diffraction integral. Analogous to that, I will briefly touch on vectorial diffraction in section B.3. Here, I will also provide a simple comparison between the Fraunhofer diffraction and a spherical vector diffraction integral, which will clearly show that the approximations of the Fraunhofer Diffraction do not hold true anymore for steep focusing, as it is often done in HI lasers. In order to estimate the computational effort that comes with a realistic calculation, I will discuss the diffraction of polychromatic light in section B.4 and conclude by reviewing some limits of the numerical evaluation of these integrals in section B.5.

B.1. Electromagnetic Waves

In order to understand the different diffraction models and their limits, one has to be aware of some rudimentary properties of light.
While geometric optics yield good approximations for a wide set of scenarios, this simple model omits the more fundamental model of light as a wave in the EM field. These limitations become apparent as soon as the scale lengths of the considered scenario approach the WaveLength (WL) of the underlying EM wave.
The behavior of the EM field is described by *Maxwell's Equations* [133, p. 20]:

$$\nabla\times\vec{E}=-\frac{\partial\vec{B}}{\partial t} \tag{B.1}$$

$$\nabla\times\vec{B}=\mu_0\vec{j}+\frac{1}{c^2}\frac{\partial\vec{E}}{\partial t} \tag{B.2}$$

$$\nabla\cdot\vec{E}=\frac{\rho}{\epsilon_0} \tag{B.3}$$

$$\nabla\cdot\vec{B}=0. \tag{B.4}$$

Here, $\vec{E}$ is the electric and $\vec{B}$ the magnetic field vector, $\vec{j}$ the electric current, ρ the charge density of the surrounding medium and c the speed of light in vacuum, while μ_0 and ϵ_0 are the vacuum magnetic permeability and the vaccum electric permittivity, respectively. In the context of focus formation in HI lasers, one are generally interested in the EM field in vacuum. Therefore, $\rho=0$, $\vec{j}=\vec{0}$. From here, the

electric wave equations for vacuum can be derived from (B.1) and (B.2):

$$\begin{aligned}
\nabla\times \vec{E} &= -\frac{\partial \vec{B}}{\partial t} \\
\Leftrightarrow \nabla\times(\nabla\times \vec{E}) &= -\frac{\partial}{\partial t}(\nabla\times \vec{B}) \\
\Leftrightarrow \nabla(\nabla\cdot \vec{E}) - \nabla^2\vec{E} &= -\frac{1}{c^2}\frac{\partial^2 \vec{E}}{\partial t^2} \\
\Leftrightarrow \nabla^2\vec{E} &= \frac{1}{c^2}\frac{\partial^2 \vec{E}}{\partial t^2}.
\end{aligned} \tag{B.5}$$

The wave equation for the magnetic field can be derived in an analogue manner.
If one assumes a monochromatic wave of the form $\vec{M}(\vec{r},t) = \vec{m}(\vec{r})\cdot e^{-i\omega t}$, where $\vec{r}$ is the location, t the time and ω the angular frequency of the wave, one obtains the time-independent wave equation from (B.5):

$$\begin{aligned}
\nabla^2\vec{M} &= \frac{\partial^2 \vec{M}}{\partial t^2} = -\frac{\omega^2}{c^2}\vec{M} \\
\Leftrightarrow \nabla^2\vec{m} &= -k^2\vec{m},
\end{aligned} \tag{B.6}$$

where k is the wave number associated with ω.

B.2. Scalar Diffraction Integrals

In order to calculate scalar diffraction, one has to formulate the full diffraction integral using the Huygens-Fresnel principle first. In this case, an isotropic spherical scalar wave is used as elemental wave. A corresponding wave, starting from the origin, can be expressed via

$$E(\vec{r},t) = E_0\frac{e^{i2\pi\left(\frac{r}{\lambda}-\frac{\lambda}{c_0}t\right)}}{r} = E_0\frac{e^{i(kr-\omega t)}}{r}.$$

Here, $\vec{r}$ is the position vector, r its absolute value and thus the distance to the wave center and λ and c_0 the WL and speed of light in vacuum, respectively. Expressed in simpler terms, one can use the wave number k and the circular frequency ω instead.

B.2.1. General Diffraction

Formulating the diffraction integral from the qualitative Huygens-Fresnel principle is a non-trivial task that has been done by Kirchoff using the Greens theorem [134]. Only the main steps will be briefly covered here to roughly understand the underlying principles.
The main idea is that the wave equation

$$\nabla^2 E = \frac{1}{c^2}\frac{\partial^2 E}{\partial t^2}, \tag{B.7}$$

which is the scalar equivalent of equation (B.5), can be solved for the interior for a compact volume V if the values and the derivative of the electric field on its boundary Σ are known.
Here, one can use the Greens theorem, which is an adaptation of Gauss's theorem (also known as

divergence theorem), where the vector field is expressed in terms of two scalar functions f and g which are at least two times continuously differentiable inside V:

$$\iiint_V f\nabla^2 g - g\nabla^2 f dV = \iint_\Sigma f\frac{dg}{d\hat{n}} - g\frac{df}{d\hat{n}}d\Sigma. \tag{B.8}$$

Now we demand f and g to be solutions to the wave equation (B.7) for the same wavenumber k. Analogous to the vectorial linear wave equation B.6, this means

$$\begin{aligned}\nabla^2 f e^{-i\omega t} &= \frac{1}{c^2}\frac{\partial^2 f e^{-i\omega t}}{\partial t^2} = -k^2 f e^{-i\omega t},\\ \Rightarrow \nabla^2 f &= -k^2 f\end{aligned}$$

for f and the same for g. Doing this, the left-hand side of equation (B.8) vanishes:

$$\iint_\Sigma f\frac{dg}{d\hat{n}} - g\frac{df}{d\hat{n}}d\Sigma = 0. \tag{B.9}$$

From here, one can directly derive the diffraction integral by choosing a special ansatz for these two functions: let f be the complex amplitude of the field and g a static spherical wave, centered at a point $\vec{x}'$ inside V:

$$\begin{aligned} f &= E,\\ g &= \frac{e^{ik\rho}}{\rho} \qquad \text{with } \rho = \|\vec{x} - \vec{x}'\|.\end{aligned}$$

Here, E will serve as the known field on the boundary Σ and g represents the sampling position inside V, for which the field strength shall be calculated.
The Greens theorem, however, only holds valid for functions that are continuously differentiable twice within V, yet g has a singularity in at $\vec{x}'$. Therefore, one has to exclude $\vec{x}'$ from the volume by introducing another boundary σ around the singularity, with inward facing normal vector (see Fig. B.1). Choosing a spherical boundary for σ is a good choice as solving the integral becomes a lot easier later on.
Having two closed boundaries, equation (B.9) is split into two components and becomes

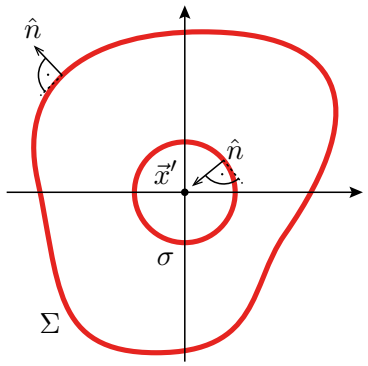

Figure B.1.: Schematic of the region definitions used for the derivation of the Kirchhoff diffraction formula.

$$\iint_\Sigma f\frac{dg}{d\hat{n}} - g\frac{df}{d\hat{n}}d\Sigma + \iint_\sigma f\frac{dg}{d\hat{n}} - g\frac{df}{d\hat{n}}d\sigma = 0.$$

Now, by taking the limit of the radius of σ towards 0, one finds that

$$\begin{aligned} 0 &= \iint_\Sigma f\frac{dg}{d\hat{n}} - g\frac{df}{d\hat{n}}d\Sigma + 4\pi f|_0,\\ \Leftrightarrow f\big|_0 &= -\frac{1}{4\pi}\iint_\Sigma f\frac{dg}{d\hat{n}} - g\frac{df}{d\hat{n}}d\Sigma,\\ \Rightarrow E(\vec{x}) &= -\frac{1}{4\pi}\iint_\Sigma E\frac{d}{d\hat{n}}\frac{e^{ik\rho}}{\rho}\bigg|_{\vec{x}'} - \frac{e^{ik\rho}}{\rho}\frac{dE}{d\hat{n}}\bigg|_{\vec{x}'} d\Sigma,\end{aligned} \tag{B.10}$$

which is the Kirchhoff diffraction integral.

B.2.2. Diffraction for Apertures

In practice, the general Kirchhoff diffraction integral is often used in specialized versions as the electric field is rarely known for the whole surface Σ. A frequently used scenario is the diffraction of an aperture where one knows the electric field over a two dimensional area of limited size and assumes that the electric field is zero elsewhere.
Here, one proposition for the geometry of the problem is to choose Σ as a sphere around the sampling position $\vec{x}$, which is clipped by the plane of the aperture [46, sec. 8.3] (see Fig. B.2).
Chosing this geometry, Σ can now be divided into three parts: A, the aperture, B, the screen surface and C, the spherical part. Therefore,

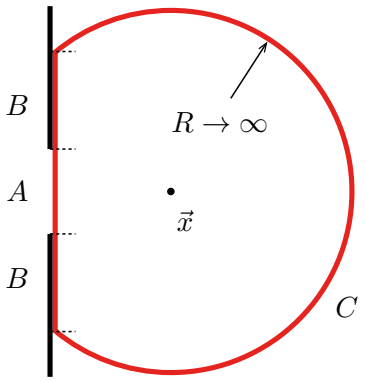

Figure B.2.: Schematic of the boundary definitions in the Kirchhoff diffraction on an aperture.

$$E(\vec{x}) = -\frac{1}{4\pi}\left[\iint_A + \iint_B + \iint_C\right]\left\{E\frac{d}{d\hat{n}}\frac{e^{ik\rho}}{\rho}\bigg|_{\vec{x}'} - \frac{e^{ik\rho}}{\rho}\frac{dE}{d\hat{n}}\bigg|_{\vec{x}'}\right\} d\Sigma.$$

Now, further simplifications can be made for each part of Σ.
For the integral over C and the radius R of the sphere, one can consider the limit of $R \to \infty$. Here, it can either be shown that the integral shrinks with R^{-2} and can thus be neglected in the limit, or a temporal argumentation can be used: if one switches on the field on Σ at a point t_0 in time, the influence of any field on C can only reach $\vec{x}$ with the speed of light and will therefore not influence the field at $\vec{x}$ until $t_0 + \Delta t = t_0 + R/c_0$. While switching on and off the field on Σ means departing from the monochromaticity of the problem, it becomes monochromatic at the limit of $R \to \infty$ again as long as the field on Σ persists for Δt. This implies that the integral part over C can be neglected for sufficiently large R [46, sec. 8.3].
The integral over B is more problematic as the electric field cannot necessarily be neglected for physical apertures. Depending of different factors, such as the material, the geometry of the aperture and the interaction of the electric field on the edges of A, the contribution of B to the overall integral may not be neglectable or even small compared to A.
Still, it is reasonable to assume that $E \approx 0$ on B except for regions in direct proximity to the the boundary to A. For sufficiently large apertures A (much larger than the WL), the field at $\vec{x}$ can therefore be approximated only by the integral over A.
Using this approach, one can estimate any field behind an aperture. One scenario that is often encountered is the diffraction of a disturbed plane wave over the aperture, which may represent a laser beam. Here, the electric field can be approximately described as

$$E(\vec{x}) = A(\vec{x}) \cdot e^{i(\vec{k}\cdot\vec{x} + \phi(\vec{x}))}.$$

A is the spatial amplitude of the wave and ϕ the spatial phase disturbance. For simplicity, the temporal oscillation was excluded as it does not influence this consideration. The derivative with respect to the

normal vector $\hat{n}$ at the aperture is then

$$\begin{aligned}\frac{\partial E}{\partial \hat{n}} &\approx \frac{\partial A}{\partial \hat{n}} \cdot e^{i(\vec{k}\cdot\vec{r})+\phi} + A(\vec{x}) \cdot i\vec{k} \cdot \vec{n} \cdot e^{i(\vec{k}\cdot\vec{x})+\phi} \\ &\approx A(\vec{x}) \cdot ik \cos(\vec{k}, \hat{n}) \cdot e^{i(\vec{k}\cdot\vec{x})+\phi(\vec{x})} = ikE(\vec{x}) \cos(\vec{k}, \hat{n}).\end{aligned}$$

Here, two assumptions have been made: first, the amplitude A was assumed to be approximately constant in normal direction to the aperture, which holds true as long as A varies slowly compared to the WL and the direction of propagation is parallel to $\hat{n}$ down to the small angle approximation. Secondly, $\phi(\vec{x})$ does not change the derivative of the oscillating term significantly by assumption, i.e. its own derivative over the aperture does not change the angle of the wave over the small angle approximation. Both assumptions are reasonable for a laser beam with a diameter of several hundreds to thousands of WLs, propagating a typical optical system.
The second term of equation (B.10) that has to be determined is the derivative of the spherical wave, originating in the sampling position $\vec{x}$, at the aperture at $\vec{x}'$ with respect to its normal vector $\hat{n}$:

$$\begin{aligned}\frac{\partial e^{ik\rho}/\rho}{\partial \hat{n}} &= \left(ik\frac{e^{ik\rho}}{\rho} - \frac{e^{ik\rho}}{\rho^2}\right) \cdot \frac{\vec{x}' - \vec{x}}{\rho} \cdot \hat{n} \\ &= \left(ik\frac{e^{ik\rho}}{\rho} - \frac{e^{ik\rho}}{\rho^2}\right) \cdot \cos(\vec{x}' - \vec{x}, \hat{n}) \\ &\approx ik\frac{e^{ik\rho}}{\rho} \cdot \cos(\vec{x}' - \vec{x}, \hat{n}) \qquad \text{for } \rho \gg \lambda.\end{aligned}$$

The full diffraction integral therefore is

$$\begin{aligned}E(\vec{x}) &= -\frac{1}{4\pi}\iint_A E \left.\frac{\partial e^{ik\rho}/\rho}{\partial \hat{n}}\right|_{\vec{x}'} - \frac{e^{ik\rho}}{\rho}\left.\frac{\partial E}{\partial \hat{n}}\right|_{\vec{x}'} d^2\vec{x}' \\ &\approx -\frac{1}{4\pi}\iint_A E(\vec{x}') \cdot ik\frac{e^{ik\rho}}{\rho} \cdot \cos(\vec{x}' - \vec{x}, \hat{n}) - \frac{e^{ik\rho}}{\rho} \cdot ikE(\vec{x}') \cos(\vec{k}, \hat{n}) d^2\vec{x}' \\ &= -\frac{i}{2\lambda}\iint_A E(\vec{x}')\frac{e^{ik\rho}}{\rho}\Big(\cos(\vec{x}' - \vec{x}, \hat{n}) - \cos(\vec{k}, \hat{n})\Big) d^2\vec{x}'. \qquad \text{(B.11)}\end{aligned}$$

B.2.3. Fresnel and Fraunhofer Diffraction

The cosine terms in equation (B.11) are called inclination factors and scale the contributions of each elemental wave with respect to the local inclination of the aperture surface. For large distances, where the inclination factors do not change considerably for any position $\vec{x}'$ within the aperture, these factors can be reduced to $2\cos(\theta)$, where θ can be thought of as a unified inclination angle.
For distances which satisfy this approximation, ρ can be approximated by

a second order Taylor expansion $\qquad \rho \approx (z - z') + \dfrac{(x - x')^2 + (y - y')^2}{2(z - z')}$

and the Fresnel approximation $\qquad \dfrac{1}{\rho} \approx \dfrac{1}{z - z'}$

as well [46, sec. 8.3], yielding the Fresnel diffraction integral:

$$E(\vec{x}) \approx -\frac{i\cos(\theta)}{\lambda}\frac{e^{ik(z-z')}}{z - z'}\iint_A E(\vec{x}')e^{ik\frac{(x-x')^2+(y-y')^2}{2(z-z')}} d^2\vec{x}'. \qquad \text{(B.12)}$$

This diffraction integral is of great use when considering the diffraction after a focusing optic. Any lens will approximately modify $\phi(\vec{x}')$ such that

$$E(\vec{x}') \approx E_0(\vec{x}') \cdot e^{-ik\frac{x'^2+y'^2}{2f}}, \tag{B.13}$$

where E_0 is the electric field before the lens and f the focal length of the lens. Inserting this into equation (B.12) while setting $z = 0$ and $z - z' = f$, one gets

$$\begin{aligned} E(\vec{x}) &\approx -\frac{i\cos(\theta)}{\lambda}\frac{e^{ikf}}{f}\iint_A E_0(\vec{x}')e^{-ik\frac{x'^2+y'^2}{2f}}e^{ik\frac{(x-x')^2+(y-y')^2}{2f}}d^2\vec{x}' \\ &= -\frac{i\cos(\theta)}{\lambda}\frac{e^{ikf}}{f}\iint_A E_0(\vec{x}')e^{ik\frac{x^2-2xx'+y^2-2yy'}{2f}}d^2\vec{x}' \\ &= -\frac{i\cos(\theta)}{\lambda}\frac{e^{ik\frac{x^2+y^2+2f^2}{2f}}}{f}\iint_A E_0(\vec{x}')e^{-ik\frac{xx'+yy'}{f}}d^2\vec{x}', \end{aligned} \tag{B.14}$$

which is called the Fraunhofer diffraction integral. One can observe two things:
First, the limit of $f \to \infty$ makes the influence of the lens vanish in equation (B.13). Therefore, the Fraunhofer diffraction integral is an approximation for the Fresnel diffraction integral at very large distances, which is why $E(\vec{x})$ is often called the Far Field (FF).
Second, equation (B.14) takes the form of a two-dimensional Fourier transform with additional factors. Therefore, the FF can be expressed as

$$E(\vec{x}) = -\frac{i\cos(\theta)}{2\pi\lambda}\frac{e^{ik\frac{x^2+y^2+2f^2}{2f}}}{f} \cdot (\mathcal{F}E_0)\left(\frac{k\vec{x}}{f}\right), \tag{B.15}$$

where $\mathcal{F}$ denotes the two-dimensional unitary angular frequency Fourier transform, which reveals that the focal spot distribution in this approximation can be interpreted as the spatial frequency decomposition of the incident electric field.

B.3. Vectorial Diffraction Integrals

Issues with the scalar diffraction model have already been touched in section B.2.2, where the conditions at the edge of an aperture and beyond could not be included. This leads to strong demands on the scenarios that can still be approximated with the scalar diffraction theory. A more complete model for diffraction of light is therefore needed to describe its propagations through smaller apertures and with larger angles.
In the following section, the vectorial diffraction model of Stratton and Chu [135] shall be described briefly, followed by a widely used simplification for spherical geometries in section B.3.2 and a comparison with the scalar model in section B.3.4.

B.3.1. General Diffraction

Stratton and Chu derive the diffraction integral in a similar manner as Kirchhoff. However, they chose a vectorial form and selected an ansatz to explicitly satisfy maxwell equations: In the scalar theory, the considered vectorial field is described in terms of two scalar functions f and g that both fulfill the time-independent wave equation. Here, Stratton and Chu choose a vector field $\vec{P} \times \nabla \times \vec{Q}$ composed

itself of two other vector fields $\vec{P}$ and $\vec{Q}$ which both are two times continuously differentiable inside a volume V. The Greens theorem in this case is

$$\begin{aligned}&\iiint_V \vec{Q} \cdot \nabla \times \nabla \times \vec{P} - \vec{P} \cdot \nabla \times \nabla \times \vec{Q}\, dV \\ &= \iint_\Sigma \left(\vec{P} \times \nabla \times \vec{Q} - \vec{Q} \times \nabla \times \vec{P} \right) \cdot \hat{n}\, d\Sigma. \end{aligned} \tag{B.16}$$

Now, similar as before, a special ansatz for $\vec{P}$ and $\vec{Q}$ can be made, choosing $\vec{P} = \vec{E}$ to be the electric field and $\vec{Q}$ as a spherical wave with arbitrary polarization $\vec{a}$ of the form

$$\vec{Q} = \vec{a} \cdot \psi \qquad \text{with } \psi = \frac{e^{ik\rho}}{\rho}. \tag{B.17}$$

Now, the vectorial Greens theorem (equation (B.16)) can be expanded accordingly, again excluding the singularity of the spherical wave in equation (B.17) by a small sphere around its origin. For scenarios, where charge and current are zero for the entire volume V and its surface Σ, the electric field at any point $\vec{x}$ inside the volume is then

$$\vec{E}(\vec{x}) = \frac{1}{4\pi} \iint_\Sigma i\omega(\hat{n} \times \vec{B})\psi + (\hat{n} \times \vec{E}) \times \nabla\psi + (\hat{n} \cdot \vec{E})\nabla\psi\, d^2\vec{x}', \tag{B.18}$$

which I will refer to as the vectorial diffraction integral from now on. Note that the unit vector $\vec{a}$ has vanished, which means that the polarization at the sampling position is, as expected, entirely defined by the surrounding electric and magnetic field. The three terms in the integral can be considered analogous to an electric current density, a magnetic current density and an electric charge distribution (from left to right) on Σ, which are entirely determined by the considered EM field.
As equation (B.18) is derived from the Greens theorem, it only holds true for an EM field that is two times continuously derivable in V and on Σ, which means that scenarios with discontinuities cannot be considered without making errors. However, Stratton and Chu find that discontinuities, e.g. the edge of an aperture, can be ”repaired” by introducing real charge and current along these edges on Σ and including these into the diffraction integral. While this consideration is elegant and enables to calculate a much wider range of scenarios, this work mainly deals with freely propagating laser beams without contact to physical apertures. Therefore, discontinuous electric fields go beyond the scope of this work and are therefore not included in this summary.

B.3.2. Diffraction for Spherical Geometry

As for the scalar model, it is useful to specialize the general diffraction integral, given by equation (B.18), to cases that are commonly encountered and introduce simplifications accordingly. In the frame of this work, the formation of the focal spot is of particular interest. A focal spot forms when an approximately spherical wave converges to it’s center. Therefore, equation (B.18) can be rewritten for the spherical case, following Kim et al. [136].
First, one has to specialize the vectorial diffraction integral to the FF case. In particular, the first Maxwell Equation (B.1) implies for a plane wave that

$$\vec{B} = \frac{1}{c}\hat{k} \times \vec{E},$$

where $\hat{k}$ is the unit vector of propagation. Furthermore, the gradient of the spherical wave can be simplified for large distances:

$$\begin{aligned}\nabla\psi = \nabla\frac{e^{ik\rho}}{\rho} &= \left(ik\frac{e^{ik\rho}}{\rho} - \frac{e^{ik\rho}}{\rho^2}\right)\cdot(-\hat{\rho}) \\ &\approx -ik\psi\hat{\rho} \qquad \text{for } \rho \gg \lambda \text{ with } \hat{\rho} = \frac{\vec{\rho}}{\rho} \text{ and } \vec{\rho} = \vec{x} - \vec{x}'.\end{aligned}$$

With that, the integral becomes

$$\begin{aligned}\vec{E}(\vec{x}) &= \frac{1}{4\pi}\iint_\Sigma i\omega(\hat{n}\times(\frac{1}{c}\hat{k}\times\vec{E}))\psi + (\hat{n}\times\vec{E})\times\nabla\psi + (\hat{n}\cdot\vec{E})\nabla\psi\, d^2\vec{x} \\ &= \frac{ik}{4\pi}\iint_\Sigma \psi\left[\hat{n}\times(\hat{k}\times\vec{E}) - (\hat{n}\times\vec{E})\times\hat{\rho} - (\hat{n}\cdot\vec{E})\hat{\rho}\right] d^2\vec{x}. \qquad \text{(B.19)}\end{aligned}$$

In order to introduce simplifications for the spherical case, it is useful to expand the cross products:

$$\vec{E}(\vec{x}) \approx \frac{ik}{4\pi}\iint_\Sigma \psi\left[(\hat{n}\cdot\vec{E})\hat{k} - (\hat{n}\cdot\hat{k})\vec{E} + (\hat{\rho}\cdot\vec{E})\hat{n} - (\hat{\rho}\cdot\hat{n})\vec{E} - (\hat{n}\cdot\vec{E})\hat{\rho}\right] d^2\vec{x}.$$

Now, choosing a reference sphere for Σ, which is approximately matched by the incoming EM wave (see Fig. B.3), it becomes obvious that $\hat{n} \approx \hat{\rho} \approx -\hat{k}$ and thus $\vec{E}\cdot\hat{n} \approx \vec{E}\cdot\hat{\rho} \approx \vec{E}\cdot\hat{k} \approx 0$.
By switching to a spherical coordinate system, centered in the center of the reference sphere, the *Debye-Wolf integral* is formed [137, 138]:

$$\vec{E}(\vec{x}) \approx \frac{ik}{2\pi}\iint_\Sigma \vec{E}e^{i\vec{k}\cdot\vec{\rho}}\, d^2\Omega \qquad \text{(B.20)}$$
$$\text{with } d^2\Omega = \sin\theta\, d\theta d\phi.$$

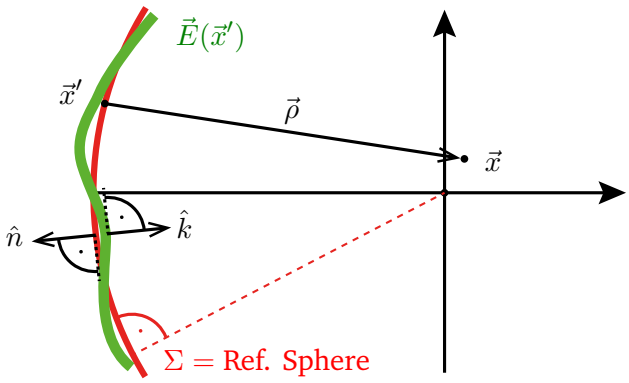

Figure B.3.: Schematic of Debye-Wolf integral geometry.

As the denominator of the once spherical waves vanished in this form due to the scaling factor of the spherical coordinates ($\rho \approx \|\vec{x}'\|$), the infinitesimal components of the integral can be thought of as plane waves. Wolf reflected that by introducing $\vec{k}\cdot\vec{\rho}$ in the exponential term, taking into account the impact of the local wave slope on the phase of the field at the sampling position. This approximation yields good results as long as the Fresnel number of the considered system is much greater than unity [139], which is given as soon as macroscopic systems are considered.

B.3.3. Calculation of the Electric Field on the Reference Sphere

Using the spherical geometry, the evaluation Debye-Wolf integral itself is straight forward. However, the electric field on the reference sphere has to be known in order to evaluate it. For the scalar Fresnel and Fraunhofer diffraction, one does just add a quadratic phase term to the incoming electric field for this purpose. In the vectorial case, this is not possible as the deflection of the polarization has to be taken into account as soon as the field is curved by the focusing optic. In general, determining the electric field on the reference sphere is a challenging problem that is best solved by raytracing [136]. In order to introduce the main concepts that I use for a comparison with the Fraunhofer diffraction later on, I will only discuss the focusing of a beam using an ideal lens here.
The calculation of the electric field on the reference sphere can be done in four main steps:

1. Projection of spherical coordinates onto the lens plane
2. Calculation of the deflection angle of the associated beam
3. Rotation of the meridional electric field vector around the deflection axis
4. Amplitude correction by multiplication with the projected area ratio

Each step is described in the following.

Step 1: projection of spherical coordinates onto the lens plane For an ideal lens with focal point $\vec{f}$, coinciding with the center of the and the reference sphere Σ of the Debye-Wolf integral, one can do an approximate ray tracing from $\vec{f}$ through any point $\vec{x}'(\theta,\phi)$ to find the corresponding projection $\vec{P}(\vec{x}'(\theta,\phi))$ on the plane of the lens (see Fig. B.4). As the electric field if often known on a plane, this projection mapping can be used to sample the field in terms of the parametrization of Σ in terms of the angles θ and ϕ. Note, however, that this is only possible as long as the wave vector of the field is approximately normal to the lens plane.

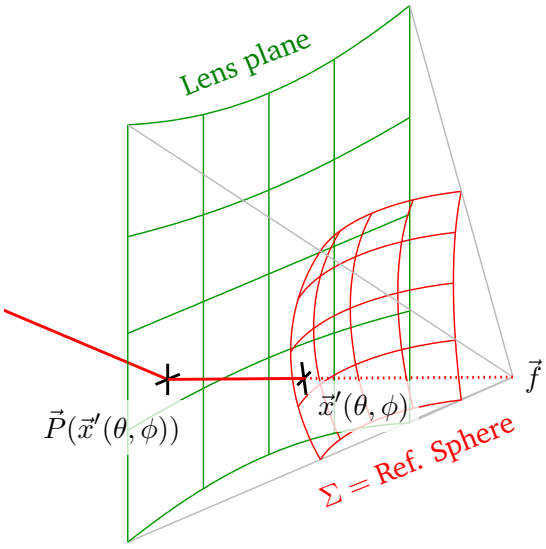

Figure B.4.: Scheme of the projection of the reference sphere onto the lens plane.

Step 2: calculation of the deflection angle of the associated beam For the given scenario, the deflection angle ψ of any ray of light through $\vec{P}(\vec{x}'(\theta,\phi))$ corresponds to the angle between the normal of the lens plane and the connection between $\vec{P}(\vec{x}'(\theta,\phi))$ and $\vec{f}$. The deflection can be expressed as a rotation around the deflection axis $\vec{\psi}$ about the angle ψ (see Fig. B.5).

Step 3: rotation of the meridional electric field vector around the deflection axis With the deflection axis $\vec{\psi}$ known, one can calculate the deflection of the incoming electric field $\vec{E}$ next. To do so, one has to split $\vec{E}$ into a sagittal component $\vec{E}_{\parallel}$ (parallel to $\hat{\psi}$) and a meridional component $\vec{E}_{\perp}$ (perpendicular to $\hat{\psi}$). This is depicted in Fig. B.6. Naturally, only the meridional part has to be rotated by ψ around $\hat{\psi}$. Now, the total deflected field vector is given by $\vec{E}_{\Sigma} = \vec{E}_{\parallel} + \vec{E}_{\Sigma,\perp}$.

Step 4: amplitude correction by multiplication with the projected area ratio The last step in the calculation is to correct for the amplitude of the electric field that changes due to the projection. As both Σ and its projection onto the lens plane are parametrized by the angles θ and ϕ, one can calculate the infinitesimal area elements A_{Σ} and A on each of them, respectively (see Fig. B.7). The ratio M of the two acts as a simple multiplier to correct for the projection. Numerically, this can be done as follows, using the smallest numerically stable value for ϵ:

$$
\begin{aligned}
d\vec{\theta}_{\Sigma} &= \frac{\vec{x}'(\theta+\epsilon,\phi) - \vec{x}'(\theta,\phi)}{\epsilon} \\
d\vec{\phi}_{\Sigma} &= \frac{\vec{x}'(\theta,\phi+\epsilon) - \vec{x}'_{\Sigma}(\theta,\phi)}{\epsilon} \\
d\vec{\theta} &= \frac{\vec{P}(\vec{x}'(\theta+\epsilon,\phi)) - \vec{P}(\vec{x}'_{\Sigma}(\theta,\phi))}{\epsilon} \\
d\vec{\phi} &= \frac{\vec{P}(\vec{x}'(\theta,\phi+\epsilon)) - \vec{P}(\vec{x}'(\theta,\phi))}{\epsilon} \\
A_{\Sigma} &= \|d\vec{\theta}_{\Sigma} \times d\vec{\phi}_{\Sigma}\| \\
A &= \|d\vec{\theta} \times d\vec{\phi}\| \\
M &= \frac{A}{A_{\Sigma}}
\end{aligned}
$$

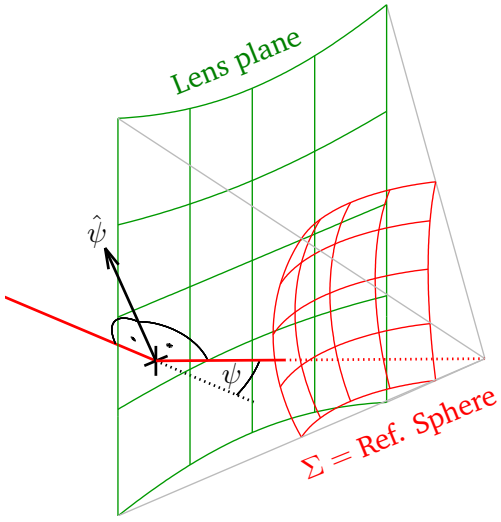

Figure B.5.: Scheme of the deflection angle of the projection.

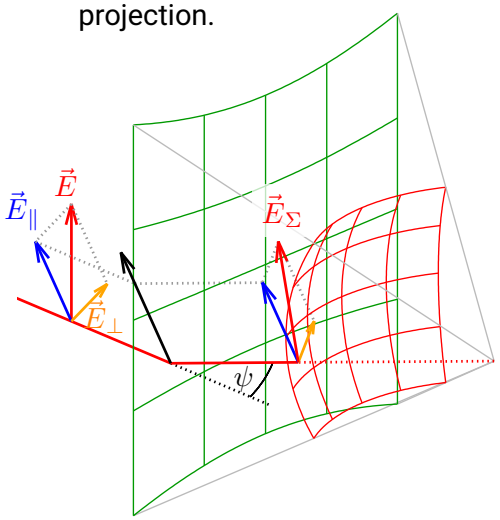

Figure B.6.: Scheme of the deflection of the electric field.

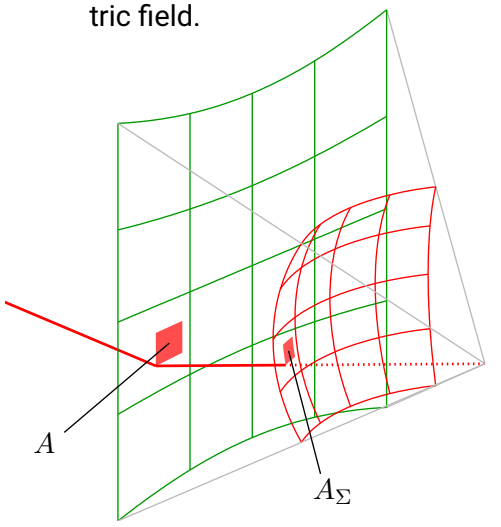

Figure B.7.: Scheme of the projected infinitesimal surface elements.

While the steps mentioned above are sufficient for the simple case of focusing a collimated beam using an ideal lens, real scenarios using e.g. thick lenses and curved mirrors can get much more complicated. The general approach of doing raytracing and projecting the polarization accordingly stays the same, but having a unified formalism can greatly aid calculating the result. A corresponding, more general approach is given by Kim et al. [136], using a matrix formalism to include not only all the steps above but also Fresnel reflection on boundary surfaces.

B.3.4. Comparison Between Scalar and Vectorial Diffraction

In section B.2, the scalar diffraction theory according to Huygens, Fresnel, Kirchhoff and Fraunhofer has been discussed alongside of its limitations. In section B.3, a more general approach in form of vectorial diffraction, following Stratton and Chu was discussed.
In many cases, scalar diffraction yields good approximations, which is useful as the Fraunhofer diffraction integral is very cheap to calculate numerically. However, one needs to be aware of the point where the approximation is not good anymore. This threshold depends on the application and the observables one is interested in.
In the frame of this work, the peak intensity of the focal spot of a focused super-Gaussian or near-tophat beam is the property of interest. In order to visualize the difference between the two diffraction models, I calculated the focal spot of a flattop beam with flat phase, focused by an ideal lens using Fraunhofer diffraction and the Debye-Wolf diffraction integral. For the latter, the electric field on the reference sphere was determined following the procedure described in section B.3.3. The result is shown in Fig. B.8.
From the graph, it becomes apparent immediately that the peak intensities of both integrals match

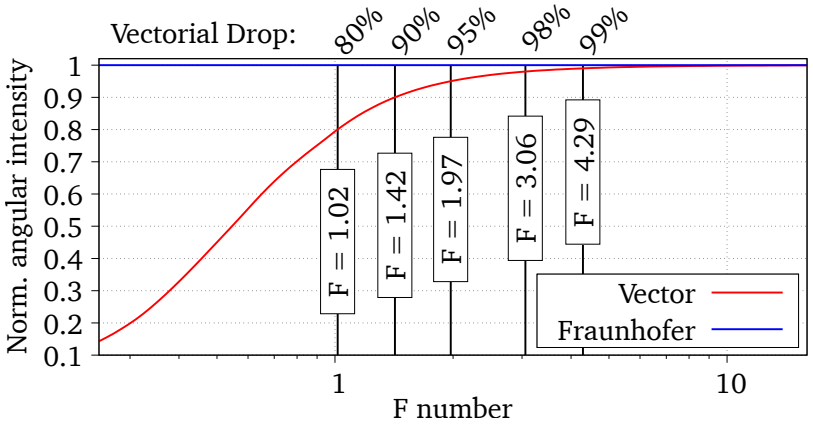

Figure B.8.: The peak angular intensity of a focused flattop beam over the F number for Fraunhofer and vectorial Debye-Wolf diffraction, normalized to the limit of large F numbers. F numbers, where the ratio of the two curves hits certain values, shown above the graph, are indicated.

exactly for large F numbers, but increasingly diverge for steeper focusing angles. Furthermore, the angular intensity of the Fraunhofer diffraction stays constant for any F number, which reflects the use of the small angle approximation.
Where these approximations loose validity, the Debye-Wolf diffraction yields lower angular intensities, taking into account distortions and depolarization at the edges of the beam. From this graph, one can see that for F numbers greater than 5, the error stays below 1 %. The focusing parabola at PHELIX, however, features an F number of up to 1.43 (depending on the beam diameter), which corresponds to

an error of almost 10 %.
This proves that the focal spot at PHELIX can only be calculated correctly by vector diffraction, which is numerically much more expensive than the Fraunhofer diffraction.

B.4. Diffraction of Polychromatic Light

Up to this point, diffraction was only treated for monochromatic light. While this can be approximately fulfilled for single mode Continuous Wave (CW) lasers, pulsed lasers feature broadened spectra. This is driven by Heisenbergs uncertainty principle, which in this case states that the product of the pulse duration and the spectral bandwidth cannot be smaller than a certain value. This value can be different, depending on the definition of the pulse duration and the spectral energy distribution, but in any case, a shorter pulse requires a larger spectral bandwidth.
In order to calculate the diffraction pattern of a polychromatic source, one has to solve the diffraction integral for each frequency in the spectrum and add up the electric fields [46, sec. 8.3]:

$$\vec{E}_{\text{poly}}(\vec{x}) = \int_S A(\nu) \cdot \vec{E}(\vec{x}, \nu) \; d\nu. \tag{B.21}$$

Here, $\vec{E}_{\text{poly}}(\vec{x})$ is the diffracted polychromatic E-field at the position $\vec{x}$, $\vec{E}_{\text{poly}}(\vec{x})$ is the diffracted monochromatic E-field for the frequency ν at $\vec{x}$, calculated using an appropriate diffraction integral, $A(\nu)$ is the complex amplitude of the spectrum, which carries amplitude and phase of the frequencies relative to each other, and S is the spectral domain that contains the entire signal.
However, (B.21) only allows to calculate $\vec{E}_{\text{poly}}$ at a single point in time. In order to retrieve the temporal evolution of the field, one can add a linear phase to the integrand, which effectively shifts the pulse in time:

$$\vec{E}_{\text{poly}}(\vec{x}, \Delta t) = \int_S A(\nu) \cdot e^{i2\pi\nu\Delta t} \cdot \vec{E}(\vec{x}, \nu) \; d\nu. \tag{B.22}$$

Here, Δt is the shift in time that shall be introduced to the pulse. Note that Δt acts on the phase of all frequencies simultaneously. The temporal evolution of the electric field is therefore reflected realistically.

B.5. Numerical Evaluation

The diffraction integrals described in sections B.2, B.3 and B.4 need to be solved numerically in many cases, including intensity estimation or simulating the focal spot as input for further simulations. As this involves discretizing the integrals, one has to be wary about the issues that can be introduced in this step. In this section, the numerical equivalents of the scalar Fraunhofer diffraction integral (equation (B.14)), the vectorial Debye-Wolf integral (equation (B.20)) and the general form of the polychromatic diffraction integral (B.22) will be presented and their properties discussed.

B.5.1. Numerical Calculation of the Fraunhofer Diffraction Integral

As described in section B.2, the Fraunhofer diffraction (equation (B.14)) can be written as a Fourier transform (equation (B.15)). The most widespread numerical equivalent is the Fast Fourier Transform (FFT), using a regular quadratic grid of complex floating point values, commonly with a resolution of $2^n \times 2^n$, where n is a real, positive integer number.

However, the FFT performs a transform into ordinary frequency space (1/m), while equation (B.15) is written for angular frequency space (rad/m). Therefore, the numerical equivalent has to be written as

$$E(\vec{x}) = -\frac{i\cos(\theta)}{2\pi\lambda}\frac{e^{ik\frac{x^2+y^2+2f^2}{2f}}}{f} \cdot (\mathcal{FFT}E_0)\left(\frac{\vec{x}}{\lambda f}\right), \tag{B.23}$$

where $\mathcal{FFT}$ denotes the FFT and λ the considered WL.
The spatial dimensions of the FF distribution can thus be calculated as follows: if the NF grid before the transform spans a width of D, each sample of the FFT spans a width of $1/D$. According to equation (B.23), the actual space spanned by each sample is therefore $\Delta d_{ff} = \lambda f/D$.
This scaling reveals that a small D increases Δd_{ff}, potentially covering most of the FF and thus concealing the structure of the focus entirely. To avoid that, one has to add a padding of zeros around the aperture. A simple estimation for the FF diameter A_{FF} in samples of the FFT grid of resolution N is

$$A_{FF} \approx \frac{N}{A_{NF}}, \tag{B.24}$$

where A_{NF} is the diameter of the aperture in the NF, given in samples. This formula stems from the fact that scaling down the field before the FFT increases the spatial frequencies contained in it and thus enlarges the FF in samples. The absolute size cannot be known as it depends strongly on the amplitude and kind of phase distortions in the NF, but for small amplitudes and distortion frequencies, equation (B.24) delivers a good rule of thumb estimation.
Besides that, one should keep in mind that the FFT interprets the input as periodical signal, meaning that the output is also periodical. As soon as the FF size approaches N, interactions between the "original" and the "copies" of the FF may occur, deteriorating the quality of the estimation. Generally, A_{FF} should be much smaller than N, while exact thresholds again depend on the scenario itself.

B.5.2. Numerical Calculation of the Debye-Wolf Diffraction Integral

Contrary to the FFT based numerical calulation of the Fraunhofer diffraction integral, one is not tied to regular grids in the case of the Debye-Wolf diffraction integral as long as proper weighting of the sampled points is done. In fact, as the grid has to be defined on the surface of a sphere, the grid cannot be regular in a cartesian sense.
Still, the sparse nature of sampling the NF does introduce artifacts which take an approximately periodic nature when a regular grid is chosen according to the spherical coordinate angles of the surface. The period p in the FF for a mean sample distance g in the NF can be estimated via simple trigonometric considerations:

$$\begin{aligned}
\cos(\alpha) &= \sqrt{1-\frac{\lambda^2}{g^2}} \\
\Rightarrow p = D\cdot\tan(\alpha) &= D\cdot\frac{\sqrt{1-\cos^2(\alpha)}}{\cos(\alpha)} \\
&= \frac{D\lambda}{\sqrt{g^2-\lambda^2}} \\
&\approx \frac{D\lambda}{g} \qquad \text{if } g \gg \lambda
\end{aligned} \tag{B.25}$$

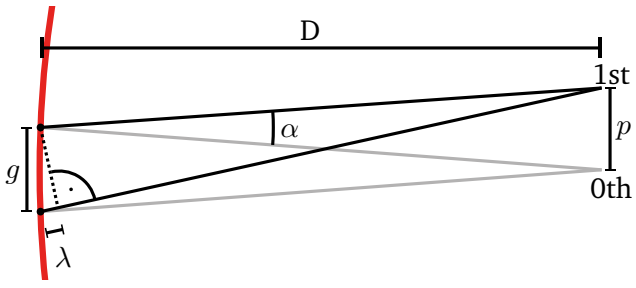

Figure B.9.: Schematic of the period estimation for the Debye-Wolf diffraction integral.

This is the same calculation as for a double slit, where the two sources interfere in a constructive manner on a screen at a distance D to the aperture. The distance p between the 0th and 1st order corresponds to the approximate distance of the replica in the probing plane of the diffraction integral. Therefore, one has to choose g such that p is much larger than the diameter of the diffraction pattern that shall be investigated.

B.5.3. Numerical Calculation of the Polychromatic Diffraction Integral

Analogous to the spatial discretization of the diffraction integrals, the spectral integral of the polychromatic diffraction, given by equation (B.22), must be expressed as sum as well:

$$\vec{E}_{\text{poly}}(\vec{x}, \Delta t) = \sum_k A(\nu_k) \cdot e^{i2\pi\nu_k \Delta t} \cdot \vec{E}(\vec{x}, \nu_k), \tag{B.26}$$

with k as a real integer number. In order to get accurate results from equation (B.26), one has to distribute ν_k on a regular grid in the frequency domain. While this will be the case in the majority of numerical calculations, one may also compensate for non-regular grids by introducing a weighting which is inversely proportional to the local sample density. In this work, however, I will only focus on the case with the regular grid.
Equation (B.26) effectively describes the addition of monochromatic waves in order to approximate the temporal behavior. For a regular grid, ν_k may be expressed as

$$\nu_k = \nu_0 + k \cdot \Delta\nu,$$

where ν_0 is a frequency offset and $\Delta\nu$ the free spectral range of the grid. Inserting this into equation (B.26), one gets

$$\begin{aligned} \vec{E}_{\text{poly}}(\vec{x}, \Delta t) &= \sum_k A(\nu_k) \cdot e^{i2\pi\Delta t(\nu_0 + k\cdot\Delta\nu)} \cdot \vec{E}(\vec{x}, \nu_k), \\ \Leftrightarrow \vec{E}_{\text{poly}}(\vec{x}, \Delta t) &= e^{i2\pi\Delta t\nu_0} \sum_k A(\nu_k) \cdot e^{i2\pi k\Delta t\Delta\nu} \cdot \vec{E}(\vec{x}, \nu_k). \end{aligned} \tag{B.27}$$

In this form, the sum is multiplied with a phase factor which evolves over time, but does not influence the amplitude. The oscillatory term inside the sum, however, returns to the same value for all k whenever $\Delta t\Delta\nu$ takes an integer value. Framed differently, the electric field evolves periodically with a period of $1/\Delta\nu$, only changed by a phase. This is intuitive as a monochromatic wave oscillates indefinitely and only discrete, monochromatic oscillations are added up in the sum.
Accordingly, $\Delta\nu$ has to be chosen such that $1/\Delta\nu$ is much larger than the expected pulse duration in order to avoid overlapping these ”pulse copies”, which are mere numerical artifacts. An example for this is shown in Fig. B.10.

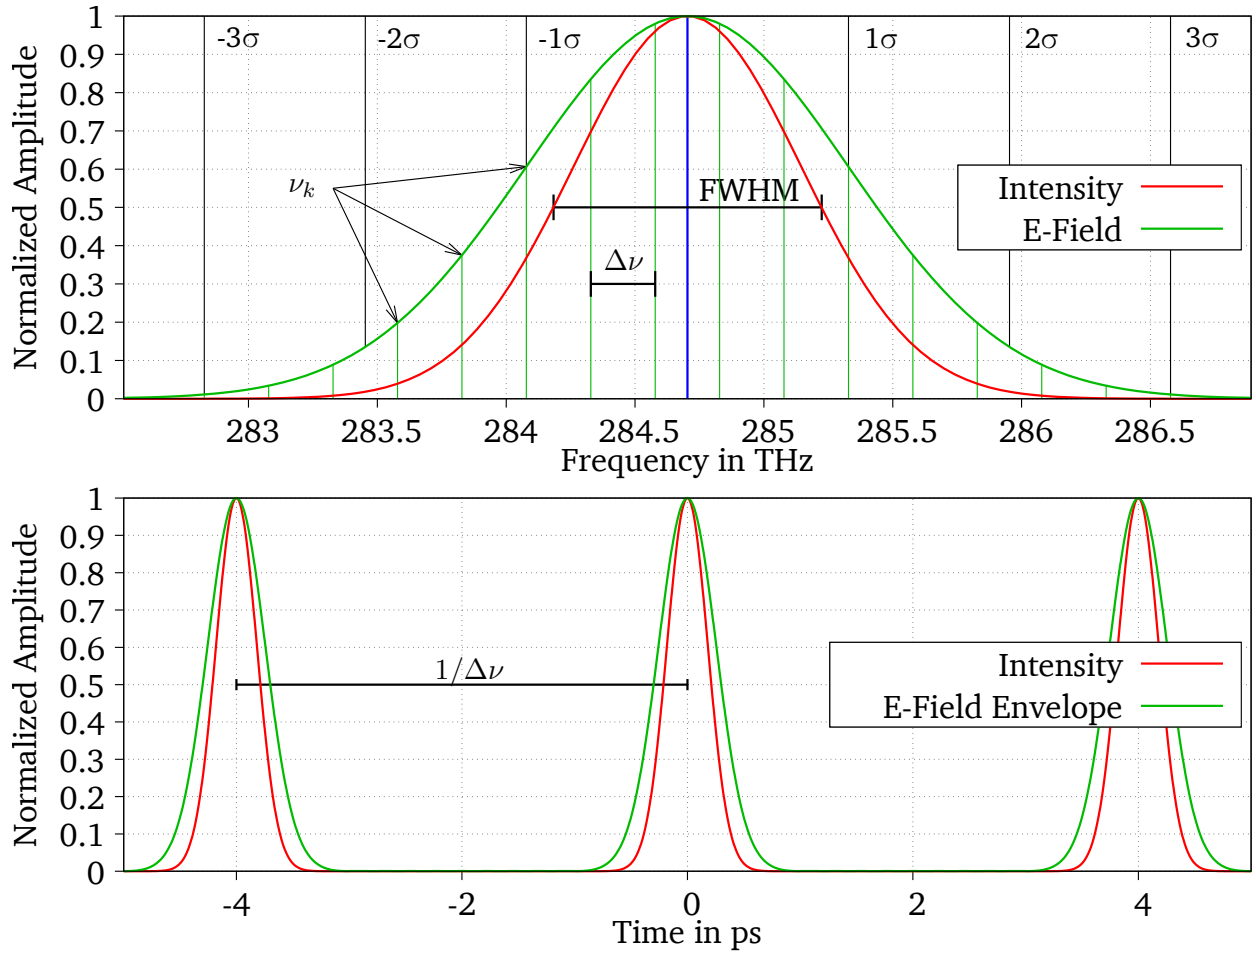

Figure B.10.: Top: A Gaussian spectrum with a central WL of 1053 nm. The red curve corresponds to the intensity envelope and the green one to the E-field envelope. The central WL is indicated in blue. The Full Width Half Maximum (FWHM) of the intensity envelope is chosen such that 16 samples with a $\Delta\nu$ of 0.25 THz fit inside a closed 3σ-environment of the electric field envelope, which yields a bandwidth of 3.85 nm.
Bottom: The temporal evolution of the pulse, calculated using (B.27), where the intensity repeats with a period ov $1/\Delta\nu$. The spectral phase was set to zero.

Conclusion

In this appendix, I have shown that a realistic calculation of the focal spot cannot be done using the Fraunhofer diffraction as the small angle approximations that are used to derive this integral do not hold true for the steep focusing angles that are common in HI laser systems. Instead, one needs to use the Debye-Wolf integral, which is a vectorial diffraction integral for a spherical geometry. Furthermore, the calculation needs to be polychromatic as soon as Spatio Temporal Couplings (STCs) are present in the beam or the bandwidth gets too large to approximate the pulse as monochromatic wave.
This is obviously a much more complex task than the calculation of the Fraunhofer diffraction image, which is just a single Fourier transform of a recorded WF. The recording of the spectrally resolved WF alone is a challenging task, while the transfer of the wave field onto the reference sphere for the diffraction integral requires a full blown ray-tracing approach, where false assumptions about the optics may lead to large errors.
Additionally, the computational costs of calculating a focal spot in this way is enormous. Considering the PHELIX pulse as an example: from equation (B.25), the spacing of the samples in the NF can be calculated for a desired spacing p of the artificial diffraction orders. If I demand the replica of the focal spots to be 30 µm apart, which is roughly ten times the focal spot size and should prevent interference effects, the spacing in the NF would be about 14 mm, which means that approximately 400 samples have to be included in the 280 mm aperture.
Accounting for the polychromatic integral, I can apply the same argument: if the temporal replica should be spaced by 5 ps, which is ten times the pulse duration, then the integral requires 20 spectral samples to include the full spectrum of PHELIX. This means that, the content of the diffraction integral needs to be calculated 8000 times for each sampling position in the focal plane.
In order to obtain a two-dimensional image of the focal spot, this calculation has to be repeated many times. For the focal spot at PHELIX, a window of 5 µm would contain the most interesting parts of the focal spot. At a resolution of 0.1 µm, this means evaluating the content of the integral 20 million times in total.
This clearly implies that the calculation itself poses a significant challenge, too, even if all the data was available. It would only become feasible with proper memory management and a General Purpose Computation on Graphics Processing Unit (GPGPU) accelerated evaluation routine. This is even more evident as the sample spacing of 14 mm in the NF cannot include high spatial frequencies, which are commonly present on milled Off-Axis Parabolic mirrors (OAPs).
In total, the lack of a spectrally resolved WF measurement and the computational demands of the diffraction integral led me to the decision to estimate the reached on-shot intensity, as shown in section 5.4 rather than simulating it.