Apollon Real-Time Adaptive Optics (ARTAO) - Astronomy-Inspired Wavefront Stabilization in Ultraintense Lasers

J. B. Ohland[1,2], N. Lebas[1], V. Deo[3], O. Guyon[3,4,5,6], F. Mathieu[1], P. Audebert[1], and D. Papadopoulos[1]

[1]*Laboratoire d'Utilisation des Lasers Intenses, Ecole Polytechnique, 91128 Palaiseau Cedex, France*

[2]*GSI Helmholtzzentrum für Schwerionenforschung, Planckstraße 1, 64291 Darmstadt, Germany*

[3]*Subaru Telescope, National Astronomical Observatory of Japan, National Institute of Natural Sciences, Hilo, HI, U.S.A.*

[4]*Astrobiology Center of NINS, 2-21-1, Osawa, Mitaka, Tokyo, Japan*

[5]*Steward Observatory, University of Arizona, Tucson, AZ, U.S.A.*

[6]*College of Optical Sciences, University of Arizona, Tucson, AZ, U.S.A.*

Abstract

Traditional wavefront control in high-energy, high-intensity laser systems usually lacks real-time capability, failing to address dynamic aberrations. This limits experimental accuracy due to shot-to-shot fluctuations and necessitates long cool-down phases to mitigate thermal effects, particularly as higher repetition rates become essential, e.g. in Inertial Fusion research.

This paper details the development and implementation of a real-time capable adaptive optics system at the Apollon laser facility. Inspired by astronomical adaptive optics, the system uses a fiber-coupled 905 nm laser diode as a pilot beam that allows for spectral separation, bypassing the constraints of pulsed lasers. A GPU-based controller, built on the open-source CACAO framework, manages a loop comprising a bimorph deformable mirror and high-speed Shack-Hartmann sensor. Initial tests showed excellent stability and effective aberration correction. However, integration into the Apollon laser revealed critical challenges unique to the laser environment that must be resolved to ensure safe operation with amplified shots.

1. Introduction

Semi-static Adaptive Optics (AO) has been an essential tool in controlling beam quality in high-energy, high-power laser facilities. These systems effectively remove static aberrations, enabling the highest intensities on target and unlocking new experimental regimes[1–5]. With the rapid advancement of laser technology, many beam quality issues have been addressed, revealing limitations that were previously unnoticed or considered irrelevant. Among these newly recognized limitations, dynamic aberrations have gained attention in recent years[6–9].

Dynamic aberrations evolve over various timescales: they range from slow drifts, e.g. due to temperature changes of optical tables, over medium timescale changes, such as cooldown effects of active media, up to fast evolution due to air turbulence or system vibrations. While pointing jitter is often mitigated using fast tip-tilt mirrors, higher-order aberrations have not yet been successfully addressed in real

time. The increasing public interest in high-energy lasers with higher repetition rates, as demonstrated by initiatives like the EU-funded Technology for High-Repetition-rate Intense Laser Laboratories (THRILL) program[1] and developments towards Inertial Fusion Energy (IFE), underscores the urgency of addressing this issue.

Given the limitations of passive mitigation strategies for dynamic aberrations, the path forward clearly involves active beam control through Real-Time Adaptive Optics (RTAO)[2]. This technology has been successfully applied for real-time wavefront corrections in fields such as astronomical instrumentation[10], ophthalmology[11], and free-space optical communication[12,13]. However, the high-energy laser community faces a significant challenge: due to the historical focus on (semi-)static AO, there is a lack of expertise in designing and constructing custom RTAO systems tailored

---

[1]https://www.thrill-project.eu/

[2]Please note that the term "AO" usually includes the real-time definition. However, in the laser community, "AO" referred to non-real-time systems for the past decades, which is why we deliberately chose use the "RTAO" abbreviation to distinguish these systems in our context.

Correspondence to: J.B. Ohland Email: j.b.ohland@gsi.de

to the specific needs of large high-power laser facilities. Consequently, no feasible commercial solutions currently exist.

In this work, we take the first step towards addressing this issue by implementing a potential RTAO solution for a specific case. This implementation provides valuable insights into system architecture, feasible technologies, and field-specific challenges. These insights will lay the foundation for the development of readily available RTAO systems in the field of scientific high-energy, high-power lasers, paving the way for a new era of beam stability in our community.

This paper is structured as follows: we begin with an introduction to the situation at the Apollon Laser in France in the next section, which was the motivation and use case that we focused on. We proceed with section 3 with a detailed description of the design considerations for our RTAO system, as well as the specific implementation in section 4. We then present the performance evaluation on a test bench in section 5 and in the actual system in section 6, while we discuss the main challenges that we encountered in the operational environment in detail in section 7. Based on these findings, we outline the future work necessary to integrate RTAO into regular operation in high-energy laser systems.

2. Background

This work was conducted at the Apollon Laser System[14,15], located close to Paris, France. Apollon is a Titanium Sapphire (Ti:Sapph) Chirped Pulse Amplification (CPA) laser system currently in its ramp-up phase, with the ambitious goal of delivering 150 J in 15 fs to the target, achieving 10 PW peak power at a repetition rate of one shot per minute. The laser pulses are initially generated in an Optical Parametric Chirped Pulse Amplification (OPCPA) frontend and then amplified through a series of five bowtie Ti:Sapph multipass amplifiers, each with increasing beam diameter.

The final amplifier in this sequence, known as "Amp300," features one of the world's largest Ti:Sapph crystals, measuring 196 mm in diameter. Amp300 is designed to bring the pulse energy up to 300 J, with the next upgrade in the pump area within 2024 already unlocking energies of 250 J. Achieving such high energies is made possible by a large beam diameter of 140 mm and a complex multiplexing sequence where currently seven pump beams pump the crystal between the four amplification passes.

While this architecture enables large pulse energies, the bowtie design of Amp300, coupled with its large beam diameter, results in over 50 m of unimaged beam propagation through laboratory air (see Fig. 1). This makes the beam extremely sensitive to air turbulence. Before any countermeasures were implemented, the Strehl ratio

Figure 1. Sketch of the Laser AMplification area (LAM) at Apollon. The last amplifier ("Amp300") was known for causing beam instabilities due to air movement in the beampath.

fluctuated between 0.2 and 0.9 within a single second, and the local wavefront slope fluctuated at up to 70 Hz. These fluctuations not only failed to meet the specifications of Apollon (maximum 10% Strehl ratio fluctuations) but also rendered the use of the Amp300 beamline unfeasible for high-intensity experiments.

Even after implementing passive measures, such as enclosing Amp300 in a tight beam housing, the Strehl ratio fluctuations still exceeded 20%, highlighting the limitations of passive beam control in certain configurations. This situation underscores the necessity for an active beam control solution, specifically a RTAO system.

It is important to note that while the architecture of Amp300 makes Apollon particularly susceptible to air turbulence, other facilities face similar challenges that limit certain types of experiments. For instance, the novel ELI NP laser system achieves remarkable beam stability due to its vibration damping system, which includes nearly 1000 springs and dampers supporting the massive 1.5-meter-thick reinforced concrete slab that forms the base of the laser hall. Nevertheless, atmospheric turbulence can still disrupt beam quality sufficiently to prevent certain types of experiments, e.g. when employing helical beams[6], which are highly sensitive to angular aberrations[16]. This issue becomes much more prominent when sudden weather changes such as thunderstorms occur, which cannot be fully compensated by the air-conditioning system.

Furthermore, glass-based high-energy lasers with increased repetition rates are currently being developed for IFE applications. The thermal load from these lasers will inevitably generate beam instabilities, further emphasizing the need for expertise in designing, implementing, and operating RTAO systems in. This motivation therefore extends beyond Apollon and permeates the entire high-energy laser community.

3. System Design and Methodology

3.1. Inspiration from Astronomy

RTAO has been an essential tool in astronomy for decades and is considered obligatory for large, earthbound telescopes in order to overcome the limitations imposed by atmospheric turbulence. Here, various RTAO schemes, such as high-resolution Single Conjugate RTAO for small fields of view, e.g. for exoplanet imaging[17], and Multi Conjugate RTAO with multiple Deformable Mirrors (DMs) for larger field of view observations[18], have been developed. These systems can handle many hundreds of control modes while operating at crossover frequencies (max. frequency that gets rejected by the loop) of hundreds of Hz (Tyson 2022[19], chapter 2). This extensive experience in astronomy provides valuable insights for implementing RTAO in high-energy laser systems. To leverage this expertise, we partnered with the Subaru Coronagraphic Extreme Adaptive Optics (SCExAO) team at the Subaru Observatory on Maunakea, Hawaii[10]. The SCExAO team, developers of the open source Compute And Control for Adaptive Optics (CACAO) software[20–22], supported us in adapting their technology to our laser requirements on various levels.

3.2. Design Considerations

Implementing RTAO in high-energy laser systems presents several unique challenges that differ from those encountered in astronomy. These differences necessitate specific adaptations in our approach:

Light Source One major difference is the light source for sensing. In astronomy, continuous light from natural or artificial guide stars (Tyson 2022[19], chapter 3.3) is typically available. However, high-energy, high-intensity laser systems operate with pulsed beams, with intervals of no light that can last for minutes during shot preparation.

One way to address this is to use a continuous pilot beam that co-propagates with the main beam, thus sampling the air turbulence and providing the RTAO loop with a signal to act upon. Like that, the loop can compensate the dynamic aberrations in the beamline up to shot delivery. The main concern using this technique is the protection of the sensor from the high fluence that occurs on shot. To mitigate this, the pilot beam can either be blocked a few milliseconds prior to the shot using a fast shutter, or the beam can be passively separated - e.g. by employing an off-spectral pilot beam that features a wavelength that can be separated from the main pulse using a dichroic mirror and filters, which we decided to du at Apollon.

Wavefront Sensor The choice of Wavefront Sensor (WFS) also required adaptation. While Pyramid Wavefront Sensors (PWFSs) are commonly used in astronomy due to their ability to deal with photon noise when running on dim light sources, as well as high spatial resolution and computationally cheap evaluation (Tyson 2022[19], chapter 5.3), the PWFS faces limitations in its ability to deal with non-common path aberrations. This is typically no big concern in astronomical instrumentation as the beam diameters remain small after splitting the light between the WFS and the scientific camera. In our field, however, the sampling is done at the largest beam diameters after the beam is amplified in order to prevent damaging optics. The successive imaging and demagnification system usually generates significant amounts of non-common path aberrations, which renders PWFSs unfeasible for this application.

For this reason, other WFS schemes that cope well with this issue have to be used while carefully tweaking the evaluation pipeline to still feature sufficient speed and throughput for RTAO applications. One example would be using a Shack-Hartmann Sensor (SHS), which features a large dynamic range, and implement a suitably fast evaluation routine, as done in this project.

Deformable Mirror Another critical consideration was the selection of the DM. In astronomical instrumentation, small, fast, high-resolution Micro-ElectroMechanical System (MEMS) mirrors are commonly used as they greatly fit this application[10,17].

However, the Laser-Induced Damage Threshold (LIDT) of these DMs is too low for high-energy, high-intensity lasers, as the thin membranes cannot withstand the stress of highly reflective dielectric coatings. At the same time, this technology cannot be scaled up easily to larger diameters at reduced fluence.

To our knowledge, the only DM technology that allows to accommodate both specialized coatings and larger diameters are piezoelectric ones. Standard piezoelectric bimorphs feature a good balance of size, speed, stroke and high damage threshold, while remaining relatively affordable[23].

RTAO Controller The demands to the computation hardware are easily derived from the requirements for latency (see Sec. 5.1) and the number of control modes and quickly exceeds the performance of a Central Processing Unit (CPU)-based system. The two main competitors for high-speed RTAO control hardware are Field-Programmable Gate Array (FPGA)[13,24,25] and Graphics Processing Unit (GPU)[20,26] based systems. Both deliver high throughput and low latency computation to achieve the required performance.

In the laser community, AO controllers were usually non-real-time application that exclusively ran on the computers CPU. Despite FPGAs potentially featuring lower latency,

the increased and hardware-specific development efforts rule this option out as a community project. For this reason, using open-source GPU code on a Real-Time Computer (RTC) is a more suitable approach.

4. Implementation

In this section, we provide details of the implementation of the Apollon Real Time Adaptive Optics (ARTAO) system. The setup is based on the design considerations that we discussed in the last section.

4.1. Pilot Beam Approach

For ARTAO, we chose the passive off-spectral pilot beam scheme. Our previous investigations showed that a wavelength of 905 nm is easy to separate from the main spectrum (730 nm - 890 nm) using dichroic mirrors and filters, while featuring a transmission of 10% up to the 1 PW compressor entrance, which is sufficient to saturate the WFS at an injected beam power of 20 mW. Moreover, the passive approach is easier to implement and more robust as there is no risk of damaging the WFS due to shutter failure.

We injected the single-mode-fiber-coupled 905 nm laser diode after the second amplifier where the beam was still small and not impaired by air turbulence. The pilot beam was then separated right before the vacuum tank of the 1 PW compressor (see section 6 for details on the sensor setup).

4.2. Real-Time Controller

As this is, to the best of our knowledge, the first time that a fully custom RTAO loop was implemented in an ultra-intense laser system, reusability in other systems was a concern. For that reason, we decided to adapt the open source Compute And Control for Adaptive Optics (CACAO) code, developed by the SCExAO team, due to its robust features, maturity, and its use of GPUs to achieve high-throughput, low latency operation. The architecture and features of CACAO are described in detail in the original publications by the SCExAO team[20–22]. In brief, CACAO can be described as a framework where computations are split into microservices. These processes are triggered by image stream updates and thus form a fully customizable AO pipeline.

For ARTAO, we extended CACAO by implementing hardware interfaces, custom WFS referencing and evaluation functions and a vectorized Wavefront (WF) management to avoid unnecessary reshaping operations.

The hardware that we used to run CACAO was inspired by the KalAO[27] system that featured specifications that were comparable to what we intended to achieve. The RTC was assembled solely of consumer grade components (see appendix A for a full list of components) and can be compared to a high-end gaming computer. As such, the assembly was relatively cheap compared to the high-speed camera and the DM while easily delivering the required performance (see Sec. 5.1 for requirement estimation).

4.3. DM Selection

As discussed before, the requirements for speed and damage threshold significantly reduce the number of choices concerning DM technology, where piezoelectric DMs are currently the most feasible options. While placing a large DM after the amplifiers as close to the WFS as possible would be the most unproblematic solution due to more direct DM imaging and less pre-compensation (foreshadowing, see Sec. 7), the increased size does not only mean increased costs, but also a lower resonance frequency for bimorph DMs. Piezostack DMs would avoid this issue, but feature a significantly lower dynamic range.

For these reasons, we decided to place a smaller bimorph DM, optimized for a 55 mm beam diameter, prior to the second last amplifier and image it to the wavefront sensor instead. The DM does therefore pre-compensate the dynamic aberrations in the main amplifier.

In this work, we were using a bimorph with 60 mm clear aperture and 96 actuators in an annular arrangement, optimized for Zernike stroke at an angle of incidence of 0°-10° (*DM9660* by Dynamic Optics, Italy). The characterization is shown in section 5.3.

4.4. Wavefront Sensor

For the WFS, we constructed a custom SHS using a Micro Lens Array (MLA) with a pitch of 250 µm and a focal length of 11 mm and a high-speed PCIe camera (*CB013MG-LX-X8G3-TG*, xiB-64 series by XIMEA), capable of streaming frames directly to the RTCs memory (or even to the GPU, if it features GPUDirect RDMA) at up to 7 kHz for a Region Of Interest (ROI) of 5×5 mm^2 (384×384 px^2) at a negligible latency. We used the taped glass version of the camera in order to remove the cover glass from the chip as this would cause interference patterns when a coherent, monochromatic light source, such as our Continuous Wave (CW) pilot beam, is employed.

We determined the distance of the MLA to the camera chip by introducing know amounts of tilt and observing the motion of the centroids on the camera image. The exact knowledge of this distance enables quantitative wavefront measurements relative to a reference. The reference was kept relative by design as the beam optimization at Apollon is done using a standard procedure with a large, slow mechanical DM (*ILAO180*, Imagine Optic), which provides the target WF to be stabilized.

For the evaluation, we implemented two routines in CACAO: a referencing routine, generating a fixed WF pupil with reference positions and a GPU-based routine that uses this reference to calculate the WF from the camera image. The latter is logically a three-step process: first, the image

is convolved with a Gaussian kernel, fitting the standard deviation of the spots of the SHS, in order to increase robustness against noise. Second, a parabola is fitted to the brightest pixel and it's closest neighbors of each smoothed spot in order to retrieve the fractional shift of the spot[28]. And third, the retrieved slopes are multiplied with a matrix that translates the gradient vector to a WF vector.

Practically, we used an adaptive windowing approach, where only the relevant pixels around each spots are copied to the GPU. since the standard deviation of the focal spots (13.9 μm) nicely matched the camera pixel size (13.7 μm) in our case, the algorithm uses a kernel size of 3×3 pixel. Under the assumption that the spot position did move less than half a pixel since the last frame, 6×6 pixels, excluding the corner pixels, need to be streamed to the GPU.

The entire operation is completed within approximately 60 μs, well within the latency budget.

5. Characterization and Testing

5.1. System Modeling

In order to verify the system's performance at a later point, we approximated the loop as a single-input-single-output system and constructed a transfer function in the Laplace domain (Tyson 2022[19], chapter 8.2). In the following, we briefly describe the components of the feedback transfer function, where the formulae are listed after the enumeration:

(1) WFS Exposure: $H_{\mathrm{exp}}(s)$
   Modeled as an integrator normalized to the exposure time T_{exp}, minus another normalized integrator delayed by T_{exp}, effectively simulating the average response within the exposure window.

(2) Frame transfer: $H_{\mathrm{trans}}(s)$
   We represented the delay due to the WFS frame transmission Represented as a simple delay corresponding to the max. possible frame rate - approximately $T_{trans} =$135 us in our case.

(3) Controller Response: $H_{\mathrm{RTC}}(s)$
   The control unit was emulated by a discrete sampling (chain of dirac impulses) every control cycle period T,

coupled with a delay T_{c}, resembling the computation time and driver latency, and a gain factor G.

(4) Leakage: H_{leak}
   As our loop ran using a leaky integrator, i.e. a small portion of the integrated signal is "thrown away" in order to avoid the accumulation of errors, the leakage per control step needed to be included in the model. For this, we modified the transfer function of an exponential decay in a way that we could use the per-step leakage multiplier λ to determine how much of the control signal is discarded every iteration.

(5) Driver Hold: $H_{\mathrm{DRV}}(s)$
   The DM driver itself was modeled by a simple zero order hold for the control cycle period T. Considering the control signal being a chain of dirac pulses with variable amplitude, this is again an averaging over the control period T.

(6) DM Response: $H_{\mathrm{DM}}(s)$
   The speed of the bimorph DM is limited by the charge time of the electrodes. For this reason, we approximate the rise time by a linear ramp. Taking the step-like signal from the driver, this is yet another averaging transfer function. The rise time T_{DM} was extracted from measurements (see next section.)

In formulae, these contributions equal:

$$H_{\mathrm{exp}}(s) = \frac{1}{sT_{\mathrm{exp}}} - \frac{e^{-sT_{\mathrm{exp}}}}{sT_{\mathrm{exp}}} = \frac{1 - e^{-sT_{\mathrm{exp}}}}{sT_{\mathrm{exp}}}$$

$$H_{\mathrm{trans}}(s) = e^{-sT_{\mathrm{trans}}}$$

$$H_{\mathrm{RTC}}(s) = G \cdot \frac{e^{-sT_{\mathrm{c}}}}{1 - e^{-sT}}$$

$$H_{\mathrm{leak}}(s) = \frac{sT}{sT - \ln(\lambda)}$$

$$H_{\mathrm{DRV}}(s) = \frac{1 - e^{-sT}}{sT}$$

$$H_{\mathrm{DM}}(s) = \frac{1 - e^{-sT_{\mathrm{DM}}}}{sT_{\mathrm{DM}}}$$

Parameter	Description	Target	Testbench	Apollon
T_{exp}	WFS exposure time, reciprocal of max. framerate	100 us	1/(7 kHz)	1/(7 kHz)
T_{trans}	frame transfer time, reciprocal of max. framerate	100 us	1/(7 kHz)	1/(7 kHz)
G	control feedback gain	0.5	variable	0.27; 0.35
T_{c}	computation + driver latency, determined experimentally*	1 ms	1.2 ms	1.0 ms
T	control period, reciprocal of control frame rate	1/(1 kHz)	1/(3.6 kHz)	1/(2.6 kHz)
λ	leakage multiplier per step	0.999	0.999	0.99
T_{DM}	DM rise time (10%-90%), determined experimentally	500 μs	486 μs	486 μs

Table 1. The parameters used in the model of the ARTAO loop for the estimation, on the testbench and in the Apollon beamline itself.
*T_{c} is dominated by the latency of the DM driver. While this is hard to separate from other effects, measurements indicate a delay of approx. 0.7 ms (see Sec. 5.3)

Figure 2. The Bode plot of the feedback transfer function $F(s)$ (left) and the error transfer function $R(s)$ (right) according to Eq. (1), using the parameter values from the "Target" column in table 1.

The parameters are listed in table 1.

The corresponding error transfer function of the loop is

$$R(s) = \frac{1}{1 + F(s)}, \qquad (1)$$

where $F(s)$ is the feedback transfer function, which in turn is the product of all transfer functions discussed above. The corresponding magnitude Bode plot (amplitude response) at frequency ω is then obtained by evaluating $R(i\omega)$. Prior to the design of ARTAO, we evaluated $R(s)$ for a set of different parameters to aid the selection of the hardware. In table 1, we listed the set of target parameters that yielded a crossover frequency of approx. 70 Hz (cutoff frequency of the measured turbulence) with suitable gain and phase margins (see[19], chapter 8.2.9), while being realistically achievable. The corresponding Bode plots are shown in Fig. 2.

We will use the same model with different parameter sets in the next section to validate our understanding of the system with experimental data.

5.2. Testbench Setup

Figure 3. Schematic of the test-bench setup, to scale.

To evaluate the functionality and performance of the ARTAO system, we set up a testbench with a simplified configuration (see Fig. 3). It consisted of a single beampath where we collimated the pilot source after exiting the fiber, an aperture fixing the beam size, an imaging beam expander to the DM

and an imaging system to the SHS afterwards. Although we did not implement a physical turbulence generator to emulate the operational conditions, we utilized CACAO's internal turbulence injection to introduce disturbances to the DM. The primary objectives of this testbench setup were to verify the system's functionality, characterize the DM, measure system latency, and determine the frequency response.

5.3. Performance Evaluation

We conducted several performance tests to validate the system.

DM Latency and Rise Time: Using a strobing technique, we evaluated the DM latency and rise time. By repeatedly poking a selected DM mode a couple of hundred times while varying the delay with respect to the camera frame, we generated a corresponding response curve (see Fig. 4 for single-mode example curves), which is the DM response, convolved with the camera exposure timeframe, and delayed by the driver delay, the camera delay and the WFS evaluation time. From this, we extracted both the latency between the command and the first available wavefront of the onset of the risetime and the 10%-90% rise time of the DM itself. We then repeated process for all control modes of the control matrix (see Fig. 5). The rise time is fairly consistent across all modes except for the lowest three ones, roughly corresponding to Tip/Tilt and Defocus. When excluding these three, the rise time is at $486\pm50\,\mu s$, following a delay of $780\pm49\,\mu s$ (including driver delay, camera exposure, frame transfer and WFS evaluation).

The lowest three modes showed notable ringing due to the DM mechanics. This ringing significantly increased the settling time to 13 ms, 8.3 ms and 1.4 ms, respectively. While this is part of a separate project, this clearly indicates the advantage of a separate pointing control system if no DM with better temporal characteristics on these modes is

available.
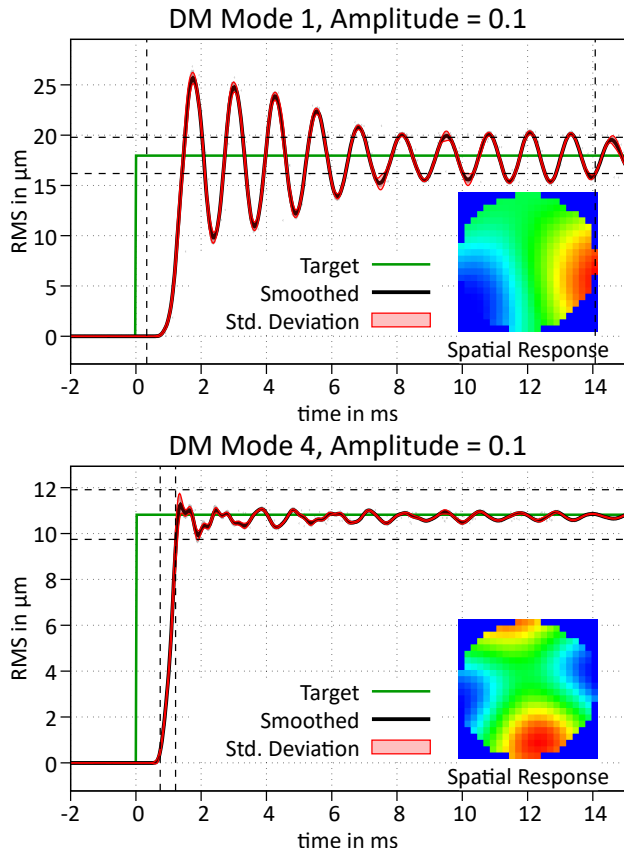
Figure 4. Step response of two DM modes over time (open loop). The dashed lines indicate the 10%-90% levels and corresponding settling times. Top: first mode, featuring severe ringing due to the mechanical DM properties. Bottom: fourth mode with a regular settling behavior. All modes from this order upwards feature a comparable settling behavior (see Fig. 5.


Figure 5. Step responses of a full set of mirror modes over time.

System Frequency Response: We measured the system's frequency response using CACAO's internal turbulence em-

ulator, which introduced disturbances on the DM by sweeping a Kolmogorov screen[29] across the actuator matrix at a constant velocity. By measuring the local wavefront jitter in both open and closed loop configurations, applying spectral smoothing, and determining the ratio of the two for different stable loop gains, we experimentally measured the frequency response of the system (see Fig. 6).

At this point, we used the model for the error transfer function (Eq. (1)) as a fit model, where the loop delay was the fit parameter. At a delay of 1.2 ms, we found excellent agreement between the model and the experimental data.

In the lowest frequency range, the deviation of the data from the simulated curves increases, while also featuring some correlation between the measurements. The latter can be attributed to the periodic nature of the Kolmogorov screen that we used to disturb the loop, while the magnitude of the errors are caused by the limited recording duration (approx. 13.65 s in 12 image cubes, 4096 frames each), which leads to poor statistics in the low frequency domain.
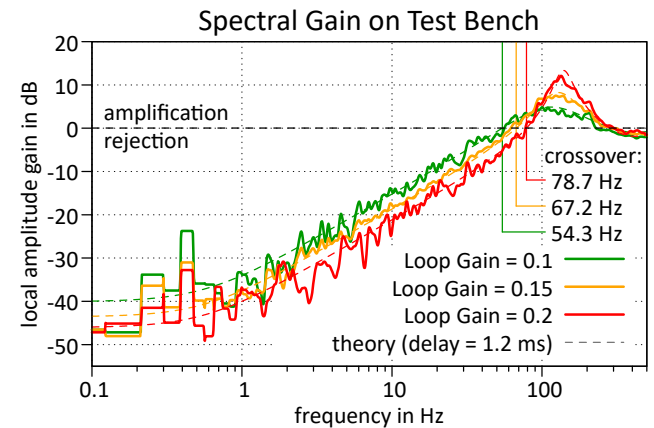
Figure 6. The AO loop gain over the frequency for different feedback gains. The dashed curves are the gains of the corresponding model, i.e. the magnitude of the error transfer function.

6. Integration with Apollon Laser System

6.1. Implementation Details

We performed the integration of the ARTAO system into the Apollon laser system in a way that did not require significant changes in the main beamline that could potentially interfere with regular operation. In the following, we will briefly describe the changes we made. The setup is shown schematically in Fig. 7.

We injected the 905 nm CW pilot beam right after the last Pockels cell shutter of Apollon, located between the second and third amplifiers. Using a reflective fiber collimator, the beam featured a diameter slightly above the diameter of the main beam. We aligned the polarization using a wave plate and injected the pilot through a dichroic mirror. The regular beam alignment references were used to ensure co-

Figure 7. Schematic setup of ARTAO in the Apollon laser chain

propagation of both the pilot and main beams.

In order to enable easy bypassing of the DM, we integrated a beam detour prior to the fourth amplifier, where the beam diameter is 55 mm, ensuring a low fluence to prevent any potential damage to the DM. The detour injection- and pickup mirrors were placed on a linear stage and could be inserted and removed whenever necessary.

Figure 8. Sketch of the diagnositc setup prior to the 1 PW compressor in side- (a) and top view (b). The main beam path is shown in red, while the pilot beam path is indicated in orange.

For the sensor setup, we implemented a dedicated beam sensor configuration right before the entrance of the 1 PW compressor (see Fig. 8). Leakage light from the main beam was focused by an f = 8 m lens and then collimated. A dichroic mirror, followed by spectral filters, separated the pilot and main beams, allowing for independent diagnostics of each. The fast SHS was placed in the pilot beam arm in the image plane of the DM and was complemented by a Far Field (FF) camera. Regular beam diagnostics, including near field (NF), FF, and a slow SHS, were placed in the main beam arm, where the image plane of the WFS matched that of the DM as well to enable direct comparison between the

WFs of the two beams.

The RTC for ARTAO was connected to the fast WFS via a 5 m PCIe copper cable and positioned close to the setup. We used a 50 m long CAT6 ethernet cable to directly connect the RTC and the DM, both of which are located in different rooms within the infrastructure.
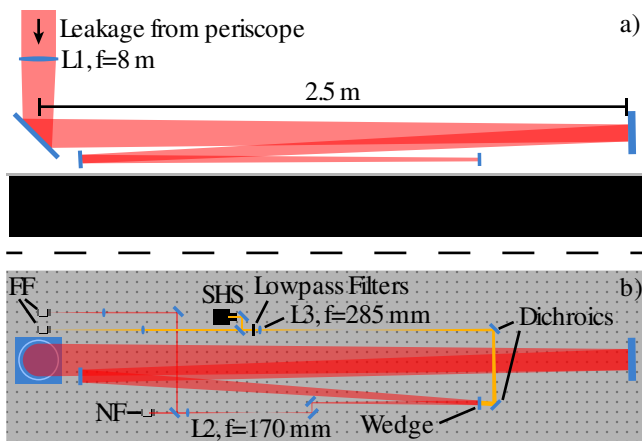
6.2. *Pilot Beam Correlation*

Using a pilot beam introduces uncertainties due to the differences in frequency and potentially alignment. In order to ensure that running an RTAO loop on the pilot beam benefits the real beam, we investigated the correlation between the two beams. To do so, we simultaneously recorded WFs on both beams using the diagnostic setup.

To determine the mapping from one WFS to the other, we constructed a sample-wise correlation matrix from the recorded sequences, where we subtracted Tip/Tilt and the mean WF beforehand (see Fig. 9). This matrix illustrates how well each WF sample in one WFS space correlates to a specific pixel in the other WFS space. We extracted the positions of maximum correlation, generating a set of point pairs, which we used to fit an affine transformation. This transformation allowed us to convert WFs between the two spaces using linear interpolation. An example frame is shown in Fig. 10, with an animated version in the supplementary material of this publication.

The analysis revealed excellent correlation between the WFs of the pilot and main beams, with errors being both of high spatial order (see Fig. 10, right column) and of significantly smaller amplitude than the full beam disturbance (less than 20%, see Fig. 11). Potential sources of these errors include imperfect synchronization (the fast SHS was not triggered, estimated at approx. 20 ms accuracy in frame selection), errors in the fitted transformation, and linear interpolation inaccuracies. While the precise attribution of these errors to the sources remains uncertain, we consider the overall error magnitude to be sufficiently small to validate the pilot beam concept for this application - especially since the

Figure 9. Bottom: the sample-wise correlation matrix between the WFS of the main beam and the pilot beam over the recorded sequence without tilt and mean WF. Top: example correlation of the main beam WF to a randomly picked location of the pilot WF (left) and vice versa (right). Point pairs for a transformation fit can be extracted from the locations of max. correlation.

Strehl ratio deficit scales quadratically with the Root Mean Square (RMS) for small aberrations.

Figure 11. RMS of the main- and the pilot beam WF, as well as the difference between the two, over a time frame of one minute. The beam was actively disturbed using a hot air source for this measurement.

6.3. Short-term Stabilization

To evaluate the short-term stabilization performance of the RTAO system, we conducted a series of experiments under
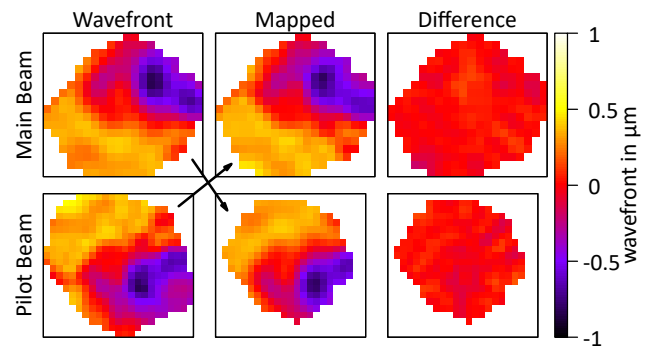

Figure 10. Example of a mapped WF between the main beam WFS (top row) and the pilot beam WFS (bottom row), where the first column is the raw WF, the second one is the mapped WF from the other WFS, respectively, and the last column is the difference between the two. Note that the main beam is smaller than the pilot beam, which is why its mapped WF is smaller in the pilots WFS space.

regular operational conditions, (i.e. all beam housings in place and air conditioning systems active). The primary goal was to achieve optimal WF stability by fine-tuning the loop parameters to the environment.

In these experiments, we mainly adjusted the framerate of the WFS, the Tip/Tilt gain and the modal gain curves. For each setting, we incrementally increased the feedback gain to the edge of stability. We then recorded a time sequence of WFs of one minute prior to closing the loop and continued to record for another minute. From the WFs, we calculated the Strehl ratio via a FFT of the measured NF, which we used as relative measure for the beam stability. This Strehl ratio obviously only applies if all static aberrations are removed up to the target area and is thus an unrealistic value from the experiment point of view, but it remains useful to build some intuition on the performance of the AO loop in terms of achievable stability. The result is shown in Fig. 12. The loop exhibits excellent performance on this short time scale, increasing the minimum Strehl ratio from 0.65 to 0.96.

Additionally, we calculated the spectral gain curve in order to compare the data to our model. The parameters that we used to obtain the theoretical curves are shown in the right column in table 1. Two deviations meet the eye: first, the rejection is significantly worse than the model for frequencies lower than 10 Hz. The reason for this is that we characterized the gain based on the actual recorded WFs, where it's portion that is not covered by the DM modes remains uncorrected. Apparently, the spatial frequencies feature a maximum in the low frequency regime, where they are additionally more visible due to the otherwise large rejection of the loop. Secondly, plugging in the actual loop parameters into the model (dashed red curve in Fig. 13) does not yield a good overlap in the higher frequency regime anymore. This is likely due to the fact that we applied a modal gain curve to the real-world loop, featuring a

Strehl Ratio, ARTAO activated at t=0

## Strehl histogram (1 min, 2.6 kHz)
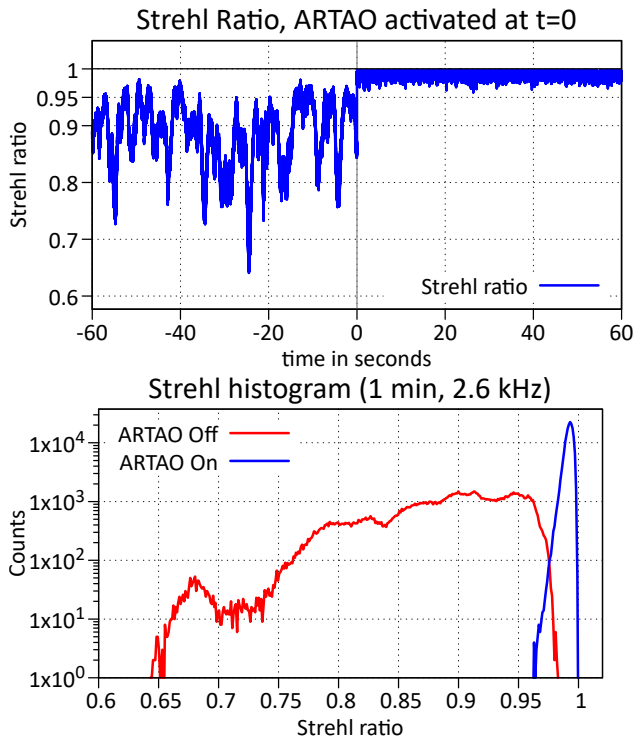
Figure 12. Top: time series measurement of the Strehl ratio (compared to the reference WF, calculated via the FFT of the measured NF) of the pilot beam, where ARTAO is activated at t=0. Bottom: the histogram of the corresponding data series.

significantly higher gain (factor 1.4) for the Tip/Tilt modes of the DM, which is obviously not reflected by the model. With an empirical upwards correction of the loop gain (0.27 to $0.35 \approx$ a factor of 1.3, which is in between the base gain and the Tip/Tilt gain), the model resembles the data well again (dashed blue curve in Fig. 13). The crossover frequency at the edge of stability was 96.8 Hz.

Pointing, ARTAO activated at t=0
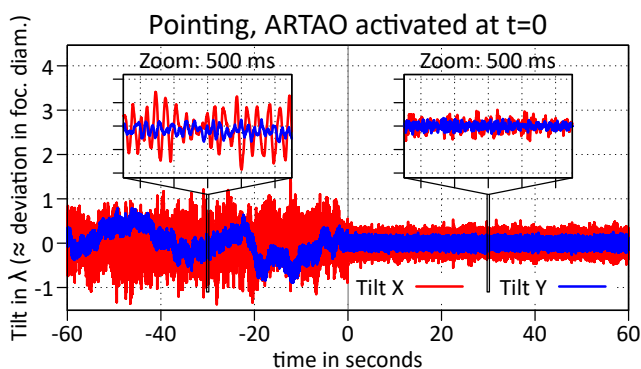
Figure 14. Time series measurement of the beam pointing of the pilot beam, where ARTAO is activated at t=0. The tilt X (red) and tilt Y (blue) curves represent the tilt-portion of the recorded WF, relative to the reference, and is given in peak-to-valley in terms of the central wavelength of Apollon, roughly corresponding to the movement of the focal spot in focal spot diameters. The inserts show a zoomed in portion to illustrate the fast oscillations in the beam pointing.

Spectral Gain in Apollon Beamline
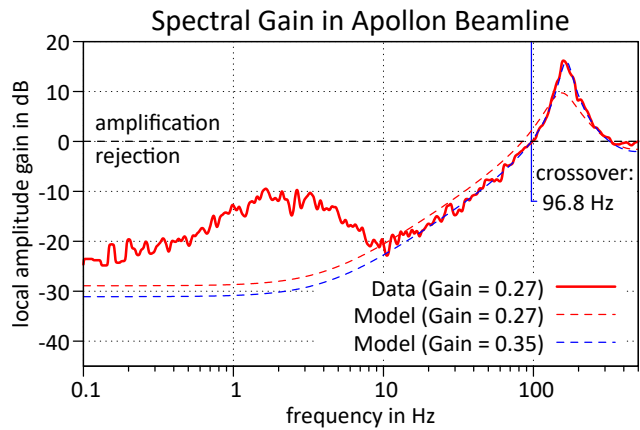
Figure 13. Recorded gain curve of ARTAO on the pilot beam WF in the Apollon beamline under regular opeartion conditions, compared with two theoretical curves (dashed lines) with parameters from table 1. The dashed red curve uses the parameters of the real-world loop, while we tweaked the feedback gain for the blue curve in order to match the data.

Furthermore, we analyzed the data with respect to the pointing stability of the beam with and without ARTAO (see Fig. 14). This plot demonstrates that the loop is able to achieve some level of compensation of the pointing instabilities, the limitations of the DM and the loop bandwidth in general become relevant here. The pointing is dominated by an oscillation at around 35 Hz, as shown in the zoomed insert in Fig. 14. This is a significant challenge considering the ringing of the first modes of the DM (see Fig. 4 and 5) and the oscillations cannot be suppressed effectively, only yielding a Peak-to-Valley rejection of 50% compared to the open loop. This underlines the necessity to implement a fast steering mirror to outsource this mode to a faster, more rigid system. While this is a project currently under its way at Apollon, this goes beyond the scope of this work.

6.4. Impact of Pumping the Amplifier Crystals

Closed Loop Pilot RMS During Amplification

Figure 15. Measured RMS of the closed loop pilot WF during a pump event on the amplifiers. The inserts are the WFs at the times indicated by the arrows.

Although we refrained from running the loop while delivering amplified shots due to the risks involved (as discussed in section 7), we investigated the impact of amplifier pumping on the loop stability: we flashed the amplifiers without a frontend beam while the ARTAO loop was closed on the pilot beam in order to observe any potential effects on the system's performance.

The results, presented in Fig. 15, show the RMS of the measured WF over time, while the ARTAO loop was closed. In the center of the plot, the amplification event occurs. At this instant, the SHS's spots became saturated due to the amplifier gain at 905 nm. While this saturation did not pose a risk to the camera, it resulted in a presumable erroneous detection of WF tilt. Approximately 2 ms later, the loop reacted to the tilt measurement, attempting to compensate, which led to a noticeable but temporary reaction in the WF. However, the disturbance was insufficient to destabilize the loop, and the WF quickly returned to its optimal state within a few tens of milliseconds - far below the repetition rate of Apollon or even the frontend of the facility.

These observations indicate that while amplifier pumping does momentarily affect the WF measurement, the RTAO system is robust enough to quickly recover, maintaining overall stability in the short term.

7. Challenges and Future Work

While the results described in the previous sections demonstrates significant progress, several challenges remain before RTAO can be considered feasible for high-power laser systems. This section outlines the key issues encountered during testing and proposes steps to address them.

7.1. Long-Term Stability Issues
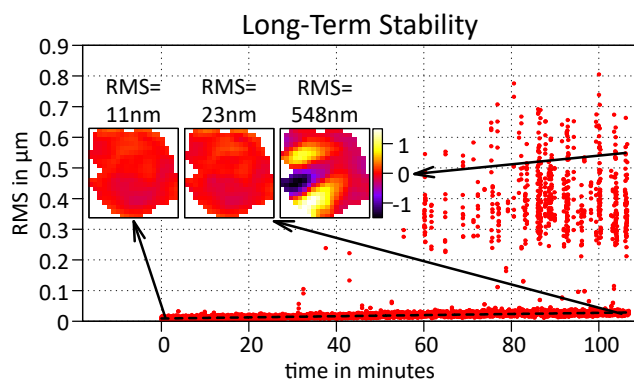
Figure 16. The WF RMS of the pilot beam under closed loop operation over an extended time frame. The insert plots are three selected WF frames from stable conditions in the beginning (left) and the end of the recording (center), as well as from a period of instability (right).

A critical limitation of the ARTAO system arises when considering long-term stability. Fig. 16 shows the evolution of the WF RMS of the pilot beam (relative to the initial

reference) over nearly two hours of closed-loop operation. While the loop remains stable for the first 30 minutes, isolated instability events occur after this point, increasing in frequency until continuous instability streaks are observed beyond the one-hour mark. The baseline stability of the wavefront also degrades slightly over time.

These instability events pose a significant risk for operation, as large DM strokes during these phases can create hotspots in the beam, potentially damaging system optics during shot delivery. The unstable wavefronts exhibit wave-like structures (Fig. 16, right inset) oscillating rapidly over time, strongly suggesting lateral misregistration between the WFS and DM[30]. This is likely due to environmental changes, e.g. temperature or humidity deforming the beamline slightly and thus shifting the DM image on the WFS over time.

This hypothesis was confirmed when minor manual adjustments to a randomly picked steering mirror outside the DM conjugation plane between the DM and the WFS realigned the DM image position, which in turn restored stability. Thus, implementing a tracking routine to maintain lateral beam alignment on the WFS is a promising countermeasure. This approach will require a sensitive analysis routine to detect beam drift early[30].

It should be noted, however, that other sources of instability are still worth consideration. This includes, e.g., a drop in illumination on the WFS, software crashes and exceedingly fast vibrations. This implies that a robust machine safety framework is necessary in any case.

7.2. Nearfield Intensity Variations

Figure 17. Imprints of artificially large DM strokes onto the NF fluence in a non-conjugate image plane, where each NF corresponds to a different set of random actuator positions.

Another issue arises from the coupling between the precompensated WF and the NF intensity. Aberrations introduced in planes not conjugated to the DM lead to precompensation effects that couple the WF to NF intensity variations. Depending on the magnitude of the aberrations, this coupling can generate hotspots, reducing the system's maximum supportable pulse energy. Artificial examples demonstrating this effect are shown in Fig. 17, where deliberate large DM strokes, approx. one order of magnitude higher than in regular operation, illustrate the potential for hotspot generation. Such amplitudes do occur for unsta-

ble loops, but could also occur for large amounts of pre-compensation, even if the loop remains stable. In this case, experiments may still be feasible at reduced energy, while the impact of the distorted NF and thus the three-dimensional intensity distribution around the focal plane on the target interaction has to be considered for each experiment individually.

The WF to NF coupling is more pronounced for smaller beam diameters due to stronger WF gradients for the same phase amplitude. Mitigation strategies therefore should focus on reducing pre-compensation stroke magnitudes. Assuming passive aberration mitigation strategies are optimized, two potential solutions emerge:

(1) Fast Post-Compensating DM: Implementing a fast, large diameter DM after the amplifier for direct aberration post-correction. This requires significant hardware development as larger DMs are slower compared to smaller ones.

(2) Cascaded Compensation: Employing a fast pre-compensating DM combined with a slower post-compensating DM in the same conjugate plane to form a cascaded control system. This would share the correction load and reduce the required pre-compensation stroke.

Both approaches require further investigation in future work, but in general, reducing and possibly eliminating the pre-compensation amplitude as far as possible is the favorable path.

7.3. Implementation of Safety Mechanisms

As discussed previously, a malfunctioning AO loop poses a severe hazard for high-power laser facilities, with potential damage to compressor gratings reaching millions of dollars, alongside significant system downtime. To safely implement RTAO for shot delivery, robust safety mechanisms are essential. A recommended starting point involves developing a comprehensive interlock system:

- Loop-Internal Interlock: A monitoring subroutine that evaluates control loop parameters in real time. Safety thresholds, such as excessive actuator values or sudden changes in DM shape, trigger an abort signal to the control system.

- Loop-External Interlock: An imaging system monitoring NF homogeneity in a critical optical plane. Violations of safety thresholds, such as localized intensity spikes or clipping, also trigger an abort signal.

- Pre-Shot DM Freeze: Before each shot, the DM is locked in its current position, halting the control loop. This interval must balance minimal dynamic aberration evolution and sufficient time for external interlocks to perform a final quality check.

7.4. Integration into Laser Control System

In addition to ensuring machine safety, integrating the RTAO system into the broader laser control system is essential for operability. While CACAO provides a powerful and flexible platform for RTAO, it requires in-depth knowledge of AO concepts, software structure and available commands, making it unsuitable for daily laser operations, where operators need to manage many subsystems simultaneously.

Developing an interface that abstracts common workflows in CACAO into few intuitive commands is critical. This interface would allow operators to execute essential tasks without detailed knowledge of the underlying system. The integration effort is an essential step toward making RTAO a practical tool for routine high-power laser operation.

8. Conclusion

This work represents the first implementation of an RTAO system in a high-power laser, demonstrating its potential to significantly enhance beam stability. This advancement enables more repeatable laser shots, reduces experimental error bars, and supports higher repetition rates.

The ARTAO system relies on an off-spectral pilot beam. The match between the WFs of the beams was found to be excellent, ensuring continuous WF control for pulsed laser systems. ARTAO operates at an acquisition rate of 2.6 kHz and a control speed of 1.3 kHz, achieving a crossover frequency of 96.8 Hz for local wavefront fluctuations. Short-term stabilization results indicate a guaranteed Strehl ratio of >0.96 for a statically fully corrected beam, compared to 0.62 without AO, and robust performance during amplifier pumping events.

The feasibility of RTAO has therefore been established. However, several issues remain which need to be resolved before RTAO can be used in daily operation: First, long-term stability is not yet given as beam drifts may lead to lateral misregistration over time, rendering the loop unstable. Second, robust machine safety mechanisms, including conjugation tracking systems and interlocks, are essential to prevent shot delivery with suboptimal beam quality, which could damage optics in the beam path. Third, simplified software interfaces are necessary as laser operators need to supervise many subsystems at once, requiring a certain level of abstraction. These next steps will be crucial for advancing RTAO technology in large laser facilities.

Once these issues are solved, RTAO will addresses critical challenges in high-power laser operations, particularly stabilizing beam quality for experiments sensitive to laser parameters, such as laser wakefield acceleration or the usage of beams with orbital angular momentum. Large laser facilities, including glass-based high-energy systems, can particularly benefit by actively mitigating temperature-induced fluctuations, enabling shorter cooldown times and improved operational efficiency.

In the long term, we expect RTAO to contribute to advancements in the entire field, including applications such as laser particle accelerators, inertial fusion energy, and integrated solutions for commercial laser systems. This may revolutionize beam stability, becoming a new standard in high-power and high-energy laser facilities.

Acknowledgements

References

1. Frédéric Druon, Gilles Chériaux, Jérôme Faure, John Nees, Marc Nantel, Anatoly Maksimchuk, Gérard Mourou, Jean Christophe Chanteloup, and Gleb Vdovin. Wave-front correction of femtosecond terawatt lasers by deformable mirrors. Opt. Lett., 23(13):1043–1045, Jul 1998. doi: 10.1364/OL.23.001043. URL https://opg.optica.org/ol/abstract.cfm?URI=ol-23-13-1043.

2. Horst Baumhacker, Georg Pretzler, Klaus J Witte, Manuel Hegelich, Malte Kaluza, Stefan Karsch, A Kudryashov, V Samarkin, and A Roukossouev. Correction of strong phase and amplitude modulations by two deformable mirrors in a multistaged ti: sapphire laser. Optics letters, 27(17):1570–1572, 2002.

3. Benoit Wattellier, Julien Fuchs, Ji-Ping Zou, Jean-Christophe Chanteloup, Heidi Bandulet, Pierre Michel, Christine Labaune, Sylvie Depierreux, Alexis Kudryashov, and Alexander Aleksandrov. Generation of a single hot spot by use of a deformable mirror and study of its propagation in an underdense plasma. JOSA B, 20(8):1632–1642, 2003.

4. Vadim Samarkin, Alexander Alexandrov, Gilles Borsoni, Takahisa Jitsuno, Pavel Romanov, Aleksei Rukosuev, and Alexis Kudryashov. Wide aperture piezo-ceramic deformable mirrors for aberration correction in high-power lasers. High Power Laser Science and Engineering, 4:e4, 2016.

5. JB Ohland, U Eisenbarth, B Zielbauer, Y Zobus, D Posor, J Hornung, D Reemts, and V Bagnoud. Ultra-compact post-compressor on-shot wavefront measurement for beam correction at phelix. High Power Laser Science and Engineering, 10:e18, 2022.

6. JB Ohland, D Posor, U Eisenbarth, V Iancu, R Ungureanu, D Ursescu, and V Bagnoud. Zernike-coefficient extraction via helical beam reconstruction for optimization (zehbro) in the far field. High Power Laser Science and Engineering, 11:e86, 2023.

7. Matteo Negro, Martino Quintavalla, Jacopo Mocci, Anna G Ciriolo, Michele Devetta, Riccardo Muradore, Salvatore Stagira, Caterina Vozzi, and Stefano Bonora. Fast stabilization of a high-energy ultrafast opa with adaptive lenses. Scientific reports, 8(1):14317, 2018.

8. Vincent Leroux, Spencer W Jolly, Matthias Schnepp, Timo Eichner, Sören Jalas, Manuel Kirchen, Philipp Messner, Christian Werle, Paul Winkler, and Andreas R Maier. Wavefront degradation of a 200 tw laser from heat-induced deformation of in-vacuum compressor gratings. Optics express, 26(10):13061–13071, 2018.

9. Jonas Benjamin Ohland. Design and setup of a post-compressor adaptive optics loop at PHELIX. PhD thesis, Technische Universität Darmstadt, 2022.

10. N Jovanovic, Frantz Martinache, Olivier Guyon, Christophe Clergeon, Garima Singh, Tomoyuki Kudo, Vincent Garrel, Kevin Newman, D Doughty, Julien Lozi, et al. The subaru coronagraphic extreme adaptive optics system: enabling high-contrast imaging on solar-system scales. Publications of the Astronomical Society of the Pacific, 127(955):890, 2015.

11. Engin Akyol, Ahmed M Hagag, Sobha Sivaprasad, and Andrew J Lotery. Adaptive optics: principles and applications in ophthalmology. Eye, 35(1):244–264, 2021.

12. Fernando P Guiomar, Marco A Fernandes, José Leonardo Nascimento, Vera Rodrigues, and Paulo P Monteiro. Coherent free-space optical communications: Opportunities and challenges. Journal of Lightwave Technology, 40(10):3173–3186, 2022.

13. Alexis V Kudryashov, Alexey Rukosuev, Vadim Samarkin, Ilia Galaktionov, and Evgeniy Kopylov. Fast adaptive optical system for 1.5 km horizontal beam propagation. In Unconventional and Indirect Imaging, Image Reconstruction, and Wavefront Sensing 2018, volume 10772, pages 250–258. SPIE, 2018.

14. K. Burdonov, A. Fazzini, V. Lelasseux, J. Albrecht, P. Antici, Y. Ayoul, A. Beluze, D. Cavanna, T. Ceccotti, M. Chabanis, A. Chaleil, S. N. Chen, Z. Chen, F. Consoli, M. Cuciuc, X. Davoine, J. P. Delaneau, E. d'Humières, J.-L. Dubois, C. Evrard, E. Filippov, A. Freneaux, P. Forestier-Colleoni, L. Gremillet, V. Horny, L. Lancia, L. Lecherbourg, N. Lebas, A. Leblanc, W. Ma, L. Martin, F. Negoita, J.-L. Paillard, D. Papadopoulos, F. Perez, S. Pikuz, G. Qi, F. Quéré, L. Ranc, P.-A. Söderström, M. Scisciò, S. Sun, S. Vallières, P. Wang, W. Yao, F. Mathieu, P. Audebert, and J. Fuchs. Characterization and performance of the apollon short-focal-area facility following its commissioning at 1 pw level. Matter and Radiation at Extremes, 6(6):064402, 10 2021. ISSN 2468-2047. doi: 10.1063/5.0065138. URL https://doi.org/10.1063/5.0065138.

15. DN Papadopoulos, Y Ayoul, A Beluze, F Gobert,

D Mataja, A Fréneaux, N Lebas, M Chabanis, M Dumergue, P Audebert, et al. The apollon laser facility upgrade to the multi-pw level. In High Intensity Lasers and High Field Phenomena, pages HTu2B–2. Optica Publishing Group, 2024.

16. Jonas B Ohland, Udo Eisenbarth, Markus Roth, and Vincent Bagnoud. A study on the effects and visibility of low-order aberrations on laser beams with orbital angular momentum. Applied Physics B, 125:1–10, 2019.

17. Julien Lozi, Kyohoon Ahn, Hannah Blue, Alicia Chun, Christophe Clergeon, Vincent Deo, Olivier Guyon, Takashi Hattori, Yosuke Minowa, Shogo Nishiyama, Yoshito Ono, Shin Oya, Yuhei Takagi, Sebastien Vievard, and Maria Vincent. Ao3k at subaru: First on-sky results of the facility extreme-ao, 2024. URL https://arxiv.org/abs/2407.19188.

18. François Rigaut and Benoit Neichel. Multiconjugate adaptive optics for astronomy. Annual Review of Astronomy and Astrophysics, 56(1):277–314, 2018.

19. Robert K Tyson and Benjamin West Frazier. Principles of adaptive optics. CRC press, 2022.

20. Olivier Guyon, Arnaud Sevin, Damien Gratadour, Julien Bernard, Hatem Ltaief, Dalal Sukkari, Sylvain Cetre, Nour Skaf, Julien Lozi, Frantz Martinache, et al. The compute and control for adaptive optics (cacao) real-time control software package. In Adaptive Optics Systems VI, volume 10703, pages 469–480. SPIE, 2018.

21. Olivier Guyon, Arnaud Sevin, Florian Ferreira, Hatem Ltaief, Jared Males, Vincent Deo, Damien Gratadour, Sylvain Cetre, Frantz Martinache, Julien Lozi, et al. Adaptive optics real-time control with the compute and control for adaptive optics (cacao) software framework. In Adaptive Optics Systems VII, volume 11448, pages 574–586. SPIE, 2020.

22. Vincent Deo, Olivier Guyon, Jared Males, Kyohoon Ahn, Brian Carcich, Sylvain Cetre, Florian Ferreira, Damien Gratadour, Janis Hagelberg, R Jansen-Clem, et al. The cacao real-time computer for adaptive optics: updates, performance, and development plans. In Adaptive Optics Systems IX], Jackson, K., Schmidt, D., and Vernet, E., eds., Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, volume 13097, page 13097120, 2024.

23. Stefano Bonora, Jan Pilar, Antonio Lucianetti, and Tomas Mocek. Design of deformable mirrors for high power lasers. High power laser Science and Engineering, 4:e16, 2016.

24. Avinash Surendran, Mahesh P Burse, Anamparambu N Ramaprakash, Jyotirmay Paul, Hillol K Das, and Padmakar S Parihar. Scalable platform for adaptive optics real-time control, part 1: concept, architecture, and validation. Journal of Astronomical Telescopes, Instruments, and Systems, 4(3):039001–039001, 2018.

25. Alexey L Rukosuev, Vadim N Belousov, Ilya V Galaktionov, Alexis V Kudryashov, Alexander N Nikitin, Vadim V Samarkin, and Julia V Sheldakova. 1.5 khz adaptive optical system for free-space communication tasks. In Free-Space Laser Communications XXXII, volume 11272, pages 331–340. SPIE, 2020.

26. Tuan N Truong, Antonin H Bouchez, Richard G Dekany, Jean C Shelton, Mitchell Troy, John R Angione, Rick S Burruss, John L Cromer, Stephen R Guiwits, and Jennifer E Roberts. Real-time wavefront control for the palm-3000 high order adaptive optics system. In Adaptive Optics Systems, volume 7015, pages 995–1002. SPIE, 2008.

27. Janis Hagelberg, Nathanaël Restori, François Wildi, Bruno Chazelas, Christoph Baranec, Olivier Guyon, Ludovic Genolet, Michaël Sordet, and Reed Riddle. Kalao the swift adaptive optics imager on the 1.2 m euler swiss telescope in la silla, chile. In Adaptive Optics Systems VII, volume 11448, pages 1460–1467. SPIE, 2020.

28. Lisa A. Poyneer. Scene-based shack-hartmann wavefront sensing: analysis and simulation. Appl. Opt., 42(29):5807–5815, Oct 2003. doi: 10.1364/AO.42.005807. URL https://opg.optica.org/ao/abstract.cfm?URI=ao-42-29-5807.

29. A N Kolmogorov. Local structure of turbulence in an incompressible viscous fluid at very high reynolds numbers. Soviet Physics Uspekhi, 10(6):734, jun 1968. doi: 10.1070/PU1968v010n06ABEH003710. URL https://dx.doi.org/10.1070/PU1968v010n06ABEH003710.

30. Berdeu, A., Bonnet, H., Le Bouquin, J.-B., Édouard, C., Gomes, T., Shchekaturov, P., Dembet, R., Paumard, T., Oberti, S., Kolb, J., Millour, F., Berio, P., Lai, O., Eisenhauer, F., Garcia, P., Straubmeier, C., Kreidberg, L., Hönig, S. F., and Defrère, D. Estimation of the lateral mis-registrations of the gravity+ adaptive optics system - perturbative method with open-loop modal correlation and non-perturbative method with temporal correlation of closed-loop telemetry. A&A, 687:A157, 2024. doi: 10.1051/0004-6361/202449311. URL https://doi.org/10.1051/0004-6361/202449311.

A. Real-Time Computer Assembly

The components that we used to construct the RTC that we used in this work are listed in Table 2. These are consumer grade components and are available at a relatively low price. The total cost for the RTC, as for the time of writing, is approximately 3,000 € (compared to tens of thousands for the SHS and DM). The entire setup was heavily inspired by the RTC of the KalAO system[27].

It is noteworthy that all available PCIe extension slots are occupied in this setup, so further extentions would require a different choice of hardware:

Component	Description	Manufacturer, Model
Processor	AM5 socket, 16 Cores, <5.5 GHz	AMD Ryzen 9 7950X
Motherboard	AM5 socket, 2 PCIe 5.0 slots from CPU	MSI MPG X670E Carbon WIFI
GPU	12 GB GDDR6X, 10240 CUDA cores, 1725 MHz	Palit GeForce RTX 3080Ti GameRock
RAM	2x16 GB DDR5-4800	Kingston FURY Beast RGB 16 TB (x2)
NVMe SSD	PCIe 4.0, 2 TB	Kingston KC3000
Network card	Intel 82576 chip, GBit Ethernet, Dual RJ45 Ports	*Unknown*
Chassis	19" rack mountable	SilverStone RM42-502-B
CPU cooler	Water cooling, 2x 240 mm radiator	SilverStone SST-PF240-ARGB-V2
Chassis fans	2x 80 mm high-throughput fans	Alphacool ES 80 mm 800-6000 rpm
Power unit	1 kW, gold standard	be quiet! Pure Power 11 FM 1000 W
Hard drive	8 TB for mass storage	Seagate BarraCuda 8 TB

Table 2. Component list for the RTC used in this work.

(1) **PCI_E1:** PCIe 5.0x8 (From CPU), occupied by GPU

(2) **PCI_E2:** PCIe 5.0x8 (From CPU), occupied by PCIe camera card

(3) **PCI_E3:** PCIe 4.0x4 (From Chipset), occupied by Ethernet card for DM

The RTC is running unter Ubuntu 20.04 LTS.