

AO for laser fusion: the challenges of controlling hundreds of high-power laser beams

Jonas B. Ohland^a, Leon Dauerer^a, Nathalie Lebas^b, Vincent Déo^c, Olivier Guyon^{c,d,e,f}, Dimitris Papadopoulos^b, and Vincent Bagnoud^{a,g,h}

^aPlasma Physics/PHELIX, GSI Helmholtz Center for Heavy Ion Research GmbH, Planckstr. 1, 64289 Darmstadt, Germany

^bLaboratoire pour l'Utilisation des Lasers Intenses, École polytechnique, Av. Fresnel, 91120 Palaiseau, France

^cSubaru Telescope, National Astronomical Observatory of Japan, 650 North A'ohoku Place, Hilo, Hawaii 96720, U.S.A.

^dAstrobiology Center, NINS, National Astronomical Observatory of Japan, 2-21-1 Osawa, Mitaka-shi, Tokyo 181-8588, Japan

^eSteward Observatory, University of Arizona, 933 N Cherry Ave, Tucson, AZ 85719, U.S.A.

^fUniversity of Arizona, 1200 E University Blvd, Tucson, AZ 85721, U.S.A.

^gInstitute for Applied Physics, Technical University of Darmstadt, Hochschulstraße 4a, Darmstadt 64289, Germany

^hHelmholtz Institute Jena, Fröbelstieg 3, 07743 Jena, Germany

ABSTRACT

Laser-driven inertial fusion requires high-energy laser drivers that are faster, more stable, and more repeatable than present landmark ignition systems. Depending on the driver concept, such systems may require hundreds to thousands of beamlines, making beam control a problem not only of stability and repeatability, but also of safety and operability. Adaptive optics is a powerful building block for this task, but it is not a complete solution by itself. This paper connects laser-damage constraints, driver architecture, aberration sources, correction placement, reference degradation, and facility-level operation into one beam-control problem. Using the Apollon Real-Time Adaptive Optics system as a case study, it argues that single-conjugate pre-compensation is insufficient for fusion-relevant drivers and that adaptive optics must become part of an integrated Adaptive Laser Architecture for sensing, correction, and automation.

Keywords: Laser Fusion, Inertial Fusion Energy, Adaptive Optics, Laser Beam Control, High Energy Lasers, High Power Lasers, Power Plant, Adaptive Laser Architecture

1. INTRODUCTION

Laser-driven Inertial Confinement Fusion (ICF) is one route toward commercial fusion energy. The December 2022 ignition result at the National Ignition Facility (NIF) marked a major milestone.¹ Still plenty of engineering challenges remain on the path towards Inertial Fusion Energy (IFE), many of which remain in the domain of the drive lasers. A suitable driver must operate with much higher repetition rate, efficiency, thermal-management capability, stability, and practical availability than present single-shot or low-repetition high-energy laser systems.

Efficiency-wise, direct-drive concepts are attractive because they can couple laser energy more directly to the target, but they also impose stricter demands on beam uniformity and repeatability.² Combined with the scale of a full ICF power plant, where hundreds to thousands of drive lasers may need to operate in lockstep, robustness and operability become primary design constraints alongside facility cost.

Further author information: (Send correspondence to J.B.O.)

J.B.O.: E-mail: j.b.ohland@gsi.de, Telephone: +49 (0)61 59 71 33 81

Adaptive Optics (AO) is directly relevant to this problem, but future fusion drivers require more than isolated deformable-mirror loops. They require an Adaptive Laser Architecture (ALA): an integrated beam-control concept combining distributed diagnostics, actuators, hierarchical correction strategy, automation, and supervisory control. In this view, AO is a major component inside a broader architecture, not a standalone answer.

This paper argues that high-energy pulsed laser beam control for fusion-relevant drivers must be treated as an integrated, system-level problem. In this frame, the most imminent unresolved issues are correction placement, reference degradation and approaches to an integrated control framework.

Section 2 introduces the laser-architecture constraints that make these issues unavoidable, while section 3 uses Apollon Real-Time Adaptive Optics (ARTAO) as a practical case study. We generalize the obtained lessons into three main beam-control challenges in section 4, from which we derive four associated technology gaps in section 5. As a response, we launched the Adaptive Laser Architecture Development and INtegration (ALADIN) project, which we briefly introduce in section 6, before Section 7 concludes the paper.

2. DRIVER-LASER ARCHITECTURE AND BEAM-QUALITY CONSTRAINTS

Large-scale high-energy pulsed laser architecture is shaped first and foremost by optical damage limits. The Laser-Induced Damage Threshold (LIDT) governs acceptable fluence and therefore beam diameter and optic size at a given pulse energy and duration. Beams must be large enough to keep the fluence below the LIDT, while optics must remain as small as cost, handling, and engineering practicality allow. Beam homogeneity is part of this constraint: a more homogeneous beam allows more usable energy before local fluences reach the LIDT.^{3,4} These constraints translate directly into two common architectural paradigms in high-energy lasers: the Master-Oscillator-Power-Amplifier (MOPA) design and relay imaging, which we briefly introduce in the following to clarify their implications for laser beam control.*

A typical high-energy architecture follows the MOPA design: an engineered seed pulse passes through a chain of laser amplifiers, between which beam diameter increases with pulse energy to keep the peak fluence below the LIDT. Large final amplifiers then become important sources of static and dynamic aberration.

Relay imaging is used simultaneously to preserve defined beam planes over long propagation. Here, the usage of Keplerian telescopes can combine beam expansion with a relay-imaging function. Figure 1 provides a simplified MOPA and relay-imaging schematic.




Figure 1. Simplified schematic of a MOPA and relay-imaging laser architecture. Beam expansion follows increasing pulse energy to limit fluence, while relay telescopes preserve selected beam planes through the amplifier chain. Schematic by J.B. Ohland, replicated with the author's permission.³

Relevant aberration sources include deformations and turbulence caused by heat deposition in and around the active media of amplifier heads, as well as cooling-induced turbulence and airflow in beam transport. Further disturbances arise from environmental drift caused by temperature or humidity changes and from mechanical vibration generated internally or transmitted through the ground.³

The effects of these aberrations during propagation are critical in high-energy lasers. The most prominent is the degradation of the target intensity distribution, that is, the point-spread function of each beam on the fusion pellet. In extreme cases, aberrations can also evolve into near-field intensity modulations that increase the risk of optical damage inside the laser chain. Large-aperture sections of the beam transport can make major

*The driver architecture considered here concerns high-energy, nanosecond-class pulsed lasers rather than ultrafast laser systems. Chirped-Pulse Amplification (CPA), which is central to many ultrafast high-intensity lasers, is therefore not required for the architecture discussed here.

contributions to dynamic aberrations other than tip and tilt, somewhat analogous to the relationship between telescope aperture and the spatial frequencies of atmospheric turbulence sampled across it.

This architecture context defines the practical question for adaptive optics: where should correction be placed, and what does the measured wavefront actually represent? The next section uses ARTAO as a concrete high-energy pulsed laser case study.

3. ARTAO CASE STUDY

The Apollon Real-Time Adaptive Optics (ARTAO) system provides a useful case study because it demonstrates real-time AO in a high-power laser hall while also exposing the limits of the chosen correction placement on its own. Apollon is a high-intensity laser facility, not a fusion driver, but the architectural overlap in MOPA and relay-imaging makes the ARTAO experience directly relevant to fusion-laser beam control. The system combined a pilot beam, a fast Shack-Hartmann Wavefront Sensor (WFS) connected through PCIe, a 6 cm diameter bimorph Deformable Mirror (DM) placed early in the beamline, and a Compute And Control for Adaptive Optics (CACAO)-based real-time controller.^{5,6}

The design was driven by practical constraints: a large, fast DM suitable for late post-correction was not available within the budget and technology envelope. For fusion scaling, such mirrors are also unlikely to be realistic for hundreds to thousands of beamlines. ARTAO therefore used a smaller DM early in the chain and relay-imaged the correction to the WFS at the end of the amplification chain for direct pre-compensation. This choice made the system compact and feasible, but it also made the result sensitive to downstream propagation. Figure 2 shows the setup.




Figure 2. Schematic of the ARTAO setup with pilot beam injection, pre-compensating DM, and wavefront sensor within the Apollon laser chain. Schematic by J.B. Ohland et al., reproduced with the author's permission.⁵

The system showed that astronomy-inspired real-time control concepts can transfer to a large laser environment. Closed-loop operation on the pilot beam reduced Strehl-ratio fluctuations at the WFS reference from roughly 30% to roughly 5% (see Figure 3).

The same experiment also showed why pre-compensation cannot be treated as a general solution. Because Apollon had insufficient relay imaging from the main amplifier onward due to geometric constraints, strong correction could produce downstream near-field shifts, warping, and intensity modulation. In other words, improving the wavefront at the WFS reference did not guarantee a safe or optimal beam state everywhere downstream. Figure 3 includes examples of such near-field modulation under strong actuator strokes.

The most important observed limitation, however, was reference degradation. The long propagation through the laser chain made the AO setup susceptible to near field distortions, with lateral misregistration caused by drift as a clear example. Such misregistration can be detected and controlled through telemetry-based estimation methods,⁷ but was not implemented at the time, imposing a severe risk to the machine since a rogue correction loop may generate hotspots on optical surfaces. Higher-order reference degradation, on the other hand, is likely to become important at larger disturbance amplitudes or in more complex amplifier environments. Proper relay imaging between aberration sources and correctors can mitigate this problem, but only if the disturbance source remains localized.⁸




Figure 3. ARTAO performance and limitations. (a) Pilot-beam Strehl-ratio time series with the loop closed at $t = 0$. (b) Histogram of the corresponding open- and closed-loop data. (c) Near-field fluence imprints produced by deliberately large random DM strokes in a non-conjugate image plane. Panels reproduced from J.B. Ohland et al. with the agreement of the authors.⁵

The ARTAO lesson is conclusive: real-time AO can work in a high-power laser hall, but single-conjugate pre-compensation is not enough. This leads directly to the broader driver problem: the imaging setup, correction placement, sensing, and control hierarchy must be designed together.

4. GENERALIZED BEAM-CONTROL CHALLENGES

On a more general level, ARTAO underlines that practical driver-laser beam control has to be approached on a higher level. The actuator pool, first and foremost, must be broader than deformable mirrors alone. It can include piezoelectric and mechanical DMs, fast steering mirrors, slow motorized mirrors, motorized lenses, bendable low-order mirrors, and motorized off-axis-parabola or telescope elements. Using this pool, propagation effects can be mitigated by correcting disturbances close to their source where possible, reducing the load on pre-compensating DMs.

These thoughts provoke the following generalized beam-control challenges:

Control Placement and Integration. Pre-compensation alone is comparatively inexpensive and easy to install, but it is sensitive to propagation effects. Correction after amplification is more direct but requires large, expensive, and likely slower components, while remaining unable to prevent near-field distortions that develop before the correction plane. Distributed control is therefore likely to be required, with suitable control logic deciding how actuator load is assigned, which requires an integrated actuator hierarchy. Correction close to the source, for example in multi-pass or multi-slab amplifier contexts, is an important related design direction.⁸ Furthermore, machine safety remains a concern, since a rogue correction loop may permanently damage the beamline.

Reference and State Inference. The second challenge is reference and state inference. A single WFS cannot locate aberration planes in a large beamline, and large-volume amplifiers can create distributed disturbances.

A pilot beam can be an effective reference, but insight remains limited: unlike astronomy, there is no practical multi-conjugate analogue based on several guide-star directions, because the laser beam behaves like a single reference direction. The disturbance planes within the full beamline are therefore not measured directly, making reference degradation difficult to model.

Defining the actual target wavefront is challenging in its own right because relevant diagnostics are either indirect or unavailable at full operating conditions due to the destructive energy levels at the target plane. A correction that is optimal at one diagnostic plane can therefore be wrong downstream. The problem is not only observability in an abstract sense; it is the concrete risk (and more often than not, a reality) that the measured reference no longer represents the beam state that matters for interaction with the fusion target.

Facility-Level Coordination. The same issue scales from one beamline to a facility. Even a single beamline is a complex, nonlocal optimization and stabilization problem because diagnostics, actuators, relay planes, and downstream propagation are coupled. In a multi-beamline driver, set points and states must be coordinated across the facility. Startup, alignment, recovery, missed-shot handling, target-position fluctuation, and faulty-beamline compensation cannot rely on manual intervention. Robustness and replicated hardware cost therefore become part of the control problem.

5. TECHNOLOGY GAPS

From these three generalized beam-control challenges in fusion power plants, we derive four technology gaps. The first two are the main emphasis of this paper because they are most directly visible in ARTAO and in the near-term control-architecture problem.

Gap 1: Online Beam Stabilization and Conditioning. This is the fast local layer that stabilizes a single beamline before facility-level scaling is considered. It may address pointing, wavefront, spatial homogeneity, and pulse shape as needed. The control strategies are not yet clear, but will likely range from classical Proportional-Integral-Derivative (PID) controllers to more elaborate schemes such as cascaded control within a defined hierarchy, predictive control for pulsed operation, smart or self-updating controllers, and other data-driven methods as useful. ARTAO supports this gap most directly because it demonstrates real-time wavefront stabilization in a high-power laser environment.⁵

Gap 2: System Modeling and Supervisory Control. This is the orchestration layer above local stabilization loops. It must represent the beamline, estimate its state, infer disturbance origins, manage loop interactions, balance actuator loads, and connect beamline control to facility control. It must also translate high-level facility commands into coordinated beamline set points. Together with Gap 1, this is among the most immediate challenges to be solved.

Gap 3: Facility Operability and Automation. A driver with hundreds or thousands of beamlines cannot depend on manual operation. It needs automated alignment, ramp-up and restart routines, health monitoring through anomaly detection and fault diagnosis, automatic error handling, and machine- and facility-safety strategies. The goal is not only better performance, but safe and efficient day-to-day operation with minimal human intervention.

Gap 4: Experiment-Informed Beam Optimization. This is the interface between beam control and fusion capsule interaction. It requires connecting beamline actuators to target-chamber diagnostics and optimizing fusion-relevant observables and ultimately reliability, stability and fusion efficiency. Possible tools include Bayesian optimization and closed-loop shot tuning, but this gap depends on more mature automation, integration, and target-plane information. It can be addressed once Gaps 1 and 2 are sufficiently mature to support autonomous beamline operation.

6. THE ALADIN PROJECT

The Adaptive Laser Architecture Development and Integration (ALADIN) project is one concrete response to the beam-control gaps identified above. It launched in June 2025, and its five-year funding period is intended to establish the basis for a sustained effort to provide fusion-ready beam-control solutions.



Its scope primarily covers the first two technology gaps. Within this scope, we will explore beam-control strategies for fusion driver lasers on a dedicated emulator testbed while developing an open-source integration tool called COntroller for Adaptive Laser Architecture (COALA). This software is intended to provide an integration layer for established distributed-control and Supervisory Control and Data Acquisition (SCADA) frameworks such as EPICS, Tango Controls, and Karabo, providing a flexible way to set up, link, and orchestrate control loops in a state-sensitive fashion.^{9–11}

The longer-term vision is a scalable beam-control architecture for hundreds of stable high-power beams, with hardware and software co-developed alongside the driver architecture rather than retrofitted after the fact.

7. CONCLUSION

Laser fusion pushes adaptive optics beyond the familiar regime. Fusion-relevant drivers are shaped by LIDT, which implies an architecture centered around the MOPA concept, including relay imaging and careful fluence management. Their aberration sources are distributed through large, energetic beamlines, and propagation can turn phase correction into intensity modulation that poses a risk to beamline integrity. Beam control must therefore be treated as an architectural problem, not only as a local wavefront-correction problem.

ARTAO shows both sides of this message. It demonstrates that real-time AO can operate in a high-power laser context, but it also shows that single-conjugate pre-compensation is not sufficient when the WFS reference degrades relative to the downstream beam state. Reference degradation is therefore the most urgent conceptual limitation for extending AO toward fusion-relevant drivers.

Future systems need an integrated beam-control architecture: diagnostics, actuators, automation, supervisory control, and realistic hardware costs must be designed together. ALADIN is an early step in this direction. The field now needs exchange with the AO community on feasible control schemes, plane-sensitive wavefront sensing, self-updating calibration, robust operation, and other relevant topics.

This endeavor relies on building a network that provides ideas, guidance, constructive critique, and joint trials on realistic high-power-laser platforms. It may become a prime example of interdisciplinary collaboration that helps bring clean fusion energy within reach. Under this vision, the ALADIN project remains open to collaboration.

USE OF GENERATIVE AI TOOLS

OpenAI GPT-5.3 to GPT-5.5 were used to assist with structuring author-provided material, as well as language and grammar polishing. The authors reviewed, revised, and independently verified all AI-assisted content and take full responsibility for the manuscript. The tools were not used to generate scientific data, analyses, or figures.

ACKNOWLEDGMENTS

The ARTAO project was funded in the framework of the Technology for High-Repetition-rate Intense Laser Laboratories (THRILL) project by the European Union's HORIZON-INFRA-2022-TECH-01 call under grant agreement number 101095207.

The ALADIN project has received funding from the German Federal Ministry for Research, Technology and Space (BMFTR) within the framework of “Fusion 2040 – Research on the Way to a Fusion Power Plant” under grant agreement No. 13F1023.

REFERENCES

- [1] Abu-Shawareb, H., Acree, R., Adams, P., Adams, J., Addis, B., Aden, R., Adrian, P., Afeyan, B. B., Aggleton, M., et al., “Achievement of target gain larger than unity in an inertial fusion experiment,” *Physical Review Letters* **132**(6), 065102 (2024).
- [2] Craxton, R. S., Anderson, K. S., Boehly, T. R., Goncharov, V. N., Harding, D. R., Knauer, J. P., McCrory, R. L., McKenty, P. W., Meyerhofer, D. D., Myatt, J. F., Schmitt, A. J., Sethian, J. D., Short, R. W., Skupsky, S., Theobald, W., Kruer, W. L., Tanaka, K., Betti, R., Collins, T. J. B., Delettrez, J. A., et al., “Direct-drive inertial confinement fusion: A review,” *Physics of Plasmas* **22**, 110501 (2015).
- [3] Ohland, J. B., *Design and Setup of a Post-Compressor Adaptive Optics Loop at PHELIX*, PhD thesis, Technische Universität Darmstadt (2022).
- [4] Danson, C. N., Haefner, C., Bromage, J., Butcher, T., Chanteloup, J.-C. F., Chowdhury, E. A., Galvanauskas, A., Gizzi, L. A., Hein, J., Hillier, D. I., et al., “Petawatt and exawatt class lasers worldwide,” *High Power Laser Science and Engineering* **7**, e54 (2019).
- [5] Ohland, J. B., Lebas, N., Deo, V., Guyon, O., Mathieu, F., Audebert, P., and Papadopoulos, D., “Apollon real-time adaptive optics: Astronomy-inspired wavefront stabilization in ultraintense lasers,” *High Power Laser Science and Engineering* **13**, e29 (2025).
- [6] Guyon, O., Sevin, A., Ferreira, F., Ltaief, H., Males, J., Deo, V., Gratadour, D., Cetre, S., Martinache, F., Lozi, J., Vievard, S., Fruitwala, N., Bos, S., and Skaf, N., “Adaptive optics real-time control with the compute and control for adaptive optics (cacao) software framework,” in [*Adaptive Optics Systems VII*], Schreiber, L., Schmidt, D., and Vernet, E., eds., *Proceedings of SPIE* **11448**, 114482N (2020).
- [7] Berdeu, A., Bonnet, H., Le Bouquin, J.-B., Edouard, C., Gomes, T., Shchekaturov, P., Dembet, R., Pau-mard, T., Oberti, S., Kolb, J., Millour, F., Berio, P., Lai, O., Eisenhauer, F., Garcia, P., Straubmeier, C., Kreidberg, L., Honig, S. F., and Defrere, 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,” *Astronomy & Astrophysics* **687**, A157 (2024).
- [8] Paliesek, T., Divoky, M., Pilar, J., Smrz, M., and Mocek, T., “Adaptive optics design for high-energy kw-class multi-slab laser amplifiers,” *High Power Laser Science and Engineering* **13**, e102 (2025).
- [9] Dalesio, L. R., Hill, J. O., Kraimer, M., Lewis, S., Murray, D., Hunt, S., Watson, W., Clausen, M., and Dalesio, J., “The experimental physics and industrial control system architecture: past, present, and future,” *Nuclear Instruments and Methods in Physics Research Section A* **352**(1–2), 179–184 (1994).
- [10] Tango Controls Collaboration, “Tango controls,” (2026). Accessed 2026-06-09.
- [11] Hauf, S., Heisen, B., Aplin, S., Beg, M., Bergemann, M., Bondar, V., Boukhelef, D., Danilevsky, C., Ehsan, W., Essenov, S., et al., “The karabo distributed control system,” *Journal of Synchrotron Radiation* **26**(5), 1448–1461 (2019).