Improved laser-plasma accelerator stability via high-bandwidth longitudinal focal position stabilization of a 100 TW-class laser system

K. Jensen,* S. K. Barber, C. Berger, C. E. Doss, F. Kohrell, and J. van Tilborg
Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA

S. J. Coleman, N. M. Cook,† J. Edelen, and J. Einstein-Curtis
RadiaSoft LLC, Boulder, Colorado 80301, USA

(Received 16 May 2025; accepted 27 August 2025; published 26 September 2025)

Laser-plasma accelerators (LPAs) offer an attractive alternative to conventional accelerators for the development of compact electron sources and next-generation light sources. Due to orders-of-magnitude larger accelerating gradients, LPAs enable the acceleration of high-brightness electron beams to ultra-relativistic energies in millimeter- to centimeter-scale distances. However, LPA stability is limited by shot-to-shot fluctuations of the driving laser system. Specifically, fluctuations in the final-focus longitudinal position result in correlated instability in LPA electron beam qualities, including total beam charge, average beam energy, and energy spread. We demonstrate active stabilization of the longitudinal focal position for a 100 TW-class laser system. This repetition-rate scalable stabilization system leverages noninvasive wave front monitoring of a copropagating, unamplified kHz pulse train to guide corrective adjustments to the focal position of a 1 Hz amplified drive laser via an upstream telescope on millisecond timescales. In this demonstration, the approach limits standard deviation fluctuations of the amplified drive beam's longitudinal focal position to ± 0.22 mm (representing at least a 53% reduction), achieved with a correction bandwidth that reaches the Nyquist frequency limit. Consequently, we observe marked improvements in both long-term and shot-to-shot LPA stability in terms of charge and spectrum.

DOI: 10.1103/j5md-sckl

I. INTRODUCTION

Laser-plasma accelerators (LPAs) [1,2] present key opportunities to develop high energy, compact accelerators for a variety of uses in industry, medicine, and basic science research. Community efforts have advanced the state of the art toward enabling the production of multi-GeV [3–6], high-quality [7–12] electron beams suitable for wide-reaching applications [13–16]. Improvements in peak energy, bunch charge, energy spread, and stability specifically motivate the use of LPAs as compact drivers for free-electron lasers (FELs) [17,18], with recent demonstrations of FEL gain [19–21] and ongoing developments at facilities worldwide [22].

Despite recent advancements in LPA technologies, achieving consistent shot-to-shot electron beam stability remains a crucial challenge to enable their integration into both near- and long-term applications [23]. As the accelerating cavity within the plasma of an LPA is re-established each shot, instabilities in LPA electron beam qualities (e.g., total charge, beam pointing, and spectral distribution) are highly correlated to fluctuations in drive laser pulse properties (e.g., pulse energy, beam pointing, focal position, and spectral aberrations). In particular, large-scale statistical studies [24] of LPA performance have identified fluctuations in the longitudinal focal position of the drive laser beam as a dominant source of instability in generated electron beam energy spectra.

Laser jitter manifests across a broad frequency spectrum but can generally be classified into three categories [25]: thermal effects and nonlinearities of the laser system (< 1 Hz), air turbulence (< 10 Hz), and mechanical vibrations (> 10 Hz). Fluctuations in the longitudinal focal position of the drive laser are largely dominated by thermal effects and nonlinearities of the laser system but may also include contributions due to air turbulence [24], thus requiring stabilization with a correction bandwidth exceeding the repetition rate of most high-power titanium-sapphire (Ti:Sa) laser systems (approximately a few Hz). Development of active stabilization techniques pushing 100 Hz to kHz feedback rates (enabling correction bandwidths fundamentally limited to half that according to the Nyquist theorem) aims to synergistically improve the stability of current LPA systems while preparing to meet

*Contact author: jensenke17@gmail.com
†Contact author: ncook@radiasoft.net

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

the demands of future high repetition rate systems [15,26,27]. Recent efforts have addressed challenges in stabilizing the drive beam's pointing and transverse focal position [28,29], resulting in improvements in electron beam pointing stability. The work presented here addresses the challenge of active correction of the drive beam's longitudinal focal position.

Here, we demonstrate a scheme for active stabilization of the drive beam's longitudinal focal position in a 100 TW-class LPA. The closed-loop stabilization system reduces longitudinal focal position fluctuations of the 1 Hz amplified drive laser by ∼53% with a correction bandwidth reaching the Nyquist frequency limit. As a result, relative fluctuations in LPA electron bunch charge are reduced by ∼45%, absolute fluctuations in average bunch energy are reduced by ∼52%, and absolute fluctuations in bunch energy spread are reduced by ∼37%. In Sec. II, we discuss the BELLA Center's Hundred Terawatt Undulator (HTU) beamline. In Sec. III, we discuss the focal stabilization system implementation and noninvasive monitoring of the longitudinal focal position. In Sec. IV, we discuss the performance of the stabilization system during both low- and high-power operation of the laser system and its corresponding effect on LPA electron beam qualities. Finally, in Sec. V, we conclude with a discussion of this demonstration and the next steps for future improvement.

II. EXPERIMENTAL CONFIGURATION

In this section, we describe the experimental layout of the BELLA Center's Hundred Terawatt Undulator (HTU) beamline, including key aspects of the laser system, general LPA operation, and electron beam characterization. A schematic of the beamline and stabilization system implementation is shown in Fig. 1 with dedicated discussion of the stabilization system in Sec. III.

High-power laser pulses are achieved via a chirped-pulse amplification scheme [30]. A Ti:Sa regenerative amplifier system provides a 1 kHz train of 800 nm, 2 mJ laser pulses. Further amplification to multi-Joule pulse energies is achieved via three multipass Ti:Sa amplifiers pumped by a 532 nm frequency-doubled neodymium-doped yttrium aluminum garnet (Nd:YAG) laser at 1 Hz frequency. Critical for high-bandwidth feedback, the unamplified

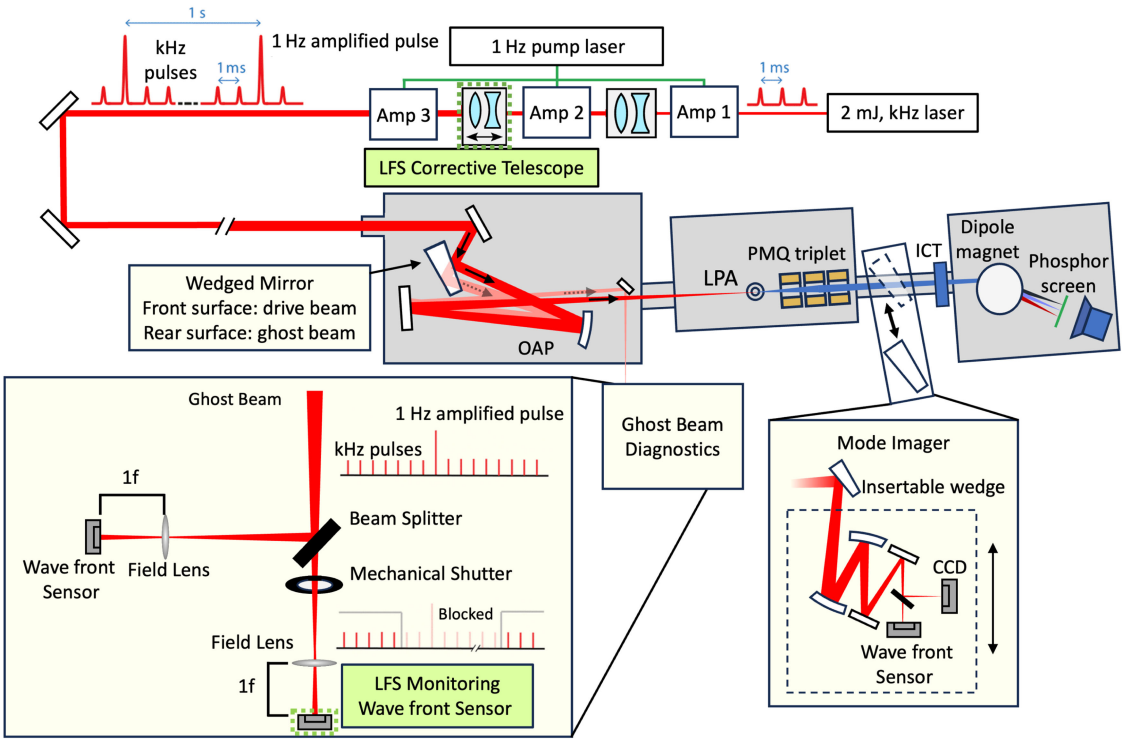


FIG. 1. Experimental configuration of the BELLA HTU beamline. A copropagating pulse train (red) of 1 Hz amplified and kHz unamplified pulses is established. The laser is focused by an OAP and the 1 Hz amplified pulses drive the LPA. LPA electron (blue) charge and spectrum are characterized via an ICT and dipole magnetic spectrometer. A wedged optic prior to the OAP creates an angularly offset, correlated copy of the drive beam focus, which is sent to the ghost beam diagnostic setup for focal position monitoring. Active stabilization of the longitudinal focal position is guided by feedback from a ghost beam WFS and performed using a corrective telescope between amplifiers 2 and 3 (labeled by light green text boxes and highlighted by green dashed outlines). An insertable wedge located between the PMQ triplet and ICT directs the laser to the translatable mode imager setup for imaging of the focal and near-focus planes. The mode imager setup allows for calibration of the true focal position with respect to the ghost beam RoC measurement.

pulses at kHz repetition rate copropagate with the 1 Hz amplified pulses. Between each amplifier stage exists a beam expander. The beam expander between amplification stages two and three was modified by adding a fast actuating stage to control the separation distance between the two lenses. After amplification and compression, the beam is focused using an $f/30$ off-axis parabola (OAP) and delivered on target with 2.5 J pulse energy, 35 fs full width at half maximum (FWHM) pulse duration, and 40 μm FWHM focal spot size for a peak intensity of $I \simeq 3.7 \times 10^{18}$ W cm^{-2} and normalized laser strength parameter of $a_0 \simeq 1.3$. Further documentation describing the laser system is provided in Ref. [28].

Laser wakefield accelerated electron beams were generated via shock-induced density downramp injection [8,31–37] by irradiating a jet-blade gas target with the 1 Hz, high-intensity drive laser pulses. The supersonic gas target has a 0.8 mm opening diameter and features an adjustable rigid blade to induce a sharp downramp transition in the longitudinal density profile. The entire jet-blade target is rigidly engineered to minimize mechanical vibrations, resulting in percent-level shot-to-shot fluctuations of the plasma density features and a density downramp location fixed to within 1σ fluctuations on the order of 10s of μm. A representative plasma density profile is presented in Ref. [21]. This density downramp feature facilitates low energy spread charge bunches through localized injection. The gas jet supplied pure helium (He) gas at a backing pressure of 200 psi (1.38 MPa).

After acceleration, the electron beams were transported to charge and spectrum diagnostics using a permanent magnet quadrupole (PMQ) triplet. Electron bunch charge is measured using a calibrated integrated current transformer (ICT) and the energy spectrum of the beam is characterized using a magnetic spectrometer.

III. LONGITUDINAL FOCAL STABILITY SYSTEM

In the context of the HTU beamline, fluctuations in the drive laser's longitudinal focal position with mm-scale total variability are observed. Short-term jitter (above 0.01–0.1 Hz), shown in Figs. 2(a) and 2(b), observed via the unamplified kHz pulse train is dominated by low-frequency contributions below 1 Hz, but with one standard deviation fluctuations exceeding ±0.5 mm and maximum deviations reaching ±2.0 mm.[1] Long-term drift (below 0.01–0.1 Hz), shown in Fig. 3, is observed to reach 5.6 ± 0.9 mm over a time span of 10.3 h. The scale of these deviations is significant as the Rayleigh range of our focus is ∼4.5 mm. Short-term jitter hinders the reliability of established LPA techniques such as localized injection via density downramp injection, see Fig. 4, and presents additional challenges for advanced concepts requiring precise temporal and spatial synchronization of multiple

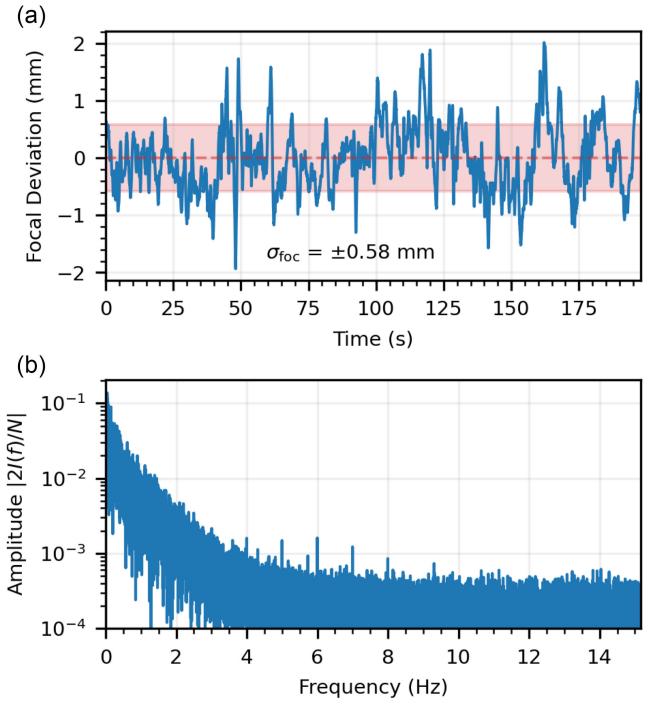


FIG. 2. Longitudinal focal position fluctuations without stabilization. Fluctuations are measured at 32 Hz using the copropagating kHz pulse train. (a) Longitudinal focal position (blue, represented as deviations from the mean position) overtime with 1σ fluctuations (red, shaded) reaching ±0.58 mm and maximum deviations reaching ±2 mm. (b) Fourier transform spectrum of the presented time series (extended over ∼15 min). Fluctuations are dominated by low-frequency contributions below 1 Hz.

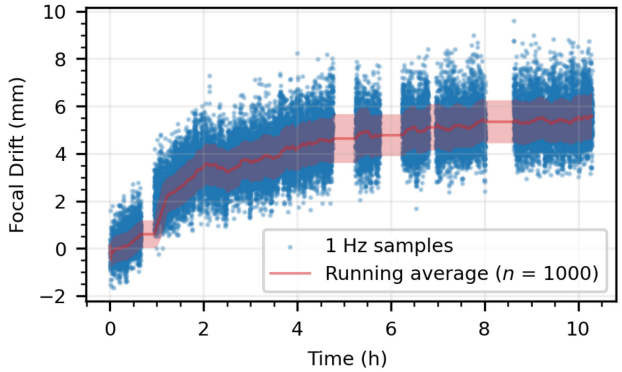


FIG. 3. Long-term drift of the longitudinal focal position of the drive beam. Focal positions (blue) were acquired near continuously over a ∼10 h period at 1 Hz. The running average (red, solid line) and standard deviation (red, shaded) for 1000 trailing samples quantified the total drift to be 5.6 ± 0.9 mm.

[1]This low-frequency dominance was further verified using a WFS capable of 88 Hz sampling. No notable spectral contributions were observed up to the 44 Hz Nyquist limit suggesting only the low-frequency contributions shown in Fig. 2 are present. This data is omitted as it was the result of a preliminary motivational study using a WFS and experimental configuration not relevant within the scope of this work.

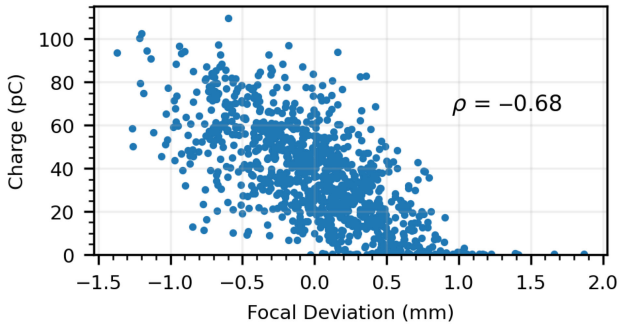


FIG. 4. Strong correlation observed between the longitudinal focal position (represented as the deviation from the mean) of the drive beam and LPA electron beam charge using density downramp injection.

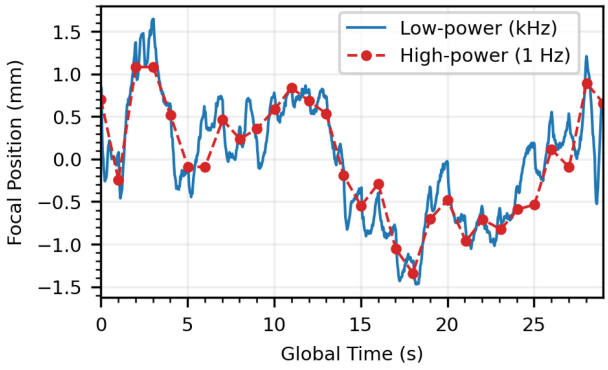


FIG. 5. Correlation between the amplified 1 Hz pulses (red, dots) and the copropagating unamplified kHz pulses (blue, solid; 32 Hz sample rate) acquired simultaneously via the two ghost beamline WFSs. In both cases, the pulses pass through the Ti:Sa crystals, which are pumped at 1 Hz. The amplified pulses correspond to the arrival time of the pump laser—initially unamplified, they become amplified upon interaction with the pumped crystal. In contrast, the unamplified kHz pulses pass through the crystals at times offset by integer milliseconds before and after the pump laser and thus remain unamplified. The corresponding correlation coefficient is 0.68. This strong correlation supports the feasibility of using the unamplified kHz pulse train as feedback to actively correct the longitudinal focal position of the amplified 1 Hz pulses.

laser pulses [7,38–41]. Similarly, long-term drift stunts the collection of large-scale statistics and general experimental operation such that the only remedy is frequent and redundant optimization of the LPA source. Fortunately, the dominance of low-frequency fluctuations in the HTU beamline, where 97% of the standard deviation arises from fluctuations below 5 Hz and 82% from below 1 Hz, suggests that active stabilization of the longitudinal focal position is feasible with only a moderate correction bandwidth. However, this still requires exceeding the Nyquist frequency limitations imposed by the 1 Hz drive laser.

Active stabilization of the 1 Hz high-power longitudinal focal position is performed by nondestructively monitoring the wave front radius of curvature (RoC) of a copropagating, unamplified kHz pulse train and performing corrective adjustments to the beam collimation into the OAP via an actuated, transmissive telescope.

A wedged fused silica mirror is used before the OAP with a 99.0% reflective coating on the front surface and a $> 99.5\%$ reflective coating on the back surface; see Fig. 1. The unattenuated front surface reflection is delivered to the target and drives the accelerator while the back surface reflection establishes a highly attenuated ($\approx 10^{-4}$), spatially separated, and correlated copy of the focus (denoted the ghost beam) for noninvasive monitoring [28,29]. A thorough investigation of the correlations between the drive beam and ghost beam is presented in Ref. [28]. The wedged mirror geometry imposes a 1.21° angular deviation of the ghost beam, allowing it to be picked off and routed outside of the vacuum chamber to a ghost beam diagnostic setup. Note that while the drive and ghost beams have near-perfect spatial overlap on the OAP, the angular offset of the ghost beam incidence introduces slight astigmatism into the ghost beam. However, it is expected that this astigmatism has negligible impact on the stabilization system performance.

Upon exit of the chamber, a beam splitter delivers the ghost beam to two diagnostic arms: (i) a high-power arm featuring a wave front sensor (WFS) (Thorlabs WFS20-7AR) triggered to capture the ghost copy of the 1 Hz amplified pulse and (ii) a low-power arm featuring a fast mechanical shutter to block the amplified pulse and second WFS (Thorlabs WFS20-7AR) for monitoring of the kHz unamplified pulses. Field lenses are placed at the beam focus on both arms to improve signal quality on the WFSs. The high-power WFS is used to measure the wave front RoC of the amplified 1 Hz pulse and identify correlations between the on-shot wave front properties and LPA electron beam qualities. The low-power WFS monitors the kHz pulse train with a 32 Hz sample rate to provide an updated measure of the wave front RoC as active feedback for the upstream corrective telescope. The longitudinal focal position of the drive beam is inferred from the ghost beam wave front RoC via the calibration in Appendix A. The feasibility of correcting the longitudinal focal position of 1 Hz amplified pulses via sampling of the unamplified kHz pulses was verified through measurement of the high degree of correlation between the two; see Fig. 5. For equipment safety, a mechanical shutter is timed to shield the low-power WFS for a 20 ms temporal window beginning ~3 ms before the 1 Hz amplified pulse. For robustness of the stabilization system, the low-power WFS is gated such that it does not acquire when the mechanical shutter is closed in order to avoid acquisition of empty frames and corresponding nonsensical corrections. Note that the mechanical shutter operation and logic gating of the WFS acquisition is consistent in the case of both low- and high-power laser operation.

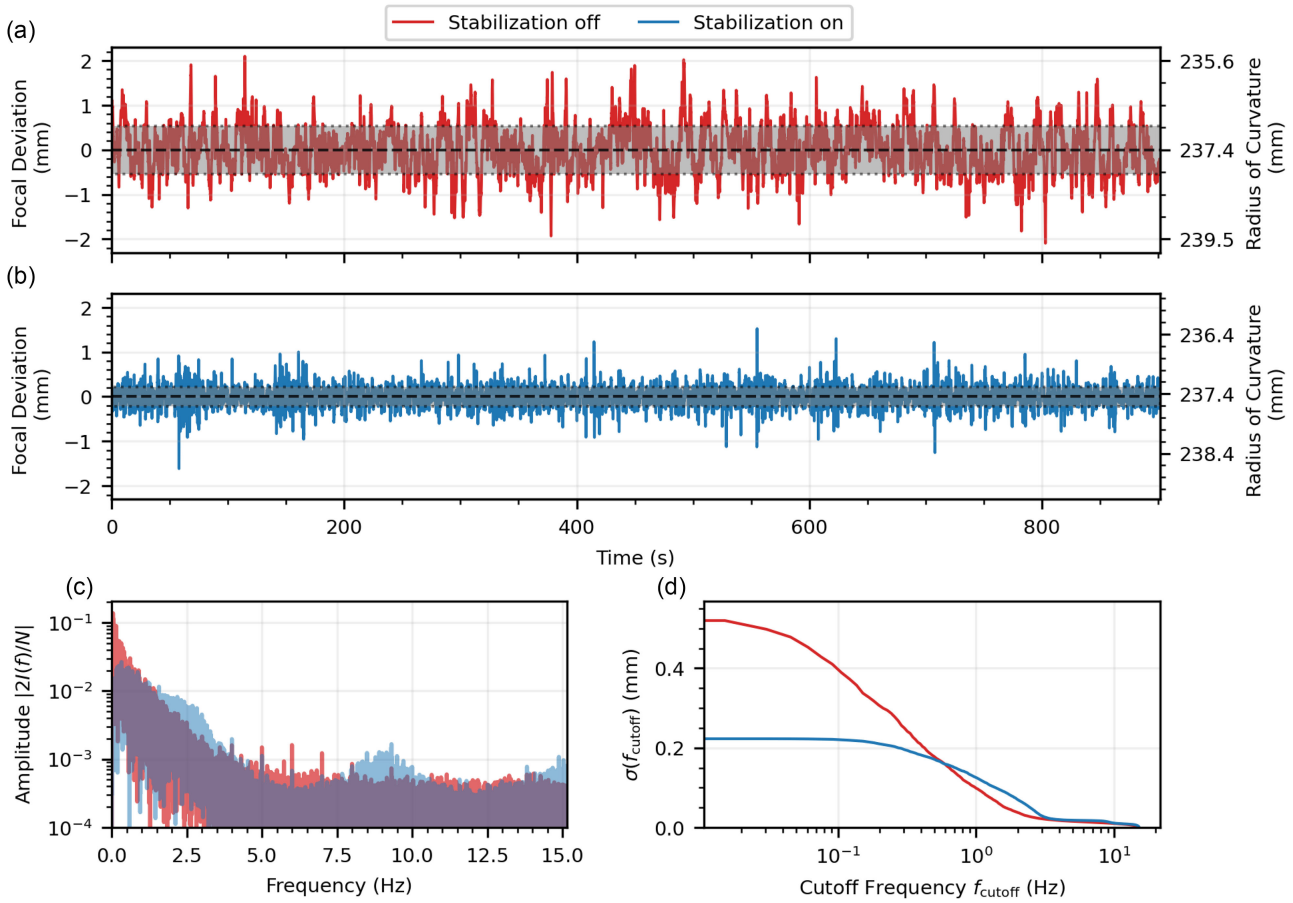


FIG. 6. Performance of the longitudinal focal stabilization system during low-power operation. Time series of the focal position (represented as the deviation from the mean) measured using the WFS on the low-power arm of the ghost beamline with (a) stabilization off (red) and (b) stabilization on (blue). The standard deviation of focal deviations is represented by the light gray band. (c) The fast Fourier transform spectrum of the presented time series. The stabilization system significantly suppresses low-frequency contributions $\lesssim$1 Hz but introduces small contributions around ~2.5, ~9, and ~15 Hz. (d) The integrated standard deviation of focal fluctuation $\sigma(f_{\mathrm{cutoff}})$ for frequencies greater than a cutoff frequency f_{cutoff}.

Calculated RoCs for the unamplified kHz pulse train are streamed to the upstream corrective telescope computer and corrective adjustments to the beam collimation are determined. As stated in Sec. II, the corrective telescope is positioned between the second and third amplification stages. This location was chosen as it is the final location of transmissive optics in the beamline due to high energy fluence of the amplified pulses after the final amplifier. The telescope consists of a concave lens with a 30.9 cm focal length and subsequent convex lens with a 50 cm focal length in an arrangement commonly referred to as a Galilean telescope. The convex lens is mounted on an actuated translation stage (Zaber X-LDA025A-AE53D12) with mm-scale range of motion and 50 mm s^{-1} maximum translation velocity. As the separation distance between the corrective telescope and the final amplification stage is short, corrections to the beam collimation have negligible impact on the laser spot size on the Ti:Sa crystal and subsequent energy extraction. Precision alignment over the full range of stage motion paired with active transverse stabilization of the laser through the amplification stages ensures minimal tip/tilt introduced during correction. This is verified by negligible transverse focal shift of the drive beam observed during translation through the full range of motion. Corrective adjustments to the telescope separation are assigned by a proportional-integral-derivative (PID) controller with parameters optimized via the Ziegler-Nichols method [42] to minimize overshoot. An independent diagnostic (the mode imager; see Fig. 1) was used to characterize the high-power focus and cross correlate the measured ghost beam RoC to the drive beam focal position in the target chamber. Further discussion regarding technical details of the stabilization system is provided in Appendix B.

IV. STABILIZATION PERFORMANCE

Testing of the longitudinal focal stabilization (LFS) system was performed in two operational modes: (i) with the main amplifiers disabled (low-power) and (ii) normal fully amplified operation (high-power). In (i) the 1 Hz drive

beam has the same energy as the kHz pulse (∼1 mJ) and in (ii) the 1 Hz drive beam has 2.5 J on target.

A. During low-power operation

Time series of measured focal positions for the unamplified kHz pulse train over a ∼15 min sample with the stabilization system off and on is presented in Figs. 6(a) and 6(b), respectively. Without stabilization, natural longitudinal focal position fluctuations exhibit a standard deviation of ±0.54 mm with maximum excursions exceeding ±2.10 mm. Active stabilization reduces the standard deviation by 29% to ±0.22 mm and reduces maximum excursion values to ±1.61 mm.

Figure 6(c) presents the normalized fast Fourier transform (FFT) spectral amplitude $|2I(f)/N|$ of focal deviations from the mean, where $I(f)$ is the FFT spectral amplitude and N is the number of samples. We note that fluctuations below 1 Hz are greatly diminished, indicating that slow variations in the focal position are being corrected by the stabilization system. A slight increase in frequency contributions is observed between 1.75 and 4 Hz, at ∼9 Hz, and at ∼15 Hz. This is further reflected in Fig. 6(d), which shows the cumulative contributions to the standard deviation as a function of a lower limit cutoff frequency f_{cutoff} given by

$$\sigma(f_{\mathrm{cutoff}}) = \sqrt{\frac{2\Delta t^2}{T}\int_{f_{\mathrm{cutoff}}}^{\infty} |I(f)|^2 \, \mathrm{d}f}, \tag{1}$$

where Δt is the time between samples and $T = N\Delta t$ is the total acquisition time. Note that the upper limit of integration in Eq. (1) is limited to the Nyquist frequency. At this time, the source of these contributions is not yet identified but is reasonably suspected to be related to movement of the corrective telescope stage. The corrective bandwidth is

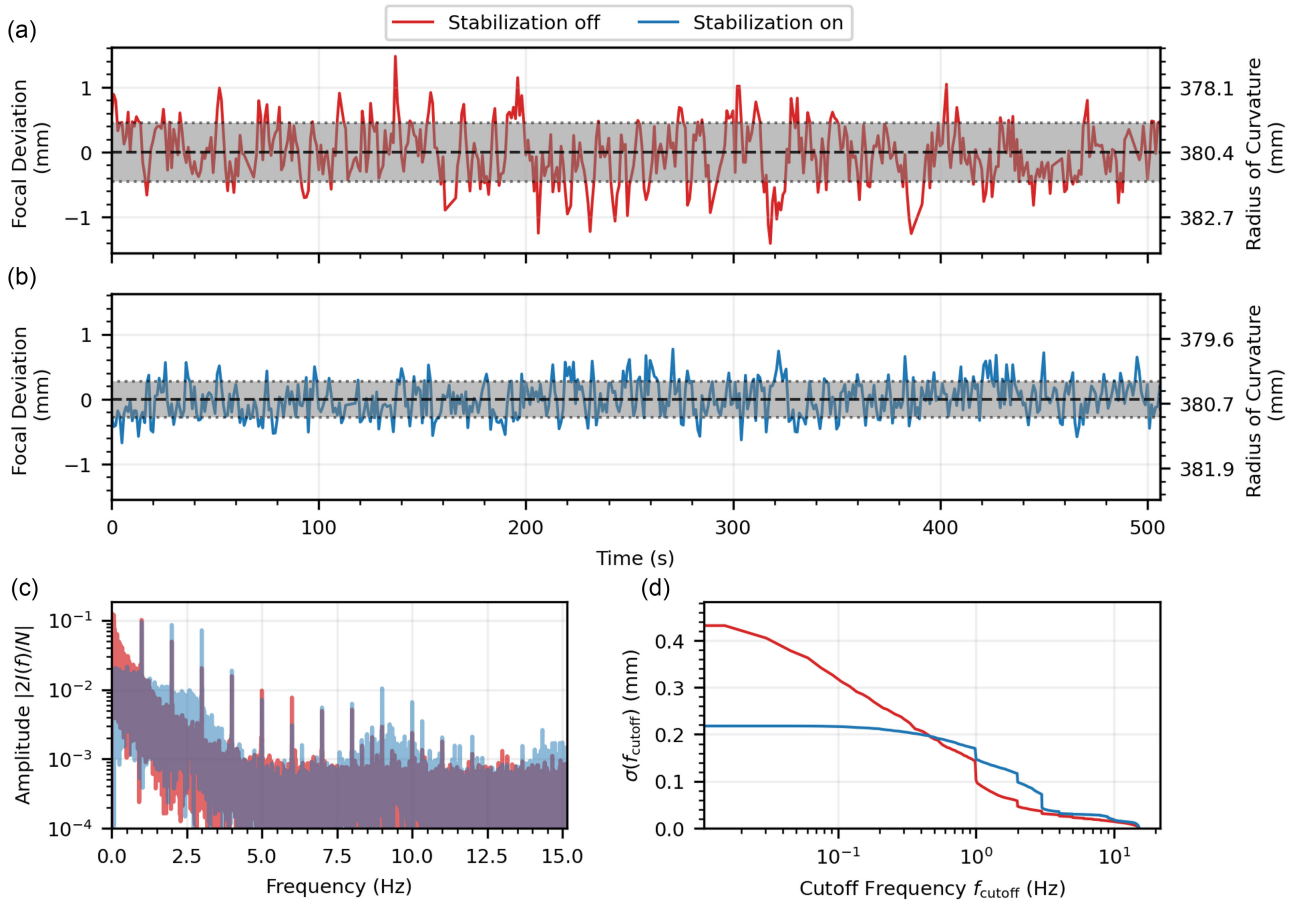


FIG. 7. Performance of the longitudinal focal stabilization system during high-power operation. Time series of the focal position (represented as the deviation from the mean) measured using the WFS on the high-power arm of the ghost beamline with (a) stabilization off (red) and (b) stabilization on (blue). The standard deviation of focal deviations is represented by the light gray band. (c) The fast Fourier transform spectrum of the simultaneously acquired focal position time series obtained via the high sample rate WFS on the low-power arm of the ghost beamline. The stabilization system suppresses low-frequency contributions $\lesssim$1 Hz but introduces small contributions around ∼2.5, ∼9, and ∼15 Hz. Spikes at integer frequencies indicate a strong 1 Hz feature related to high-power operation but independent of stabilization. (d) The integrated standard deviation of focal fluctuation $\sigma(f_{\mathrm{cutoff}})$ for frequencies greater than a cutoff frequency f_{cutoff}.

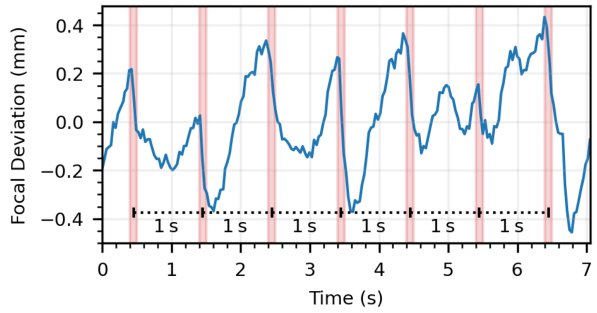


FIG. 8. Suspected thermal lensing effects of amplification. A time series of the focal position for the unamplified kHz pulses during high-power operation (blue) shows a repetitive sudden upstream retraction (negative focal deviation, highlighted by red shaded regions) and subsequent downstream push (positive focal deviation) of the focal position at 1 Hz frequency. A sudden retraction of the focal position may be caused by thermal lensing in the Ti:Sa amplifier crystals, while the slow, downstream push of the focal position may be due to relaxation of this effect between 1 Hz pump cycles.

estimated to be 0.59 Hz by noting the intersection of the $\sigma(f_{\text{cutoff}})$ curves.

B. During high-power operation

Time series of measured focal positions for the amplified 1 Hz ghost pulses over a 500 s sample window with the stabilization system off and on are presented in Figs. 7(a) and 7(b), respectively. Without stabilization, natural longitudinal focal position fluctuations exhibit a standard deviation of ± 0.46 mm with maximum excursions exceeding ± 1.48 mm. Active stabilization reduces the standard deviation fluctuations by 53% to ± 0.22 mm and limits maximum excursions to ± 0.22 mm. The stabilized performance is nearly identical to that observed during low-power operation, despite a slightly lower percent reduction in standard deviation fluctuations. Rather, this difference is attributed to smaller unstabilized fluctuations during high-power operation, likely a consequence of the lower repetition rate and smaller sample size.

While measurements of the amplified laser are restricted to 1 Hz, we can also monitor the unamplified laser to further characterize the spectral response. Figure 7(c) shows the normalized FFT spectral amplitude for focal deviations of the unamplified kHz pulse train while amplifying at 1 Hz. Figure 7(d) shows the cumulative standard deviation fluctuations due to frequencies ranging from a cutoff frequency to the Nyquist frequency limit. As in the case of low-power operation, fluctuations below 1 Hz are largely suppressed while a slight increase in spectral contributions between 1.75 and 4 Hz, at ∼9 Hz, and at ∼15 Hz is observed. Indicated by the intersection of $\sigma(f_{\text{cutoff}})$ curves, a slight decrease in correction bandwidth to 0.46 Hz is observed during high-power operation. Nonetheless, the correction bandwidth achieved is significantly better than what could be achieved using the 1 Hz pulse train alone.

In contrast to the unamplified case, the FFT of the kHz pulse train for the amplified system exhibits large spikes at integer frequencies; see Fig. 7(c). The time series of focal deviations during amplification in Fig. 8 shows a pattern of sudden upstream focal retractions (negative deviations) followed by slow focal drifts back downstream (positive deviations) with 1 Hz repetition. This behavior is not suspected to be related to the stabilization system as it occurs whether stabilization is on or off. Rather, the process of amplifying the laser at 1 Hz introduces changes to the thermal lensing in the amplifier crystals that strongly affect the low-power pulses arriving immediately after the amplified pulse. This causes discrete jumps in the RoC after the amplified pulse that manifest as 1 Hz (and subsequent harmonic) peaks in the frequency spectra. Notably, these rapid periodic spikes do not impact the stability of the amplified laser. However, they do contribute to the cumulative standard deviation calculations and correction bandwidth estimation.

C. Impact on LPA performance

LPA operation during active stabilization shows marked improvement in the stability of LPA electron beam qualities in comparison to operation without stabilization. Electron

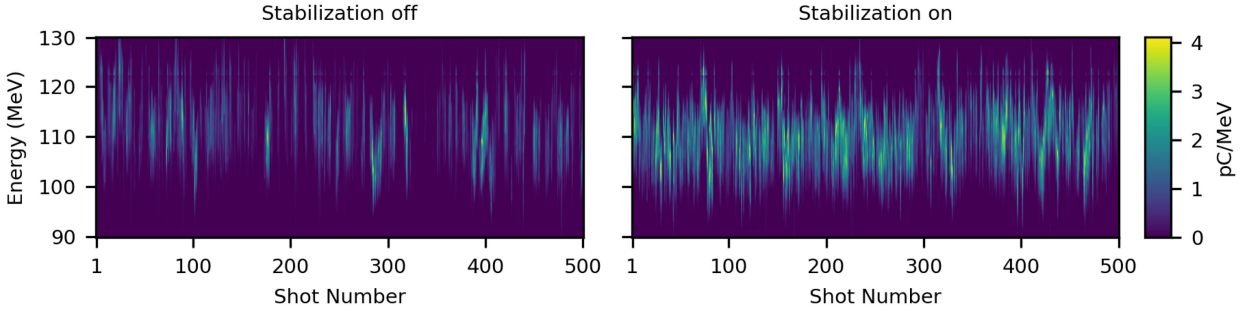


FIG. 9. Comparison of LPA electron beams over 500 consecutive shots with stabilization off (left) and on (right). Stabilization reduces fluctuations in shot-to-shot beam charge and central beam energy. Note that a faint dark line can be observed at ∼120 MeV. This is an artifact of a transition point in our diagnostic imaging setup and can be ignored.

TABLE I. Quantitative effects of stabilization on LPA performance.

	Stabilization off		Stabilization on	
Quantity	All shots	Shots ≥ 5 pC	All shots	Shots ≥ 5 pC
Missed shots (< 5 pC)	145 (29%)		32 (6%)	
Charge (pC)	19.4 ± 17.9	26.8 ± 16.2	40.4 ± 20.4	43.1 ± 18.3
Average energy (MeV)	120.3 ± 18.9	111.8 ± 4.3	111.1 ± 9.1	109.7 ± 3.8
Energy spread (%, rms)	5.3 ± 3.0	4.1 ± 0.7	4.3 ± 1.9	3.9 ± 0.4

energy spectra over 500 consecutive shots with and without stabilization are shown in Fig. 9 and corresponding electron beam properties are displayed in Table I. The total bunch charge observed without stabilization is 19.4 ± 17.9 pC but improves to 40.4 ± 20.4 pC with stabilization. Not only do we observe a factor of ~2 increase in the average bunch charge, likely due to simply remaining closer to the optimal longitudinal focal position over more shots, but the relative 1σ fluctuation is reduced from ~92% to ~50% of the mean. Additionally, the percent root-mean-squared (rms) energy spread of the electron bunches was $5.3\% \pm 3.0\%$ without stabilization and $4.3\% \pm 1.9\%$ with stabilization; demonstrating a 19% reduction in mean energy spread and 37% reduction in standard deviation. Finally, the average bunch energy is measured to be 120.3 ± 18.9 MeV without stabilization and 111.1 ± 9.1 MeV with stabilization.

A critical figure of merit for characterizing LPA performance is the number of missed shots, where a missed shot is defined as one with a total bunch charge of less than 5 pC. Over a span of 500 consecutive shots, 145 missed shots (29%) were observed without stabilization while only 32 missed shots (6%) were observed with stabilization. While it may be unsurprising, this characterization is insightful as some of the observed electron beam statistics with and without stabilization converge when the missed shots are removed; see Table I. For example, without stabilization, the average beam energy including all shots is measured at 120.3 ± 18.9 MeV, but this improves to 111.8 ± 4.3 MeV when missed shots are excluded. This is now comparable to the stabilized case with missed shots removed, where the average energy is 109.7 ± 3.8 MeV. Still, performance with stabilization maintains an advantage across all quantities even when removing missed shots. Convergence of beam statistics is most prominent for quantities related to bunch spectral content while the discrepancy in average bunch charge and relative charge fluctuations remains significant. This suggests that by removing missed shots, we are likely rejecting shots with near-maximal focal deviation resulting in significantly reduced electron injection and retaining those with more minor focal deviations.

It is important to note that the transverse stabilization system (Ref. [29]) was operated in a slow-drift correction mode applying corrections to the transverse focal position at ~0.1 Hz. Thus, while the transverse focal position was not stabilized shot-to-shot, it was maintained between datasets ensuring valid comparison. However, given the mm-scale gas target size and findings from Ref. [29] showing that transverse laser fluctuations primarily affect electron beam trajectory with minimal impact on charge and spectral quality, the shot-to-shot fluctuations are not expected to have significantly influenced the measured electron beam properties.

V. CONCLUSION

Active stabilization of the longitudinal focal position for a 100 TW-class laser system was demonstrated and resulted in marked improvement in LPA shot-to-shot stability and corresponding electron beam figures of merit, particularly in terms of reducing missed shots, increasing total beam charge, and reducing fluctuations in average beam energy and energy spread. Furthermore, active stabilization of the longitudinal focal position enables stable LPA operation over many hours without the need for reoptimization.

Further investigation to characterize and optimize the stabilization and online monitoring system is expected to enable enhanced performance. In its current state, the system suppresses dominant low-frequency (<1 Hz) fluctuations of the focal position but introduces low amplitude higher-frequency contributions around ~3, ~9, and ~15 Hz. The source of these fluctuations is currently unidentified but is suspected to be related to either mechanical vibrations of the corrective telescope motion or misleading aliasing in the FFT. Additionally, sampling of the wave front RoC enabled the observation of potential 1 Hz thermal lensing effects of the amplifier stages during high-power operation. Further investigation is required to identify the significance of these effects on the stabilization of the 1 Hz high-power focal position.

Multiple paths forward can be taken toward improving this stabilization system. Preliminary PID-loop simulations suggest significant enhancements to the overall system performance can be made by increasing the correction rate. Immediate improvements to the system's correction bandwidth may be possible by utilizing alternative acquisition modes native to the deployed WFS, enabling WFS sampling frequency to be increased from 32 Hz to a maximum

of 880 Hz. However, faster operations require reducing spatial resolution of the WFS, improving throughput but increasing error in the absolute setpoint and readback during operations. Additional mechanical improvements are required to enable the stage to operate with such high-speed corrections. Combining the current stage with a faster corrector, such as a piezo system, should permit high-speed correction. Similarly, deployment of a WFS and translation stage with higher precision and higher fidelity may enable a higher degree of correction efficiency.

Long-term improvements to the system envision machine-learning (ML) correction models deployed on field programmable gate array (FPGA) hardware [43]. ML-based correction enables improved modeling of the measured wave front and RoC while also improving the speed and accuracy for which corrective moves are determined based on a trained model of the system. In addition, ML techniques such as convolutional neural networks and Kalman filtering have shown success in denoising and filtering experimental datasets. FPGA-based deployment circumvents data transfer and processing bottlenecks in a conventional control system to allow performance speeds exceeding 5000 updates per second. Ultimately, these modifications are necessary to scale this stabilization system, and those similar, to support the kHz repetition rates desired for future LPAs.

ACKNOWLEDGMENTS

This work was supported by the U.S. Department of Energy (DOE), Office of Science, the Office of Basic Energy Sciences, the Office of High Energy Physics, under Contracts No. DE-SC0021680 and No. DE-AC02-05CH11231.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX A: CALIBRATION OF ROC TO FOCAL POSITION

The LFS system's telescope optics were scanned over a range of separation distances while simultaneously measuring either the ghost beam RoC or the longitudinal focal position of the drive beam. Figure 10(a) shows how the ghost beam RoC varies as a function of the telescope separation. A second order polynomial was used to map this relationship. Figure 10(b) shows how the focal position of the drive beam changes as a function of the telescope configuration. For each telescope configuration, the focal position was mapped by scanning the mode imager imaging plane around the focus, fitting a Gaussian to the signal intensity, and defining the focal location as the mean of the Gaussian fit. Uncertainty in the focal position was negligible ($\lesssim 0.1\%$) and thus omitted from the plot. From these relations, we can cross calibrate the relationship between the measured ghost RoC and on-target focal position of the drive beam. Figure 10(c) shows this relationship for both the low- and high-power ghost beams.

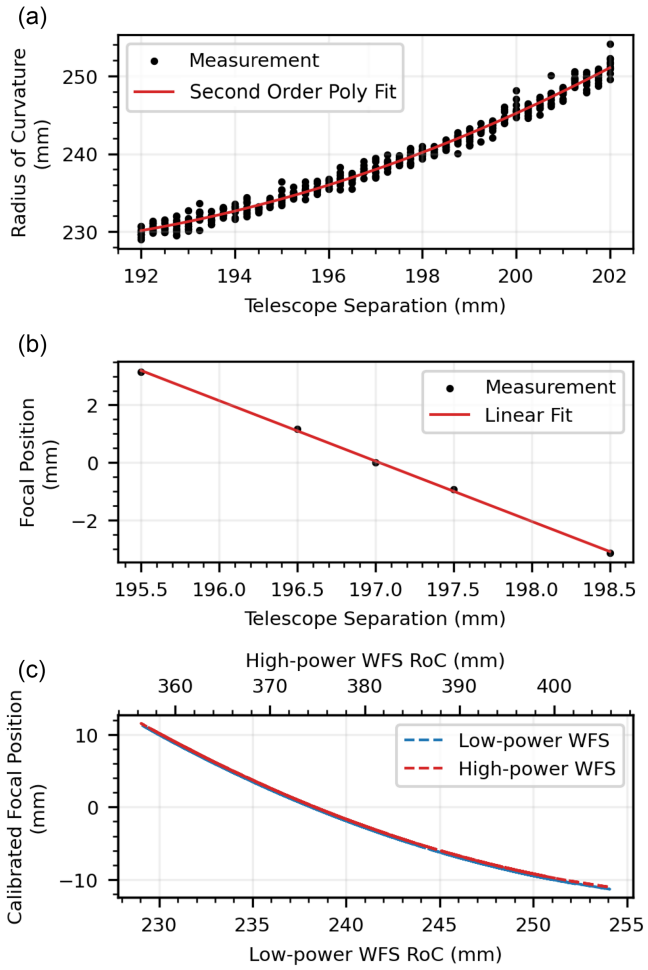


FIG. 10. Calibration of ghost beam RoC to drive beam focal position. (a) Relationship between low-power ghost beam RoC and the actuated telescope position. (b) Relationship between the main beam focal position and the actuated telescope position. (c) The calibrated focal position (with respect to the arbitrary mode imager stage position) and ghost beam RoC values.

APPENDIX B: TECHNICAL DETAILS OF THE LFS SYSTEM

The control system was developed based on a closed-loop PID controller using RoC calculations from a Thorlabs WFS20-7AR Shack-Hartmann WFS as input. The PID parameters were optimized via the Ziegler-Nichols method [42] while minimizing overshoot. The WFS supports a range of different operational modes, which balance acquisition rate and spatial resolution. For this demonstration, the WFS was operated in the so-called "normal"

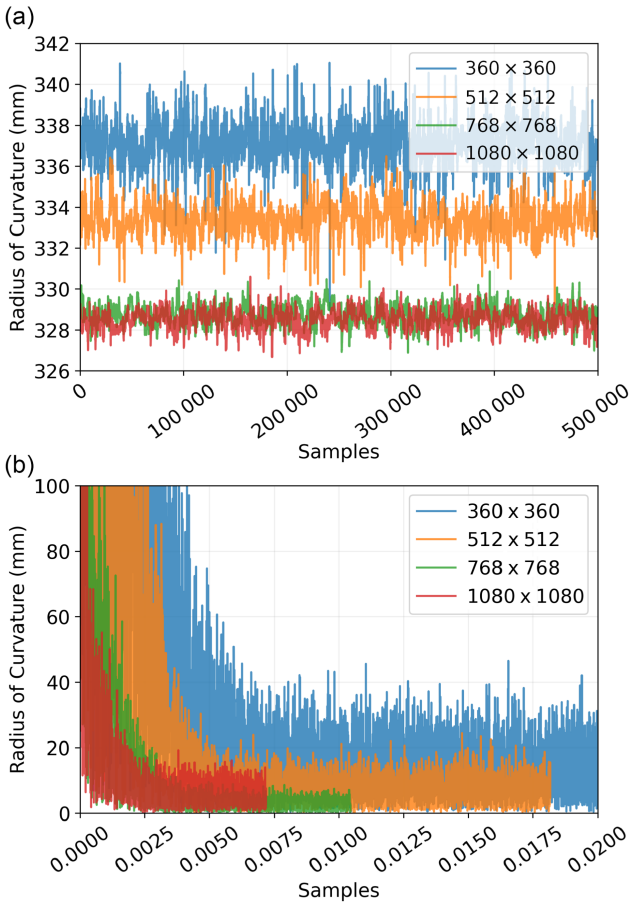


FIG. 11. RoC measurements for varying WFS resolution settings. (a) Sensor resolution vs measured RoC in time. (b) Sensor resolution vs measured RoC in frequency. Systematic reductions in measured RoC variance amplitude and frequency are observed with increasing pixel count.

acquisition mode with a sample rate of 32 Hz, 8 ms exposure time, and the highest available spatial resolution capturing the full 47 × 35 microlens spot array. When the Thorlabs WFS20 operates in "normal" mode, the entire captured image is sent to the connected computer. Alternative notable operation modes (assuming 1 ms exposure for maximum sampling rate of this camera) support reducing the spatial resolution to an 11 × 11 microlens spot array in exchange for a boost in sampling rate to 116 Hz for "normal" acquisition mode operation and up 880 Hz using a "high-speed" acquisition mode. When the Thorlabs WFS20 operates in "high-speed" mode, a built-in front-end processes the image data and sends only the calculated spot centroid coordinates. Figure 11 shows a series of RoC measurements for differing sensor resolutions. The RoC calculation gains a static error as the resolution decreases, which would be an effect of both beam position and less available data for calculating the RoC.

The motion stage for this proof-of-principle telescope design is a Zaber X-LDA025A-AE53D12 moving magnet linear stage. Relatively fast default acceleration and

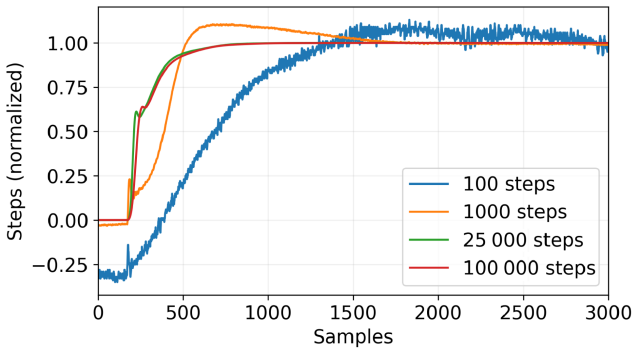


FIG. 12. Stage response based on requested number of steps to move in normalized values. Each sample corresponds to 0.1 ms, with each step representing 1 nm.

velocity of the motion stage resulted in higher noise levels in focal position during preliminary characterization, likely due to properties of the load seen by the stage. Using a manufacturer-provided toolkit, the system was tuned to achieve a decrease in operational limits and a subsequent smoothing of the motion profile. These changes, combined with the utilization of available motion averaging features, provided for the best RoC stability.

Several factors determine our correction bandwidth using this particular stage, including serial communication speed, the stage's internal control loop bandwidth, and settling time. The stage response to requested moves of various step sizes can be seen in Fig. 12, i.e., the stage moves toward the target position much quicker when the requested move requires a larger number of steps. As the stage's positioning control loop is an internal PID (separate from the stabilization system PID control loop), the control loop parameters can be tuned to adjust the settling time and overshoot. Using the default, tuned parameters [44], it takes about 15–20 ms for the stage to get relatively close to the requested position. To circumvent bandwidth limitations arising from the stage's physical movement and communication protocol, a "last-message-first" protocol is implemented (discussed further below). Essentially, if a new move request is made while a move is in progress, the ongoing move is interrupted and the target set point is immediately updated.

Acquisition and control settings are controlled via a custom Python library, which leverages the Thorlabs driver Application Programming Interface (API) to manage acquisition, processing, and control of the WFS. The same library employs a custom API for the Zaber stage, using the device's native ASCII protocol to enable asynchronous control during operations. Discrete components communicate with each other using a custom User Datagram Protocol (UDP) allowing for diagnostic data to be sent from the sensor to multiple endpoints for logging, online monitoring, and stage control.

Data transport and receive overhead was minimal due to the use of custom UDP messages with no acknowledgment

(ACK). This eliminated concerns of not meeting timing due to a lost link, while also allowing for a "last-message-first" processing priority by dropping other messages. Unlike many control and monitoring applications, the priority is processing data as fast as it arrives with no need for guaranteed delivery. This also allowed for the easy addition of signal monitoring and logging on other networked workstations besides the stage and loop controller hardware.

An important benefit of this system design is that stabilization is dependent only on relative corrections to the wave front RoC. Thus, the target RoC is defined purely as a setpoint for the control loop and the critical factor for robust stabilization is the motion profile of the stage over the correction distance. Not only does this allow for flexibility in fixing the focal position with respect to an experimental configuration, but also increases flexibility of WFS acquisition as accounting for static offset in RoC measurements becomes trivial.

[1] T. Tajima and J. M. Dawson, Laser electron accelerator, Phys. Rev. Lett. **43**, 267 (1979).

[2] E. Esarey, C. B. Schroeder, and W. P. Leemans, Physics of laser-driven plasma-based electron accelerators, Rev. Mod. Phys. **81**, 1229 (2009).

[3] B. Miao, J. E. Shrock, L. Feder, R. C. Hollinger, J. Morrison, R. Nedbailo, A. Picksley, H. Song, S. Wang, J. J. Rocca, and H. M. Milchberg, Multi-GeV electron bunches from an all-optical laser wakefield accelerator, Phys. Rev. X **12**, 031038 (2022).

[4] C. Aniculaesei *et al.*, The acceleration of a high-charge electron bunch to 10 GeV in a 10-cm nanoparticle-assisted wakefield accelerator, Matter Radiat. Extremes **9**, 014001 (2023).

[5] A. Picksley, J. Chappell, E. Archer, N. Bourgeois, J. Cowley, D. Emerson, L. Feder, X. Gu, O. Jakobsson, A. Ross, W. Wang, R. Walczak, and S. Hooker, All-optical GeV electron bunch generation in a laser-plasma accelerator via truncated-channel injection, Phys. Rev. Lett. **131**, 245001 (2023).

[6] A. Picksley, J. Stackhouse, C. Benedetti, K. Nakamura, H. E. Tsai, R. Li, B. Miao, J. E. Shrock, E. Rockafellow, H. M. Milchberg, C. B. Schroeder, J. van Tilborg, E. Esarey, C. G. R. Geddes, and A. J. Gonsalves, Matched guiding and controlled injection in dark-current-free, 10-GeV-class, channel-guided laser-plasma accelerators, Phys. Rev. Lett. **133**, 255001 (2024).

[7] J. Faure, C. Rechatin, A. Norlin, A. Lifschitz, Y. Glinec, and V. Malka, Controlled injection and acceleration of electrons in plasma wakefields by colliding laser pulses, Nature (London) **444**, 737 (2006).

[8] H.-E. Tsai, K. K. Swanson, S. K. Barber, R. Lehe, H.-S. Mao, D. E. Mittelberger, S. Steinke, K. Nakamura, J. van Tilborg, C. Schroeder, E. Esarey, C. G. R. Geddes, and W. Leemans, Control of quasi-monoenergetic electron beams from laser-plasma accelerators with adjustable shock density profile, Phys. Plasmas **25**, 043107 (2018).

[9] W. T. Wang, W. T. Li, J. S. Liu, Z. J. Zhang, R. Qi, C. H. Yu, J. Q. Liu, M. Fang, Z. Y. Qin, C. Wang, Y. Xu, F. X. Wu, Y. X. Leng, R. X. Li, and Z. Z. Xu, High-brightness high-energy electron beams from a laser wakefield accelerator via energy chirp control, Phys. Rev. Lett. **117**, 124801 (2016).

[10] X. Xu, F. Li, F. S. Tsung, K. Miller, V. Yakimenko, M. J. Hogan, C. Joshi, and W. B. Mori, Generation of ultrahigh-brightness pre-bunched beams from a plasma cathode for x-ray free-electron lasers, Nat. Commun. **13**, 3364 (2022).

[11] F. Foerster *et al.*, Stable and high-quality electron beams from staged laser and plasma wakefield accelerators, Phys. Rev. X **12**, 041016 (2022).

[12] O. Lundh, J. Lim, C. Rechatin, L. Ammoura, A. Ben-Ismaïl, X. Davoine, G. Gallot, J. P. Goddet, E. Lefebvre, V. Malka, and J. Faure, Few femtosecond, few kiloampere electron bunch produced by a laser-plasma accelerator, Nat. Phys. **7**, 219 (2011).

[13] F. Albert and A. G. Thomas, Applications of laser wakefield accelerator-based light sources, Plasma Phys. Controlled Fusion **58**, 103001 (2016).

[14] B. Hidding, O. Karger, T. Königstein, G. Pretzler, G. G. Manahan, P. McKenna, R. Gray, R. Wilson, S. M. Wiggins, G. H. Welsh, A. Beaton, P. Delinikolas, D. A. Jaroszynski, J. B. Rosenzweig, A. Karmakar, V. Ferlet-Cavrois, A. Costantino, M. Muschitiello, and E. Daly, Laser-plasma-based space radiation reproduction in the laboratory, Sci. Rep. **7**, 42354 (2017).

[15] C. Schroeder *et al.*, Linear colliders based on laser-plasma accelerators, J. Instrum. **18**, T06001 (2023).

[16] A. Borghini, L. Labate, S. Piccinini, C. M. V. Panaino, M. G. Andreassi, and L. A. Gizzi, FLASH radiotherapy: Expectations, challenges, and current knowledge, Int. J. Mol. Sci. **25**, 2546 (2024).

[17] K. Nakajima, Towards a table-top free-electron laser, Nat. Phys. **4**, 92 (2008).

[18] A. R. Maier, A. Meseck, S. Reiche, C. B. Schroeder, T. Seggebrock, and F. Grüner, Demonstration scheme for a laser-plasma-driven free-electron laser, Phys. Rev. X **2**, 031019 (2012).

[19] W. Wang, K. Feng, L. Ke, C. Yu, Y. Xu, R. Qi, Y. Chen, Z. Qin, Z. Zhang, M. Fang, J. Liu, K. Jiang, H. Wang, C. Wang, X. Yang, F. Wu, Y. Leng, J. Liu, R. Li, and Z. Xu, Free-electron lasing at 27 nanometres based on a laser wakefield accelerator, Nature (London) **595**, 516 (2021).

[20] M. Labat *et al.*, Seeded free-electron laser driven by a compact laser plasma accelerator, Nat. Photonics **17**, 150 (2023).

[21] S. K. Barber, F. Kohrell, C. E. Doss, K. Jensen, C. Berger, F. Isono, Z. Eisentraut, S. Schröder, A. J. Gonsalves, K. Nakamura, G. R. Plateau, R. A. Van Mourik, M. Gracia-Linares, L. Labun, B. M. Hegelich, S. V. Milton, C. G. R. Geddes, J. Osterhoff, C. B. Schroeder, E. H. Esarey, and J. Van Tilborg, Greater than 1000-fold gain in a free-electron laser driven by a laser-plasma accelerator with high reliability, Phys. Rev. Lett. **135**, 055001 (2025).

[22] C. Emma, J. Van Tilborg, R. Assmann, S. Barber, A. Cianchi, S. Corde, M. E. Couprie, R. D'Arcy, M. Ferrario, A. F. Habib *et al.*, Free electron lasers driven by plasma

accelerators: Status and near-term prospects, High Power Laser Sci. Eng. **9**, e57 (2021).

[23] M. Galletti, R. Assmann, M. E. Couprie, M. Ferrario, L. Giannessi, A. Irman, R. Pompili, and W. Wang, Prospects for free-electron lasers powered by plasma-wakefield-accelerated beams, Nat. Photonics **18**, 780 (2024).

[24] A. R. Maier, N. M. Delbos, T. Eichner, L. Hübner, S. Jalas, L. Jeppe, S. W. Jolly, M. Kirchen, V. Leroux, P. Messner, M. Schnepp, M. Trunk, P. A. Walker, C. Werle, and P. Winkler, Decoding sources of energy variability in a laser-plasma accelerator, Phys. Rev. X **10**, 031039 (2020).

[25] G. Genoud, F. Wojda, M. Burza, A. Persson, and C.-G. Wahlström, Active control of the pointing of a multi-terawatt laser, Rev. Sci. Instrum. **82**, 033102 (2011).

[26] L. Gizzi, P. Koester, L. Labate, F. Mathieu, Z. Mazzotta, G. Toci, and M. Vannini, Lasers for novel accelerators, J. Phys. Conf. Ser. **1350**, 012157 (2019).

[27] L. Kiani *et al.*, High average power ultrafast laser technologies for driving future advanced accelerators, J. Instrum. **18**, T08006

[28] F. Isono, J. van Tilborg, S. K. Barber, J. Natal, C. Berger, H.-E. Tsai, T. Ostermayr, A. Gonsalves, C. Geddes, and E. Esarey, High-power non-perturbative laser delivery diagnostics at the final focus of 100-TW-class laser pulses, High Power Laser Sci. Eng. **9**, e25 (2021).

[29] C. Berger, S. Barber, F. Isono, K. Jensen, J. Natal, A. Gonsalves, and J. van Tilborg, Active nonperturbative stabilization of the laser-plasma-accelerated electron beam source, Phys. Rev. Accel. Beams **26**, 032801 (2023).

[30] D. Strickland and G. Mourou, Compression of amplified chirped optical pulses, Opt. Commun. **56**, 219 (1985).

[31] S. Bulanov, N. Naumova, F. Pegoraro, and J. Sakai, Particle injection into the wave acceleration phase due to nonlinear wake wave breaking, Phys. Rev. E **58**, R5257 (1998).

[32] H. Suk, N. Barov, J. B. Rosenzweig, and E. Esarey, Plasma electron trapping and acceleration in a plasma wake field using a density transition, Phys. Rev. Lett. **86**, 1011 (2001).

[33] C. G. R. Geddes, K. Nakamura, G. R. Plateau, Cs. Toth, E. Cormier-Michel, E. Esarey, C. B. Schroeder, J. R. Cary, and W. P. Leemans, Plasma-density-gradient injection of low absolute-momentum-spread electron bunches, Phys. Rev. Lett. **100**, 215004 (2008).

[34] K. Schmid, A. Buck, C. M. S. Sears, J. M. Mikhailova, R. Tautz, D. Herrmann, M. Geissler, F. Krausz, and L. Veisz, Density-transition based electron injector for laser driven wakefield accelerators, Phys. Rev. ST Accel. Beams **13**, 091301 (2010).

[35] A. J. Gonsalves, K. Nakamura, C. Lin, D. Panasenko, S. Shiraishi, T. Sokollik, C. Benedetti, C. B. Schroeder, C. G. Geddes, J. Van Tilborg, J. Osterhoff, E. Esarey, C. Toth, and W. P. Leemans, Tunable laser plasma accelerator based on longitudinal density tailoring, Nat. Phys. **7**, 862 (2011).

[36] A. Buck, J. Wenz, J. Xu, K. Khrennikov, K. Schmid, M. Heigoldt, J. M. Mikhailova, M. Geissler, B. Shen, F. Krausz, S. Karsch, and L. Veisz, Shock-front injector for high-quality laser-plasma acceleration, Phys. Rev. Lett. **110**, 185006 (2013).

[37] K. K. Swanson, H.-E. Tsai, S. K. Barber, R. Lehe, H.-S. Mao, S. Steinke, J. Van Tilborg, K. Nakamura, C. G. R. Geddes, C. B. Schroeder, and others, Control of tunable, monoenergetic laser-plasma-accelerated electron beams using a shock-induced density downramp injector, Phys. Rev. Accel. Beams **20**, 51301 (2017).

[38] L.-L. Yu, E. Esarey, C. B. Schroeder, J.-L. Vay, C. Benedetti, C. G. R. Geddes, M. Chen, and W. P. Leemans, Two-color laser-ionization injection, Phys. Rev. Lett. **112**, 125001 (2014).

[39] B. Hidding, J. B. Rosenzweig, Y. Xi, B. O'Shea, G. Andonian, D. Schiller, S. Barber, O. Williams, G. Pretzler, T. Königstein, F. Kleeschulte, M. J. Hogan, M. Litos, S. Corde, W. W. White, P. Muggli, D. L. Bruhwiler, and K. Lotov, Beyond injection: Trojan horse underdense photocathode plasma wakefield acceleration, AIP Conf. Proc. **1507**, 570 (2012).

[40] B. Hidding *et al.*, *First Measurements of Trojan Horse Injection in a Plasma Wakefield Accelerator* (CERN, Geneva, 2017), pp. 1252–1257.

[41] E. Esarey, R. F. Hubbard, W. P. Leemans, A. Ting, and P. Sprangle, Electron injection into plasma wakefields by colliding laser pulses, Phys. Rev. Lett. **79**, 2682 (1997).

[42] J. G. Ziegler and N. B. Nichols, Optimum settings for automatic controllers, Trans. ASME **64**, 759 (1942).

[43] J. Einstein-Curtis, S. J. Coleman, N. M. Cook, J. P. Edelen, S. Barber, C. Berger, and J. van Tilborg, Online correction of laser focal position using FPGA-based ML models, J. Phys. Conf. Ser. **2420**, 012074 (2023).

[44] Zaber X-LDA025A-AE53D12 default settings, https://www.zaber.com/device-settings?device=X-LDA025A-AE53D12/&peripheral=N/A&version=7.39&protocol=ASCII (2024).