SUPPLEMENTARY INFORMATION:

Drift-free, 11 fs relative pulse delay stability in dual arm PW-class laser system

Andrei B. Nazîru,[1,2] Ştefan Popa,[1,2] Ana-Maria Lupu,[1,2] Dan Gh. Matei,[1] Alice Dumitru,[1,2,3] Dmitrii Nistor,[1,3] Antonia Toma,[1] Lidia Văsescu,[1] Ioan Dăncuş,[1] Claudiu A. Stan[4] and Daniel Ursescu[1,2,*]

[1] Extreme Light Infrastructure - Nuclear Physics, National Institute for Physics and Nuclear Engineering - Horia Hulubei, 30 Reactorului str., Măgurele, 077125, Ilfov, Romania

[2] Doctoral School of Physics, University of Bucharest, Atomiştilor 405, Măgurele, 077125, Ilfov, Romania

[3] Engineering and Applications of Lasers and Accelerators Doctoral School, National University of Sciences and Technology Politehnica Bucharest, Bucharest 060042, Romania

[4] Department of Physics, Rutgers University-Newark, New Jersey 07102, USA

*daniel.ursescu@eli-np.ro

1. Generation and characterization of the light pulses

The High Power Laser System (HPLS) at ELI-NP uses independent closed-loop feedback systems for each arm to achieve the shortest duration of the delivered pulses [1]. The temporal compressors of the HPLS use a leaky mirror for beam steering that transmits less than 1% of the light towards diagnostics benches where the pulse properties are characterized. Pulse durations are measured using a self-referenced spectral interferometry (SRSI) device [2] (Fastlite, Wizzler). The information from the SRSI is fed to acousto-optic programmable dispersive filters (AOPDF) [3] (Fastlite, Dazzler) located before the first amplifiers in each arm. Then, the feedback loop adjusts the dispersion coefficients of the AOPDF to achieve the shortest possible pulse duration. With feedback control, a compressed pulse duration of ~23 fs is achieved, with at most 10% variation, for all output configurations (100 TW, 1 PW, 10 PW) [1].

The pump pulse used in the experiment had the shortest achievable duration of ~23 fs. The probe pulse was first set to this minimum duration using the feedback loop. Then, we added a group delay dispersion (GDD) of ~10000 fs^2 by tuning the second order dispersion parameter with the AOPDF, which stretched the probe pulse to 1.1 ps.

The size of the plasma mirror was estimated as the size of the burned area left on the tape after ablation. This burnt area was roughly circular with a size of approximately 150 μm, and was stable both within one run and between runs. Figure S1(a) shows consecutive spots from one run, and Figure S1(c) shows individual spots from other runs.

The size of the probe pulse was estimated by imaging the probe beam at the location of the target, and its image is shown in Fig. S1(b), at the same scale as the images of the burnt spots. The probe pulse had a ~40 μm high-intensity core, and wings that extended over a region up to ~120 μm wide, thus fitting completely inside the area of plasma mirror. The proper alignment of the beams and the stability of alignment against beam pointing fluctuations were confirmed by the full blocking of the spectral intensity after the spectral cutoff edge (see Fig. S2).

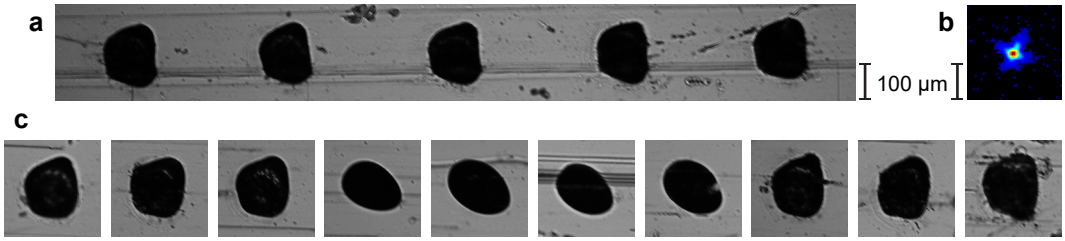

Figure S1: Beam sizes at the sample. The size of the plasma mirror was estimated as the size of burned spot on the tape, and the size of the probe beam by imaging. All images have the same scale. (a) Consecutive burned spots from a run. (b) Image of the probe beam at the tape. (c) Burned spots from different runs, slightly different due to changes in alignment and focusing.

2. Delay calibration

The calibration of the delay measurements consists in determining the relationship between the cutoff wavelength and an introduced delay. Fig. S2 displays several transmitted spectra of the long pulse, taken during a 300 s long scan of the delay stage over 600 µm of optical path length difference (OPLD). The OPLD variation rate was 2 µm/s, hence the scanned temporal delay window was 2 ps wide. For the starting spectrum (t = 0 s), the short pulse arrived after the end of the long pulse, and the entire spectrum was transmitted. The other three spectra were measured at 120 s, 170 s, and 200 s respectively. These spectra are narrower and have a cutoff wavelength; the missing part of the spectra corresponds to the spectral components in the long pulse that arrived after the plasma mirror turned on. Figure S3 shows all the spectra in the scan as a color-coded 2D spectrogram that also includes the cutoff wavelengths. Figure S4 displays the scanned delays as a function of cutoff wavelengths. These data were fit with a fifth-order polynomial function, which was then used to convert cutoff wavelengths into relative delays.

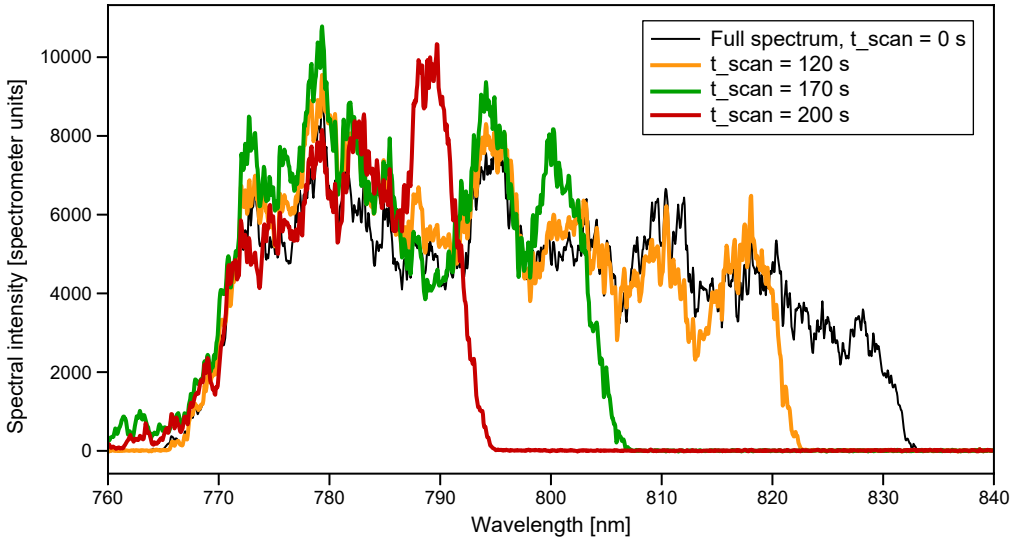

Figure S2. Transmitted spectra of the long pulse during the calibration scan. During the scan parts of the spectra are blocked, starting at a cutoff wavelength, due to the short pulse generating a plasma mirror that reflects the late-arriving part of the chirped long pulse.

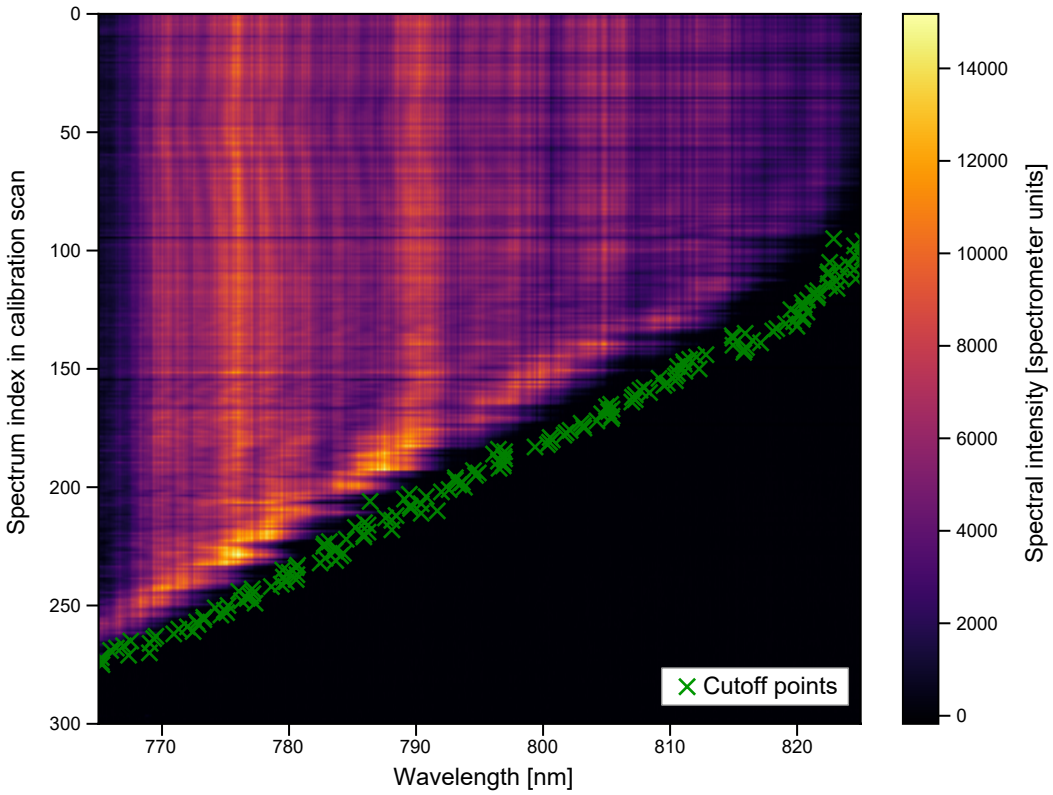

Figure S3. Full calibration data. The spectrogram shows all spectra collected in a calibration scan of the delay line over 300 µm (600 µm optical path length difference, corresponding to 2 ps delay) in optical path steps of 2 µm (corresponding to 6.67 fs). The cutoff points, shown as green crosses, were calculated as the wavelength where the spectral intensity drops below 10 % of the maximum intensity of the spectrogram.

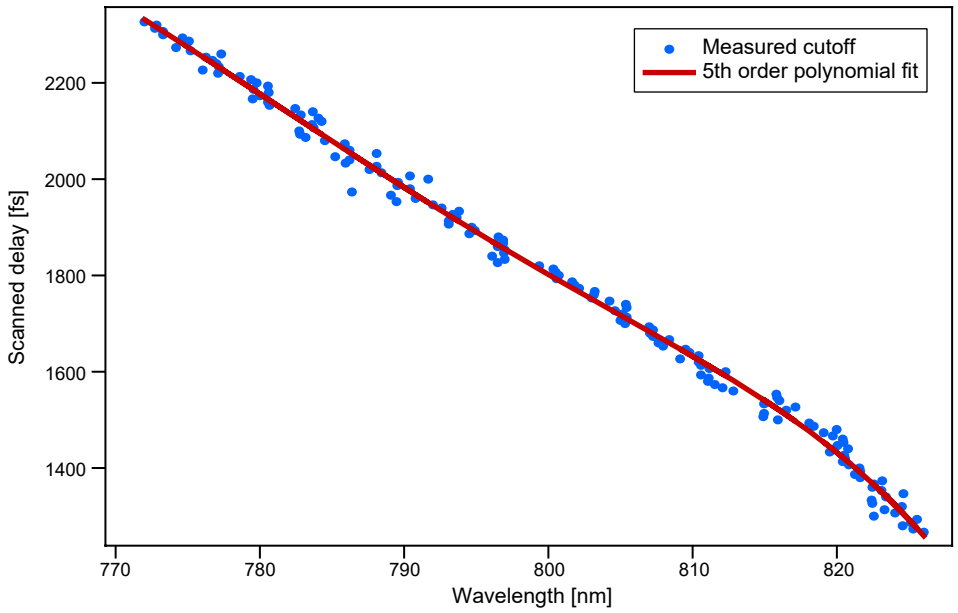

Figure S4. Calibration of the relative delay as a function of the cutoff wavelength. The cutoff wavelengths were fitted with a fifth order polynomial, which was then used to convert the cutoff wavelength to the relative time delay.

3. Stability of the experimental conditions

The stability of pulse spectra and energy were evaluated using the diagnostics data recorded automatically by the HPLS. The spectra for the pump pulse (plasma mirror) and the probe pulse (stretched) were recorded by spectrometers that sampled the beam after the amplifiers A.2 on both arms.

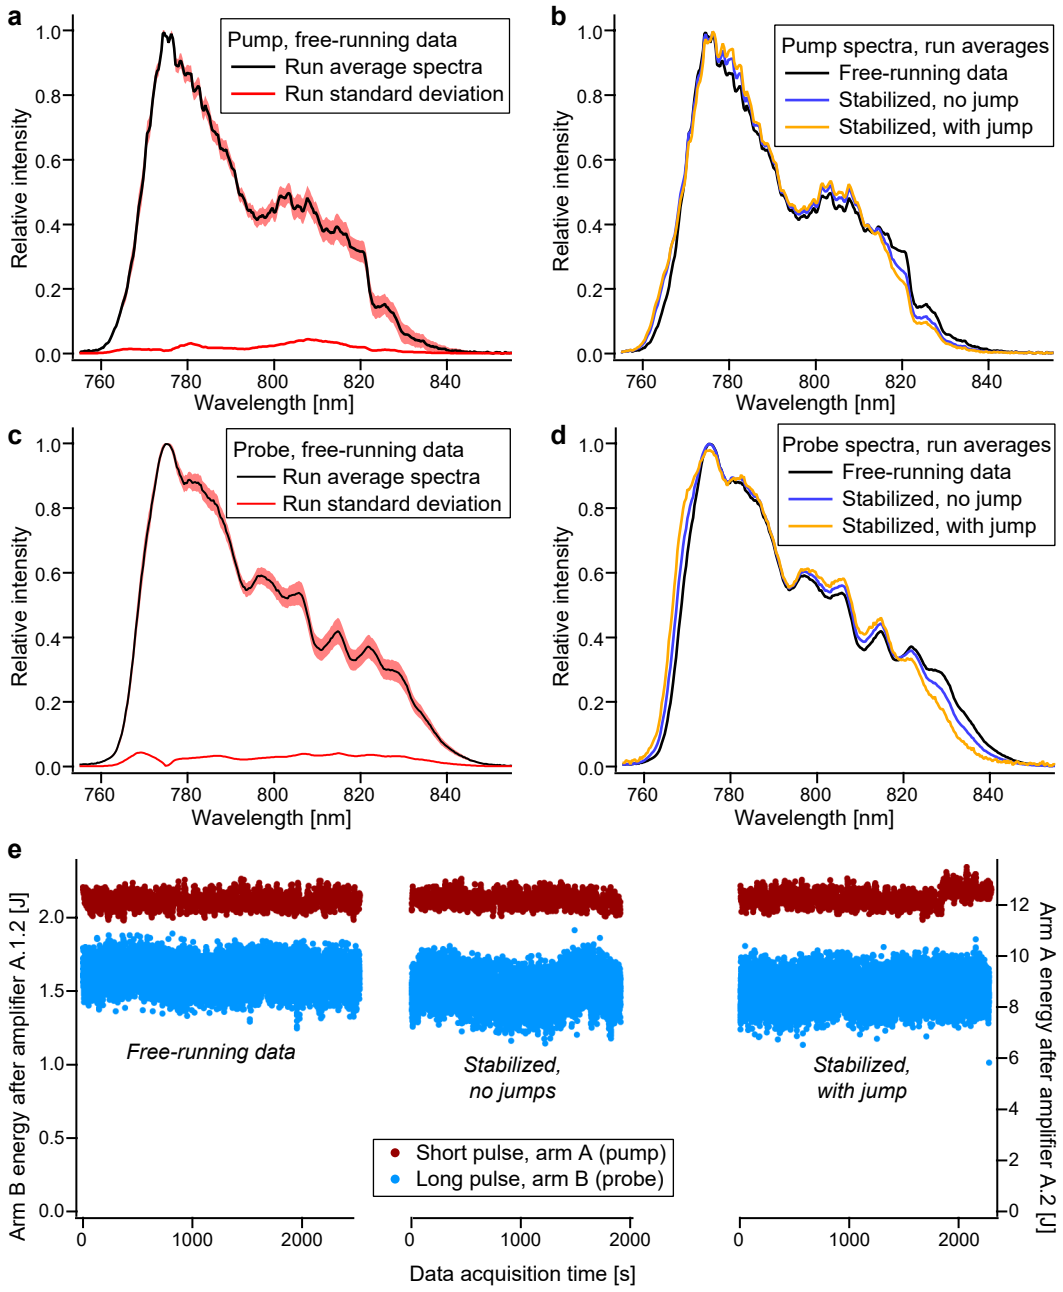

Figure S5. Stability of the pump and probe pulses during the experimental runs. (a-d) Stability of the spectra within a run, and variations of spectra between runs. In panels (a,c), the standard deviation is represented both as a stand-alone curve and as an uncertainty band around the average. (e) Stability of the pump and probe pulse energies during all experimental runs reported here.

The energies were measured by pulse energy meters located after the final stages of amplification for each pulse: amplifier A.2 for the pump, and amplifier A.1.2 for the probe. The data corresponding to a given experimental run was selected from the HPLS data based on the recording time interval of each run.

To evaluate the spectral stability, we calculated at each wavelength point, using all spectra in a run, the average intensity and the standard deviation of intensity. Figure S5(a-d) shows the average spectra and the standard deviation within a run, and the average spectra for all the runs, with all traces normalized to the maximum spectral intensity of the average. Over the range of recorded spectral cutoffs (778 to 791 nm), the relative standard deviation, *i.e.*, the standard deviation divided by the mean, was below 0.05.

The pulse energies of the full beam are shown in Fig. S5 (e-f). Within the runs, the relative standard deviation of the pump pulse energy varied between 0.02 and 0.025, and the relative standard deviation of the probe pulse energy varied between 0.05 and 0.063. Although the probe pulse energy had a larger standard deviation, the variations of the pulse energy have only a minor impact on the delay measurement, because we normalized the spectra before extracting the spectral cutoff.

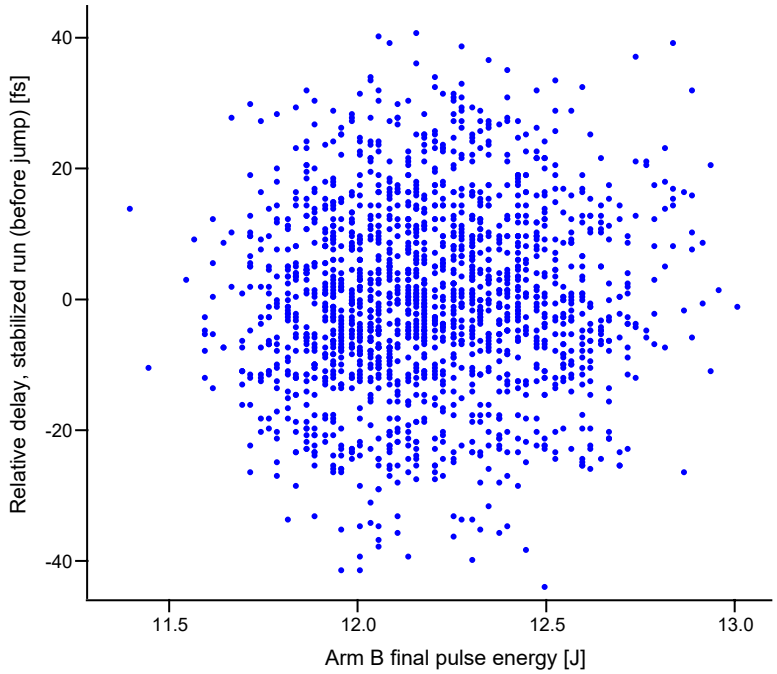

Figure S6. Measured delay and pump pulse energy for the last run, indicating a weak or negligible correlation.

For the last run, the large jump in the delay (see Fig. 4(c) in the main text) was correlated with a change of the HPLS arm B parameters, which included the pump pulse energy (see Fig. S5(e)). We used this event to synchronize the HPLS pulse energies to the delay measurements based on their individual time stamps. The relation between delay and the pump energy, for data before the jump, is shown in Fig. S6. There is no discernible correlation between the pump pulse energy and the measured delay between pulses, and a linear fit of this data has a very low coefficient of determination, $r^2 = 0.005$. This indicates that when the pump pulse energies varied by several percent, the possible delay in the formation of the plasma mirror varied less than the precision of delay measurements.

4. The turn-on time of the plasma mirror

Figure S7 shows the histograms of the turn-on time of the plasma mirror for the three runs. We defined the turn-on time as the time in which the transmission of the mirror decreased from 90% to 10% of the transmission at the local maximum closest to the spectrum cutoff. The cutoff was initially estimated as the intercept of the cut spectrum at 10% of the global maximum of the cut spectra. Since the raw spectra had modulations that could lead to a local

maximum along the edge (see the spectra at 120 s and 170 s in Fig. S2), the local maximum was determined using a running average of the raw spectra.

To check that applying the running average did not affect systematically the calculated turn on time, we alternatively determined the local maximum as the maximum in the range from 50 fs to 150 fs of the initial cutoff. The histograms calculated with the alternate method were similar to the ones shown in Fig. S7.

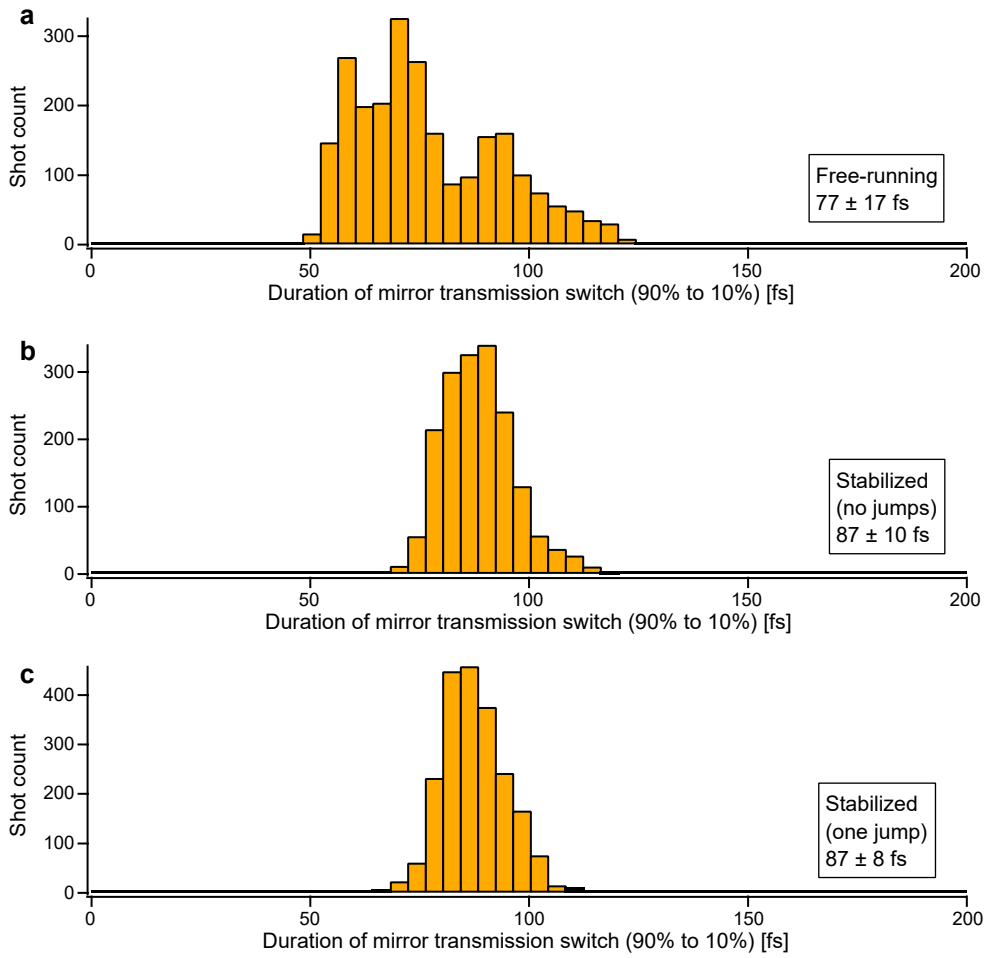

Figure S7. Histograms of the turn-on times of the plasma mirror for all runs.

5. Noise and resolution of delay measurements

Figure S8(a) shows the impact of window averaging on the free-running data. Smoothing with a window size of 16 points removes most of the very short term variations, and shows that the delay fluctuations have an additional type of variation: relatively rapid drifts with a rate on the order of 1 fs/s and an amplitude around 20 fs.

These rapid drifts are also detectable as variations of the standard deviation of data from the averaging window, as shown in Fig. S8(b). While such variations are expected as a statistical property, for a window size longer than ~10 points, the local standard deviation is approximately constant in some regions. This indicates that not only the local average of the delay but also its standard deviation can have relatively stable values on a timescale of ~10 seconds.

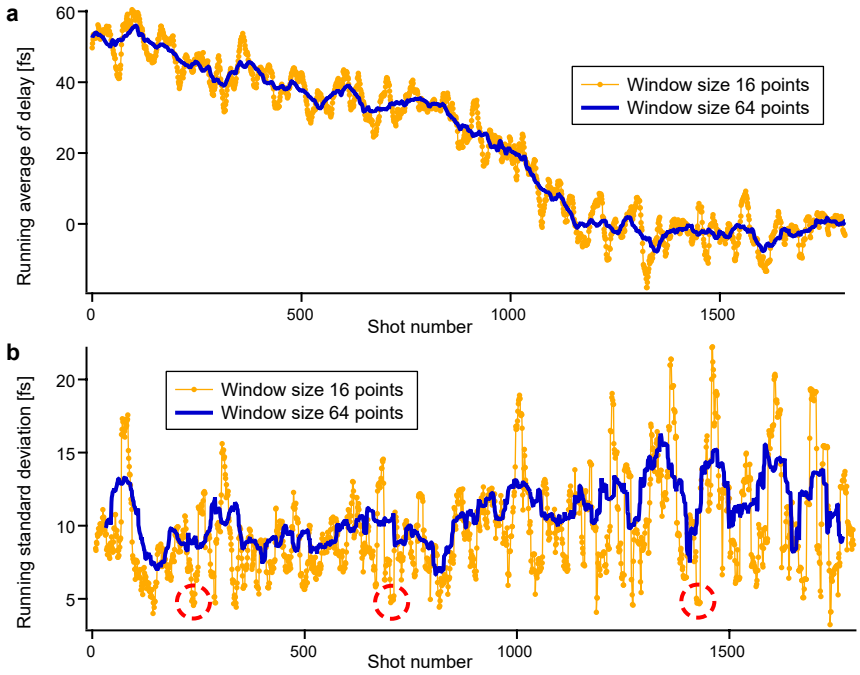

Figure S8. Analysis of delay fluctuations. (A) Running averages of the delays for window sizes of 16 and 64 points. (B) Running values of the standard deviation for the same window sizes. The red dashed circles highlight instances where the running standard deviation was around 5 fs, for 9 or 10 consecutive shots.

The smallest relatively stable standard deviation calculated with a 16-point window is approximately 5 fs, and three such instances are indicated in Fig. S8(b). Since such a behavior seems improbable if the random noise due to measurement uncertainty is larger than 5 fs, we estimated that 5 fs is an upper limit for the uncertainty, and thus for the precision, of the delay measurements.

8. Stabilization loop

The measured time delay between the two arms was used as the process variable in a proportional-integral (PI) control loop. The loop was set to stabilize the delay to a user-defined set point, by moving the translation stage. It was implemented as a computer program that calculated and sent to the translation stage a corrected value for the stage position. The adjustable PI loop variables were the proportional and integral gains, which were optimized during the measurements using the Ziegler- Nichols method [4].

The Allan variance and the power spectral density (PSD) of the free-running and the stabilized data are shown in Fig. S9. Stabilization reduced the noise of the delay over a broad frequency range and reduced substantially three peaks in the PSD of the free-running data. The virtually stabilized data (discussed later in Section 9) displayed a substantial reduction in the low-frequency noise, but enhanced slightly the high-frequency noise.

During the experimental stabilization runs, we determined the cutoff level as 10% of the maximum intensity of the transmitted part of the spectrum, with the 10% threshold chosen empirically to minimize the noise of delay measurements. Compared to more refined procedure used later in data analysis (using the average of 40% and 60% levels), the 10% criteria can increase by up to a few femtoseconds the delay noise caused by spectral intensity modulations. The larger noise due to the 10% threshold did not prevent stabilization.

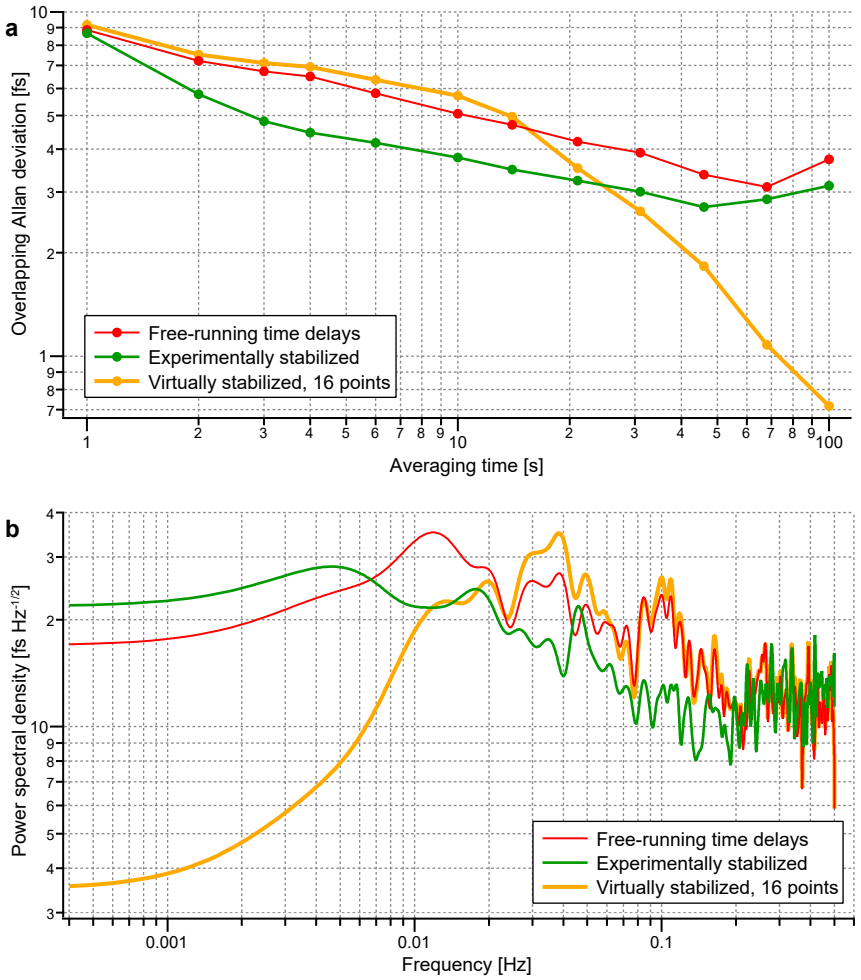

Figure S9. Frequency-domain properties of the delay fluctuations for the free-running, experimentally stabilized, and virtually stabilized data. (a) The Allan variance. (b) The power spectral density.

9. Virtual tests of delay stabilization

The active stabilization implemented experimentally does not represent the best possible stabilization performance, because, as shown in Fig. 4 of the main text, the recovery of the delay after a large jump was slow. This was due to the slow response time of the PI loop, whose order of magnitude can be estimated as the ratio of the proportional gain constant to the integral gain constants, which were tuned using Ziegler-Nichols method. The ratio of gain constants was around 50–60 s, which is consistent with the fitted exponential decay time of 129 s (Fig. 4(c), main text). Another limitation of the stabilization was due to a delay in the spectrometer readout. The data from the spectrometer was written as separate files for each shot in text format, and the readout of the data involved the directory listing of all recorded files in an experimental data set. During the experiment, we observed that the time stamp of the most recently written file was delayed by up to 3 s relative to the current computer time. This delay in the data readout led to cases where the applied delay correction did not use information from the most recent pulses, which degraded the performance of the feedback loop.

To explore the potential of feedback loops for stabilizing the delay more rapidly, we performed virtual stabilization tests on the data recorded without stabilization. We used as input the experimentally recorded data because it contains the jitter, drift, and jumps of the delay specific to the HPLS system. For the virtual tests, we assumed that the delay measurements and the delay corrections had precise and accurate values. We also assumed

that the control loop (spectrometer reading, correction calculation, correction application) was faster than the interval between pulses. Such a speed seems feasible for the 10 PW and 1 PW outputs (1/60 Hz and 1 Hz repetition rates, respectively), and may be possible for the 100 TW output (10 Hz repetition rate).

We found that a basic proportional-integral algorithm is sufficient for a substantial improvement of the response time. The algorithm we tested operates on measured delays, but it can operate on other variables that are related to the delay, such as the cutoff wavelength. The virtual algorithm calculates the average, t_{AVG}, of a number N_{DA} of the most recent delay measurements, and calculates the difference between t_{AVG} and the desired stabilized value, t_{SET}. The corrected value of the next time delay is $t_{CORR} = t_{RAW} + t_{SET} - t_{AVG}$, where t_{RAW} is the delay that would be measured without applying a correction, and t_{CORR} is the delay that would be measured if the stabilization loop sends to the delay stage a time correction equal to $t_{SET} - t_{AVG}$.

The virtual stabilization is illustrated in Fig. S10. We stabilized the free-running data set shown in Fig. 4(a) in the main text, which includes a drift of ~50 fs over 1800 shots (Fig. S8). Since we did not have a raw data set that contained a large jump in the delay, we generated simulated jump data by increasing by 100 fs all delays in the second half of the free-running data, as shown in Fig. S10(a). We chose $N_{DA} = 16$, because this setting represents a compromise between a faster response (which requires smaller N_{DA}) and a smaller standard deviation of the stabilized data (which decreases with larger N_{DA}). The reference time delay t_{SET} was the average of the first 16 delays in the free-running data.

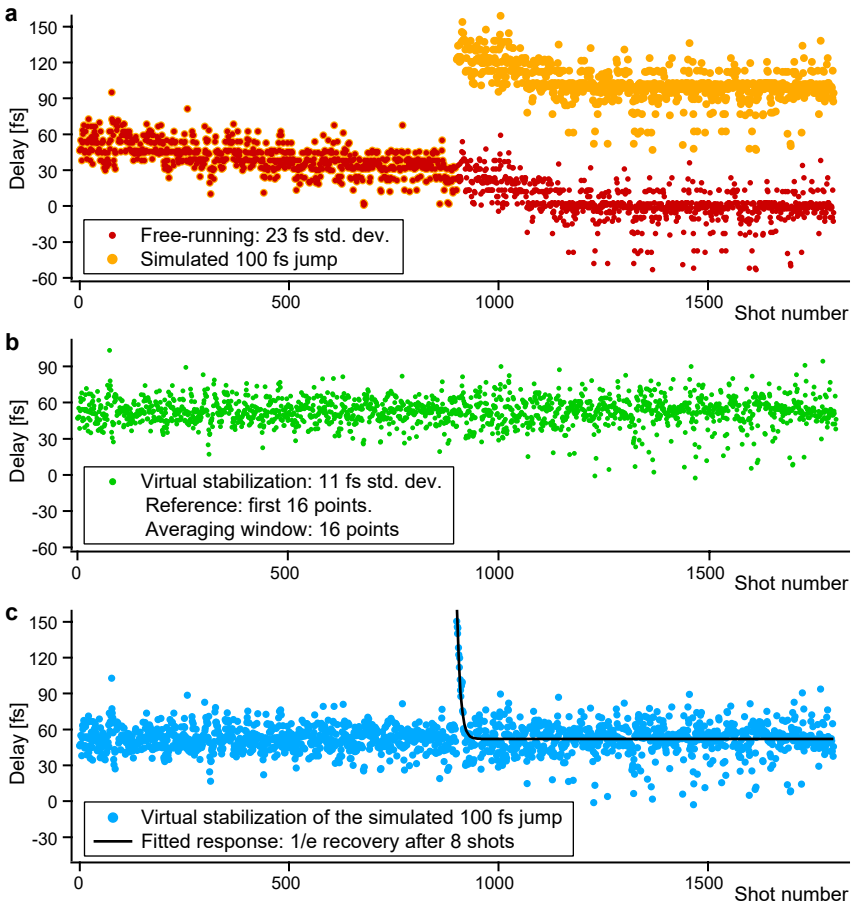

Figure S10. Virtual stabilization of the time delays. (a) The inputs are an experimental free-running data set, and a simulated jump data set obtained by adding 100 fs to the second half of the free-running data set. (b) Stabilization of the free-running data set, showing the removal of the drift in the free-running data. (c) Stabilization of the delay jump data set. After the jump, the delay was fitted with a decaying exponential.

For the free-running data set, the virtual stabilization corrected the drift, and the standard deviation of the corrected data was the same as the one achieved experimentally. Due to the small value of N_{DA}, the feedback loop was more responsive and the stabilized data in Fig. S10(b) did not display noticeable modulations. For the delay jump data set, the corrected data, shown in Fig. S10(c), exhibited a very rapid stabilization, with a fitted 1/e exponential decay time of 8.4 shots (or 8.4 seconds), more than 10 times shorter than obtained experimentally.

References

1. F. Lureau, G. Matras, O. Chalus, C. Derycke, T. Morbieu, C. Radier, O. Casagrande, S. Laux, S. Ricaud, G. Rey, A. Pellegrina, C. Richard, L. Boudjemaa, C. Simon-Boisson, A. Baleanu, R. Banici, A. Gradinariu, C. Caldararu, B. De Boisdeffre, P. Ghenuche, A. Naziru, G. Kolliopoulos, L. Neagu, R. Dabu, I. Dancus, and D. Ursescu, "High-energy hybrid femtosecond laser system demonstrating 2 x 10 PW capability," High Power Laser Science and Engineering 8, e43 (2020). DOI: https://doi.org/10.1017/hpl.2020.41
2. T. Oksenhendler, S. Coudreau, N. Forget, V. Crozatier, S. Grabielle, R. Herzog, O. Gobert, and D. Kaplan, "Self-referenced spectral interferometry," Applied Physics B-Lasers and Optics 99, 7-12 (2010). DOI: https://doi.org/10.1007/s00340-010-3916-y
3. P. Tournois, "Acousto-optic programmable dispersive filter for adaptive compensation of group delay time dispersion in laser systems," Optics Communications 140, 245-249 (1997). DOI: https://doi.org/10.1016/s0030-4018(97)00153-3
4. J. G. Ziegler and N. B. Nichols, "Optimum settings for automatic controllers," Journal of Dynamic Systems Measurement and Control-Transactions of the ASME 115, 220-222 (1993). DOI: https://doi.org/10.1115/1.2899060