LCLS-II Technical Note

High-repetition rate laser heater shaping for femtosecond slicing at LCLS-II

LCLS-II TN-21-12

10/28/2021

Joseph Duris, Sergio Carbajo, Randy Lemons, Agostino Marinelli

Abstract:

In this technical note, we present plans for laser heater shaping of the electron beam at high repetition rates for LCLS-II. This will allow selective beam manipulation for several femtosecond pulse slicing, subfemtosecond pulse slicing, and subfemtosecond pulse trains. This note details the setup and simulations for the first application.

1 Introduction

XFELs use a high-brightness electron beam as their lasing medium by amplifying spontaneous shot noise to a high power as the electron beam propagates through an undulator magnet. This property enables shaping the produced X-rays pulses by shaping the electron beam. X-ray FELs typically control the pulse duration by controlling the length of the electron beam or passing parts of a dispersed electron beam through a slotted foil to spoil the transverse emittance of the electron beam at various longitudinal positions (Ding, et al., 2015). This technique is impractical for the LCLS-II for two main reasons. First, high repetition rates may lead to unacceptable radiation losses and damage to the foil. Second, LCLS-II is envisioned to multiplex to different beamlines for simultaneous operation of different experiments, yet spoiling via masks in a dispersive region would affect all delivered beams. Different experiments may require different pulse durations, necessitating a selective and controlled method of beam shaping capable of multiplexing beams with different properties.

Laser heater temporal shaping is one solution to this problem as it has been shown to selectively heat parts of the electron beam and produce X-rays pulses as short as 11 fs (Marinelli, et al., 2016). Furthermore, by switching the laser on and off in time with the RF kickers, undulators in different beamlines can produce different pulse durations. Typically, a laser heater is used before the bunch compressors to heat the electron beam enough to suppress the microbunching instability which may otherwise spoil the electron beam before it reaches the undulator (Huang, et al., 2010). However by temporally shaping the laser heater pulse, different parts of the electron beam maybe heated enough to even suppress the FEL lasing in the undulators downstream. The energy spread in the undulators needed to suppress the FEL lasing must greatly exceed the FEL resonant energy bandwidth which is typically on the order of 0.1% for soft X-ray FELs or at least a few MeVs for a few GeV electron beam energy. This necessitates an induced energy spread of at least several tens of keV after the laser heater. The bunch compressors reduce the electron beam's duration by a factor of 100, allowing picosecond manipulations of the laser heating the beam to control the lasing on a scale of tens of femtoseconds.

2 High-repetition rate laser heater

The roughly 2 km transport line from linac to undulators makes LCLS-II especially susceptible to the microbunching instability, and the LCLS-II laser heater plays a crucial role suppressing microbunching (LCLS-II Conceptual Design Report, 2014). LCLS-II will operate at 1 MHz so will use a high repetition rate Amplitude Systems Tangerine laser to drive UV generation for the cathode as well as the laser heater (Gilevich, et al., 2020). The laser is transported from the laser room via vacuum to the laser heater undulator (9 periods, each 5.4 cm) several meters downstream of the injector where it is focused to a waist of 180 um to heat the 100 MeV electron beam. The RMS energy spread in keV induced by the laser heater (with power specified in kW) is given by $\sigma_{\Delta E} = 0.86\ keV\ \sqrt{P_{laser}[kW]}$.

Since the laser heater is so critical to LCLS-II, a separate dedicated laser is planned for shaping the beam to avoid interrupt basic operations. The shaping laser is a Light Conversion Carbide CB3-40W 1030 nm laser capable of producing 0.4 mJ pulses at repetition rates of up to 1 MHz for an average power of 40 W. The 8 nm bandwidth supports FWHM pulse durations as short as 250 fs, which could induce modulations on the electron beam with durations comparable to the FEL cooperation length after acceleration and (100x) compression. The laser has a pulse picker to deliver pulses to two outputs during operation; however, as these pulses are separated in time by the reciprocal of the repetition rate, they cannot be overlapped in time.

Instead, LCLS-II will use an Amplitude Systems Tangerine laser for baseline microbunching suppression via laser heating. These two separate lasers may be combined so that shaped pulses from the one laser achieves pulse shaping goals while the Tangerine laser suppresses the microbunching instability.

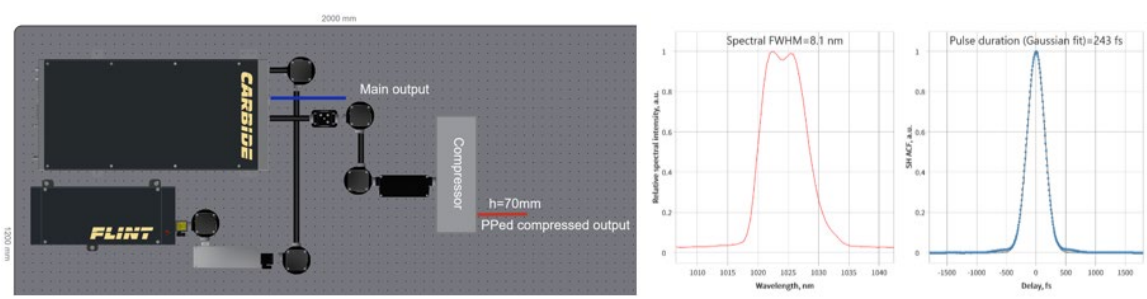

Figure 1 Layout of the shaping laser, spectrum and pulse autocorrelation at 1 Mhz, 40 W.

3 Femtosecond slicing via pulse stacker

3.1 Pulse stacking with a Michelson interferometer

Temporal control may be achieved by stacking copies of the same pulse displaced in time via a Michelson interferometer (Greco, Molesini, & Quercioli, 1995) as shown in Figure 2. The fast axis of a half waveplate (HWP) is rotated by 22.5 degrees relative to the horizontal linear polarization of an input pulse to equally split the transmitted and reflected power from a polarizing beam splitter (PBS). Normal incident mirrors

(M0) return the pulse in each arm to the PBS, and one of these mirrors is placed on a linear stage to control the delay between the two pulses.

During its return trip to the PBS, the polarization of each pulse is rotated by 90 degrees via two passes through a quarter waveplate (QWP) with an axis at 45 degrees relative to the horizonal. This change in polarization causes the originally transmitted pulse to reflect and the reflected pulse to transmit through the PBS so that both pulses are sent to the 45-degree mirror (M45) which directs the power to a half waveplate. Each pulse is orthogonally polarized and equal in intensity, so the final HWP is oriented 22.5 degrees from the horizontal to maximize the projection of each pulse onto the horizontal plane to match the transverse oscillation plane of the undulator. The vertical polarization of each pulse is reflected to a dump by a polarizing beam splitter. The total power of the laser may be controllably attenuated with the addition of another waveplate and polarizer after the final one (not shown in the figure).

By rotating the final HWP angle φ, the power in each pulse can be controlled as $P_{\pm} = (1 \pm \cos(4\,\varphi))\,P_0/4$ where P_0 is the power of the incident laser pulse. This enables control of the fraction of power in each pulse $P_{\pm}$ and gives a net efficiency of 50% for the total power transmitted from the interferometer (the bulk of the rest of the power is in the wrong polarization).

For delays comparable to the input pulse duration, the stacked pulses will interfere with each other, leading to ripples in the power profile with a period given by $T = 2\pi/\omega_1\delta t$ where δt is the delay and the chirp is given by $\omega_1 = \sqrt{\Delta t^2\Delta\omega^2 - 16\,ln\,2}/\Delta t^2$ with Δt and $\Delta\omega$ being FWHM measures of the pulse in the time and frequency domain, respectively. Additionally for larger delays, it is necessary to combine these pulses with a long duration laser to suppress the microbunching instability. This also results in temporal ripples as shown in Figure 2. These ripples are imprinted in the energy modulation after the laser heater but are washed out after acceleration and compression and result in increased energy spread.

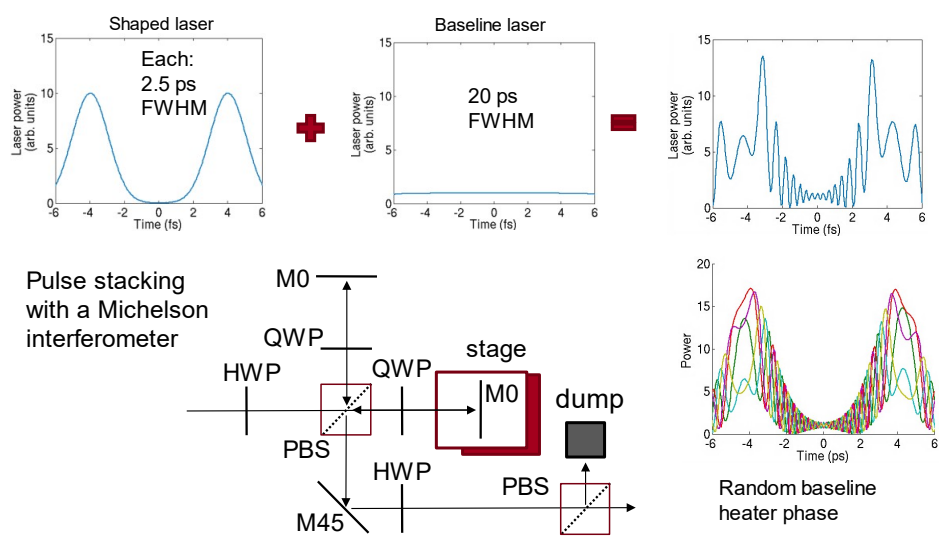

Figure 2 A 2.5 picosecond laser pulse is equally split and delayed by 8 ps with a Michelson interferometer. Upon stacking with the baseline laser, the pulse profile has fringes separated by hundreds of femtoseconds.

3.2 Beam dynamics and FEL simulations

In order to assess the impact on the X-ray FEL pulse, we simulate the entire FEL beamline from the output of the photoinjector to the LCLS-II soft X-ray undulators with Elegant (Borland, 2000), and then simulate the FEL interaction with Genesis 1.3 (Reiche, 1999). Figure 3a shows the laser profile resulting from stacking two 2.5 ps FWHM duration pulses separated by 6 ps combined with a stretched 20 ps FWHM pulse. All pulses have identical spectra peaked at 1030 nm and with an 8 nm FWHM bandwidth, but the stretched pulse is given a random CEP phase since it is assumed to come from a different laser. The electron beam used in the simulation has a charge of 100 pC, an energy of 100 MeV, and was simulated from Astra simulations with 10 million macroparticles of the LCLS-II injector with a Gaussian laser profile (Neveu & Ding, 2021). An elegant simulation of the electron beam using the "LSRMDLTR" simulation element (Figure 3b) shows a 2 to 3 ps region of low energy spread beam flanked on either side by regions of large energy spread where the laser power was large. This heated charge is smeared out by R56 in the bunch compressors within the linac while the cold core is compressed to > 1 kA (Figure 3c). This 20 to 30 fs slice of beam has a current that's more than twice the rest of the beam, and longitudinal impedance in the long transport line from the linac to the undulators linearly chirps this short beam. A Genesis 1.3 FEL simulation was performed with the LCLS-II soft X-ray focusing and undulator lattice tuned to lase at 700 eV. Figure 3d shows the X-ray pulse duration is comparable to the cold part of the beam just after saturating in 10 undulators. The same simulation was repeated for various delays between the two stacked pulses, and the pulse energies and durations for each delay shown in Figure 3e are the result of averaging over several simulations with different initial electron beam shot noise.

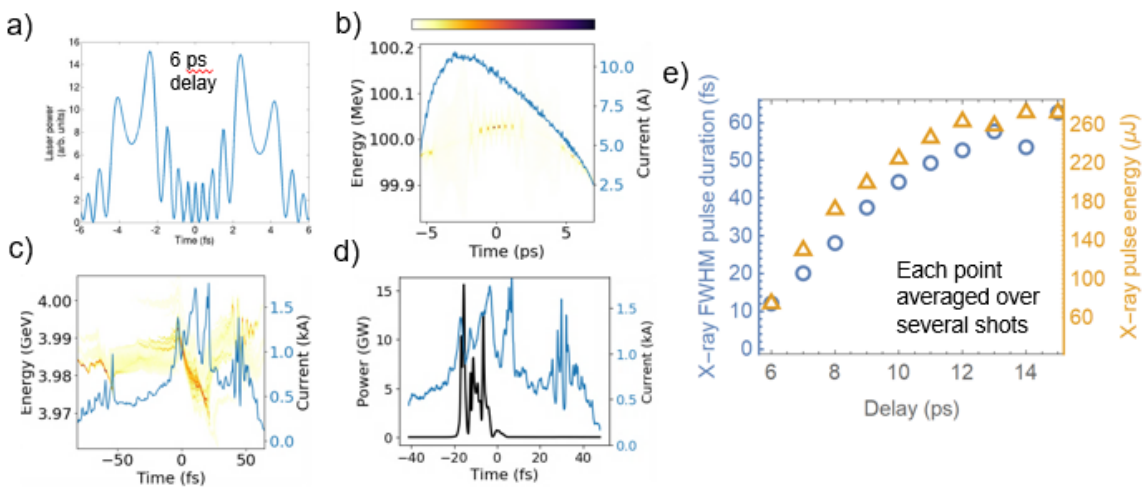

Figure 3 Simulations of slicing a 100 pC beam via pulse stacking. a) Two 500 kW, 2.5 ps phase-locked pulses are stacked 6 ps apart on top of a 50 kW, 20 ps randomly phased pulse. b) The heated electron beam profile after the laser heater has a ~2 ps cold core. c) After acceleration and compression, the cold beam core is about 20 fs whereas the heated regions have large energy spread and lower current. d) The X-ray FEL pulse duration at saturation is comparable in duration to the cold core. e) The X-ray pulse energy and duration (each averaged over several shots) varies nearly linearly with interpulse separation. The head is to the left in each panel.

As the pulse separation delay is decreased below 6 ps, ripples begin to fill in the temporal notch in the laser pulse profile, and parts of the head and tail of the beam may be unheated leading to other parts of the beam capable of lasing (see Figure 3c). These can be mitigated by attenuating the long laser and lengthening each stacked pulse to reduce the width of the temporal notch to further reduce the size of the cold electron beam core after the laser heater, but at this point space charge impedance of the long transport line from linac to undulators chirps and stretches the cold part of the electron beam.

A lower charge beam from the injector (Neveu & Ding, 2021) can help mitigate these problems in two ways. First, the shorter beam in the laser heater is easier to cover with shorter pulses, allowing narrower temporal notches while covering the edges of the beam (Figure 4b). Second, the lower beam current reduces the space charge induced chirp and resulting stretch in the long transfer line to the undulators (Figure 4c), enabling X-ray pulses with durations shorter than 10 fs (Figure 4d). The response of the FEL pulse energy and pulse duration are reasonably linear in response to the laser pulse stacker delay (Figure 4e), and the minimum pulse duration with at least 20 uJ of pulse energy is about 4 fs. The 100 pC beam is sufficient for pulse durations as short as 10 fs whereas the 20 pC beam is good for pulses between 4 and 14 fs. These pulse durations may be estimated online from X-band deflecting cavity streaked temporal measurements of the region of beam with energy lost to the X-ray field.

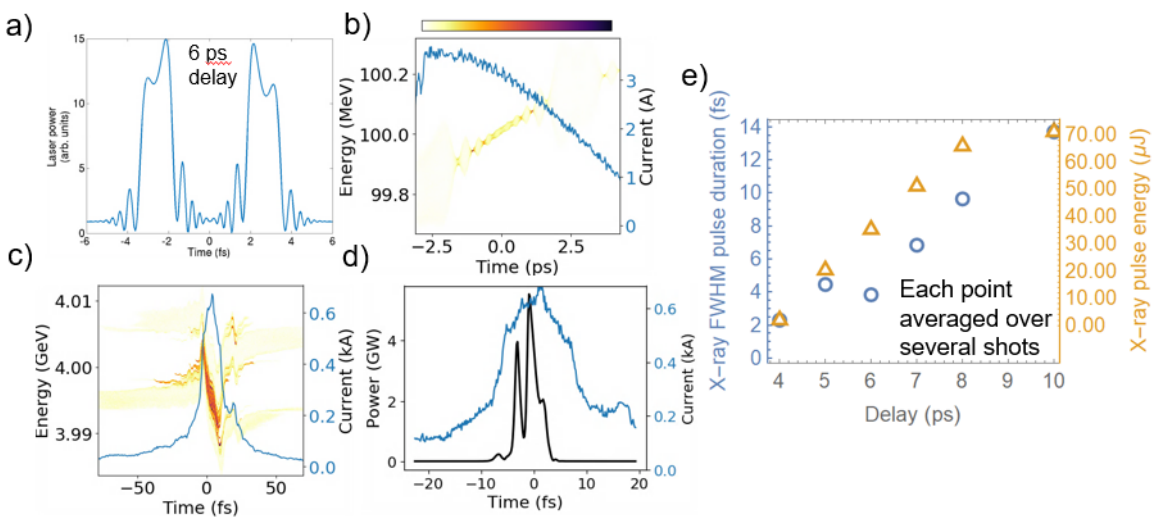

Figure 4 Simulations of slicing a 20 pC beam via pulse stacking. a) Two 500 kW, 1.5 ps phase-locked pulses are stacked 6 ps apart on top of a 50 kW, 20 ps randomly phased pulse. b) The heated electron beam profile after the laser heater has a ~2 ps cold core. c) After acceleration and compression, the cold beam core is about 13 fs. d) The X-ray FEL pulse duration at saturation les than 10 fs. e) The X-ray pulse energy and duration (each averaged over several shots) varies nearly linearly with interpulse separation. The beam's head is to the left in each panel.

4 Slicing via spectral filtering

An alternative approach to creating a notched beam is by applying a notch filter to a linearly chirped laser

pulse. Volume Bragg gratings (VBG) are tunable narrow band notch filters capable of coring out a linearly stretched pulse. Alternatively, a diffractive grating can be used in a zero delay compressor in order to present the Fourier plane of a chirped laser for direct manipulation of the spectrum. We compare the effectiveness of these below.

4.1 Notch filtering with a volume Bragg grating (VBG)

Volume Bragg gratings (or volume holographic gratings) are gratings recorded into the bulk of photo-thermo-refractive (PTR) glass. A Bragg mirror has an index of refraction with a sinusoidal modulation along a direction within its bulk (in this case, normal to the surface) $n(z) = n_1 + \Delta n_1 \sin(2\pi z/\Lambda)$, where n_1 is the bulk index of refraction, Δn_1 is the perturbation amplitude, and Λ is the grating pitch. Waves scattering coherently from the sinusoidal gradient in the index of refraction satisfy the Bragg condition $m\,\lambda_B = 2\,n_1\,\Lambda\cos\theta_1$, where m is an integer representing the order of the diffraction and θ_1 is the angle of the incident radiation within the bulk of the material relative to the surface normal. This internal propagation angle θ_1 is related to the incident angle θ_0 of the radiation from air (index of refraction n_0) by Snell's law $n_1 sin\,\theta_1 = n_0\,sin\,\theta_0$. Unscattered waves are unaffected by the VBG. By varying the incident angle θ_0, the central wavelength of the scattered radiation can be controlled. The peak reflectivity R_B and bandwidth $\Delta\lambda_B$ of the Bragg resonance are determined by the depth d and the bulk properties of the grating by the relations $R_B = tanh^2(\pi\,\Delta n_1\,d\,/\lambda_B\,cos\,\theta_1)$ and $\Delta\lambda_B = 2\,n_1\,\Lambda\cos\theta_1\sqrt{\frac{\Delta n_1{}^2}{n_1{}^2\cos^4\theta_1} + \frac{4\,\Lambda^2}{d^2}}$. Thus, thinner gratings offer wider scattering bandwidths and still have high efficiencies if thicker than $d > \lambda_B\,cos\,\theta_1/\Delta n_1$, or >100 wavelengths as Δn_1 is typically ~0.001 (OptiGrate, 2020) (Glebov, 2008). Furthermore, gratings can be stacked to achieve wider notches.

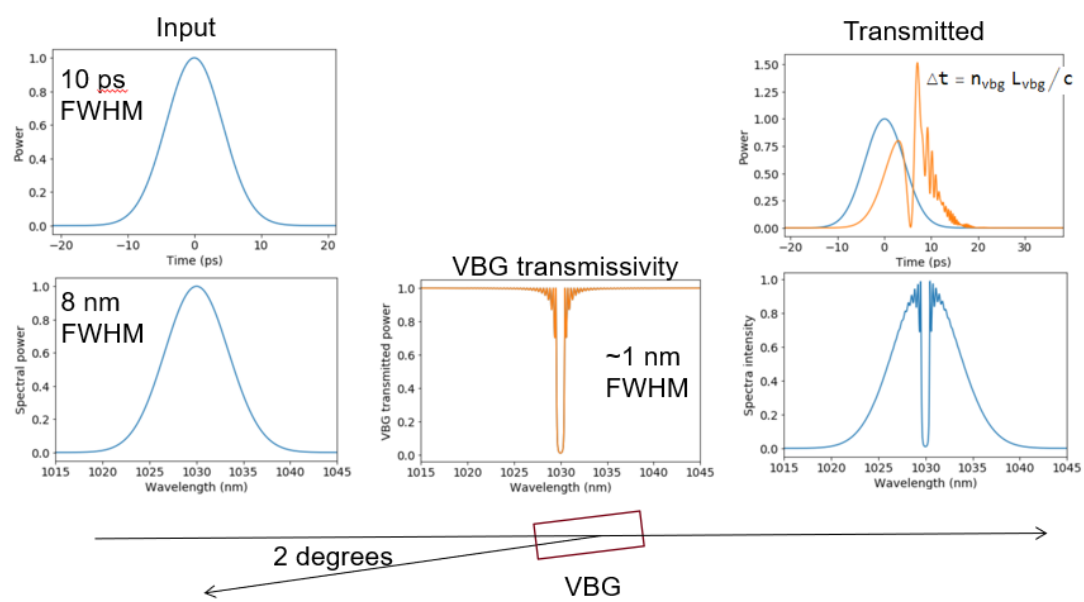

Figure 5. A linearly chirped pulse is cored when passed through a 1 nm VBG notch filter. The notch can be blue shifted by increasing the radiation's incidence angle. The pulse's head is to the left in each panel.

To estimate the effect of the VBG on a linearly chirped laser pulse passing through it, we use an ideal VBG model (Jacobsson, 2008). The laser is a linearly chirped, 1030 nm, Gaussian pulse, as shown in the left side of Figure 4, with 8 nm FWHM bandwidth and 10 ps FWHM duration sufficient to cover the ~6 ps electron beam in the laser heater. The modeled VBG has an index of refraction $n_1 = 1.48$, amplitude of sinusoidal index variation $\Delta n_1 = 0.001$, grating pitch $\Lambda = 348$ nm for a Bragg resonance of 1030 nm at normal incidence, and a thickness of $d = 1$ mm for a FWHM bandwidth of 1 nm as shown by the transmittance of the VBG in the middle plot of Figure 4. The filtered pulse, show on the right side of Figure 4 has a FWHM notch of width 2.8 ps and is delayed relative to the input pulse by the time spent traversing the VBG.

One benefit of the large peak transmissivity of the VBG is that most of the laser power may be used efficiently to heat electrons (<10% of the power is reflected by the VBG) if the duration of the input laser is comparable to the electron beam duration within the laser heater. This makes it suitable for slicing the shortest pulses. However, the spectral notch width cannot be readily controlled without stacking multiple VBGs. Instead, the temporal notch width can be controlled by varying the input pulse duration for variable electron beam slicing, placing a lot of laser power off the electron beam and making it inefficient for controlling the FEL pulse duration.

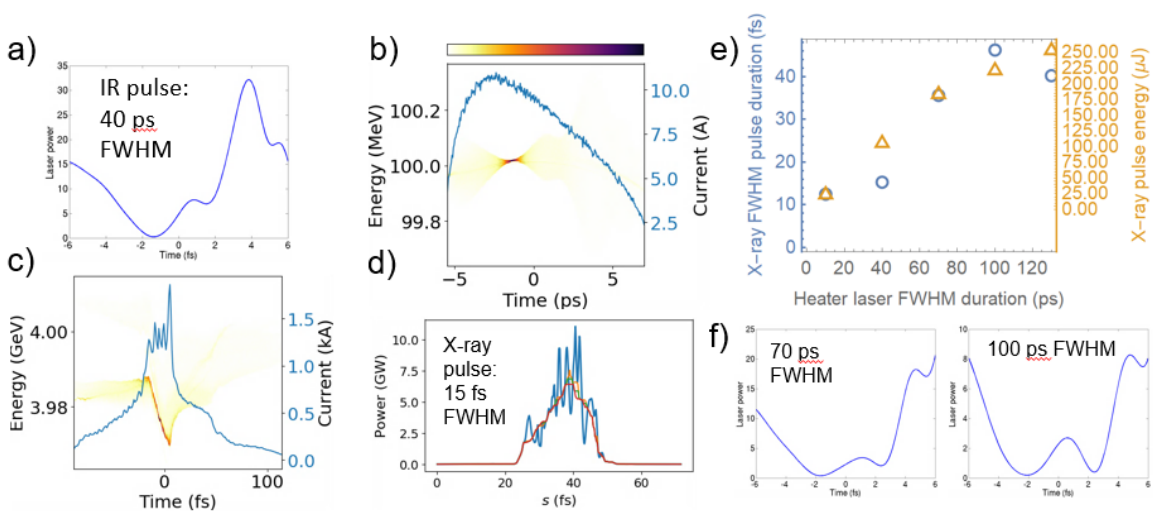

Figure 6. Simulations for VBG spectrally filtered pulses. a) A 1nm FWHM notch filter from the VBG puts a few ps hole in the spectrum of a 0.8 MW, 40 ps FWHM linearly chirped Gaussian pulse. b-c) Longitudinal phase space at the end of the laser heater and start of the SXR undulator line. d-e) FEL pulse profile and response to the input pulse duration as the input peak power is held fixed. f) Further stretching the input laser results in isolated regions of low laser power. The head is to the left in each panel except for panel d, for which the head is to the right.

Elegant and Genesis simulations repeated with the same parameters as previously described in Section 3.2 are shown in Figure 5. Whereas the X-ray pulse energy varies linearly with the input laser pulse duration (while keeping the input peak power fixed), the X-ray pulse FWHM does not (Figure 6e). This is due to the ripples in the filtered pulse's time domain leaving two isolated regions of cold beam for large input pulse durations (see Figure 6f). That makes this approach fine for the shortest sliced beams (~10 fs X-ray power profile), but it does not lend itself to an easily understandable knob for controlling the FEL pulse duration.

Note that the pulse energy of a 0.8 MW peak power, 10 ps FWHM laser is about 8 uJ per pulse, or an average of 8 W at 1 MHz repetition rate.

4.2 Notching with a zero-dispersion stretcher

To better control the spectral notch for filtering the beam, a zero-dispersion Martinez stretcher may be used to expose the Fourier plane of the laser to introduce filters. By adjusting the telescope between two gratings to image one plane at the other, the group delay dispersion of the stretcher is eliminated. A mask can then be inserted normal to the dispersive direction to mask a part of the spectrum. If the mask is wedged, then the width of the spectral notch is adjustable. The temporal notch may then be controlled by varying the duration of the input laser or the width of the spectral notch. The former leads to inefficiency from the laser being longer than the electron beam while the latter's efficiency is limited by the power removed by the mask.

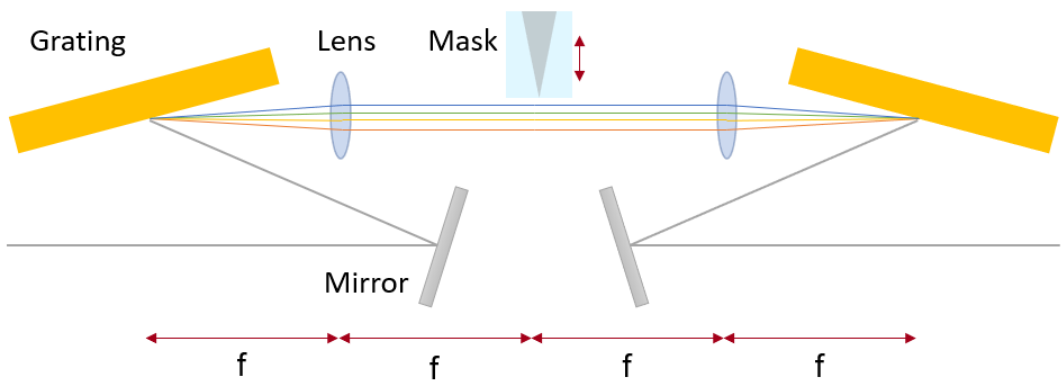

Figure 7. A Martinez stretcher with telescope configured so that the image of one grating coincides with the position of the other (two focal length f lenses) results in zero group delay dispersion (GDD). A partially silvered wedge mask can be inserted in the dispersive plane to shape the spectra.

Elegant and Genesis simulations of the pulse shaping are shown in Figure 8. A partially transmissive mask (e.g. silver deposited on glass) blocks 90% of the peak spectral intensity (Figure 8a), forming a 4 ps wide notch in the time domain (Figure 8b), and 58% of the pulse energy is transmitted through the mask. This leaves some power in the few fs wide temporal notch to suppress the microbunching instability for the cold part of the beam without the need for combining with another laser which would otherwise cause ripples in the time domain (see Figure 2). We also scale the laser pulse energy to keep the laser power in the middle of the temporal notch equal to 50 kW. Figure 8c shows a 2 ps cold core bounded by regions of enhanced energy spread after the laser heater, while Figure 8d shows that the cold core is compressed to 20-30 fs after acceleration and compression. The resulting power profile after lasing to saturation in an untapered soft X-ray undulator line is 20-30 fs long. The X-ray pulse energy and duration are controllable by varying the input laser heater duration, and the shortest pulse obtained is about 10 fs. Whereas the energy and duration of the shortest pulses are comparable to the simple pulse stacking setup before, the X-ray pulse properties vary less linearly than with the pulse stacker.

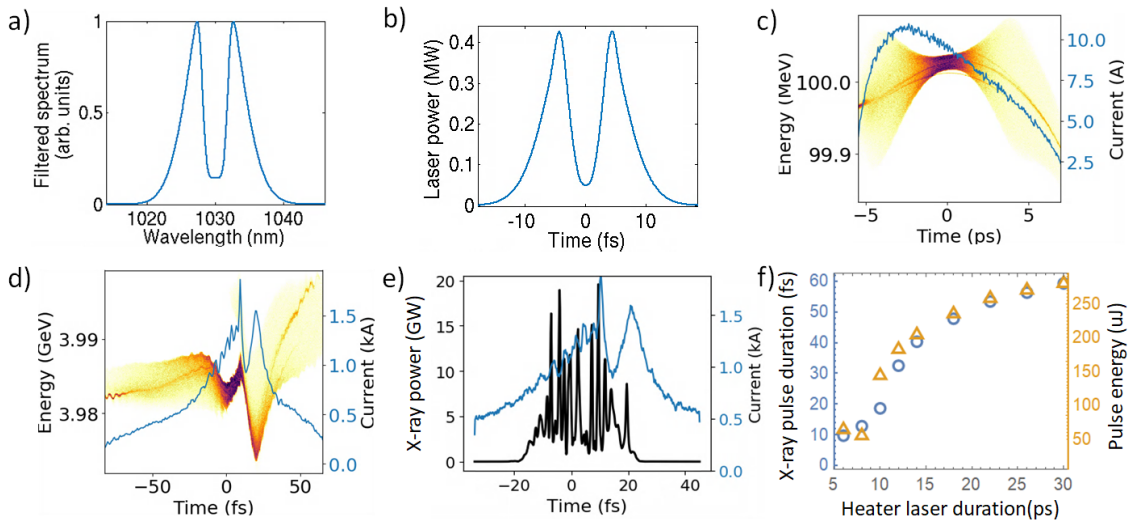

Figure 8. Simulations with partial spectral filtering via a dispersionless Martinez stretcher. a) and b) shows the 4 nm spectral notch and temporal notch in an 8 ps pulse with 8 nm bandwidth. The laser pulse energy is scaled to keep the laser power in the center of the notch equal to 50 kW. The head is to the left of each figure.

4.3 Pairs of pulses

This setup is flexible enough for temporal shaping of two pulses by placing two notch masks the in the dispersed spectral plane of the Martinez stretcher as shown in Figure 9. Nonlinearities in the injected beam and transport lead to different pulse durations for the same width spectral notches depending on their positioning. The width of the notch may be varied to control the electron beam as measured by the X-band deflecting cavity, and the positioning of each may be controlled to control the delay between each pulse. In this way, pairs of pulses for pump-probe experiments may be made.

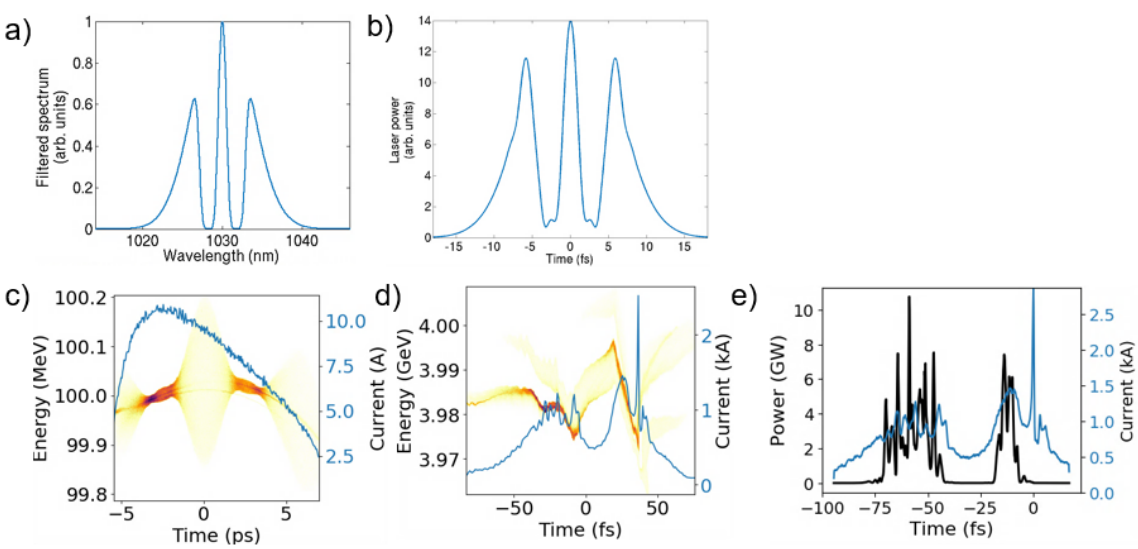

Figure 9. Simulations of two X-ray pulses created via spectral filtering with a dispersionless Martinez stretcher. Two temporal notches are created by masking the spectrum to create two notches. The beam's head is to the left in each panel.

5 Vacuum window damage threshold

The picosecond laser pulses with high peak powers and high repetition rates incident on the vacuum injection window may lead to material damage, altering the injected laser, and possibly necessitating an invasive downtime for a vacuum window replacement. To assess the threshold laser fluence which damages the vacuum window, the Carbide laser (see Section 2) was focused with a 75 cm focal length lens to a waist of approximately 260 µm. The focus was gentle enough that no breakdown in air was observed during any of the tests. All tests were performed with the Carbide laser operating at a repetition rate of 100 kHz and compressed with the external compressor. Average power incident on the test vacuum window was measured via bleed through from a dielectric mirror positioned after the window.

The first tests placed the window downstream of the focus at a position where the laser waist was 0.51 mm as measured by observing 92% power transmitted through an iris with 1.16 mm diameter aperture. The maximum average power incident on the window was 33.4 W for a peak average intensity of 20.7 W/mm^2. The damage test was repeated for three pulse durations (2.4 ps, 1.0 ps, 0.3 ps) on the same spot on the window for peak intensities up to 680 MW/mm^2. No change in linearity of transmitted power, fluorescence of the window, or damage marks left on the window were observed.

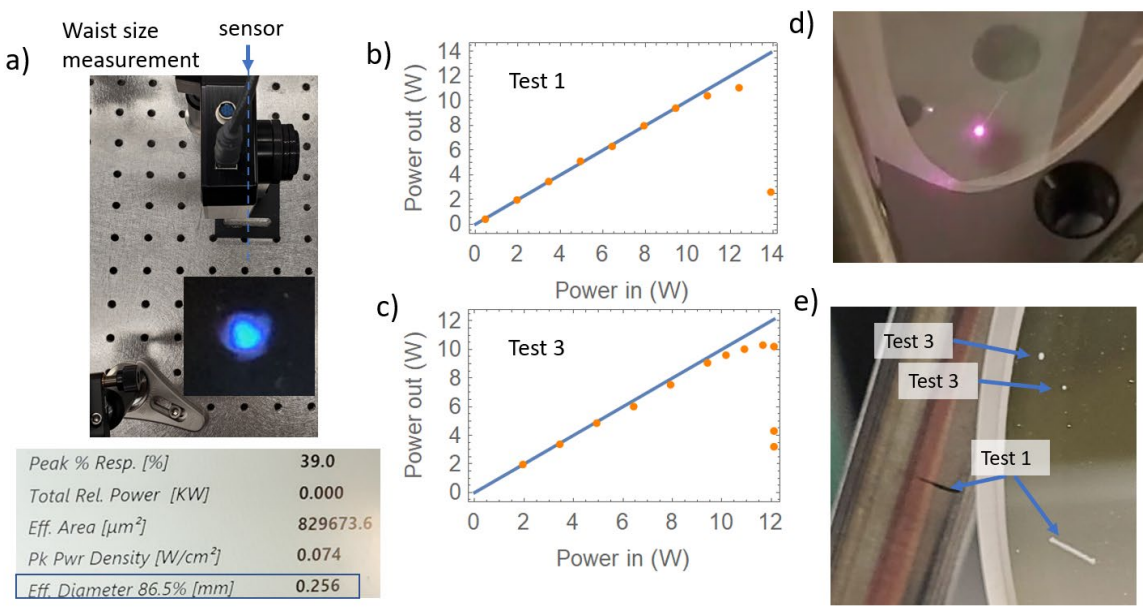

Figure 10. Vacuum window laser damage test. a) A D4 sigma beam size of 0.26 mm was measured with a camera placed at the window position (near the focus). b-c) Measured window transmitted laser power vs input laser power. d) Image of the test spot for the first test. e) Visible damage after the 3 tests.

For the second set of tests, the window was placed at the focus where a laser waist of 0.13 mm was measured with a camera after significant attenuation with the laser controls and a polarizer. For these tests, the repetition rate was 100 kHz and pulse duration was set to 1.0 ps. For tests on three different regions of the window, significant decreases in transmittance through the window were observed as the power was

increased above 10 W average incident power. This indicates that the intensity threshold for damage of the vacuum window is 380 W/mm^2 (3.5 GW/mm^2 peak intensity). The damage appeared localized to the surface of the incident side of the window (see Figure 9e) which is consistent with a previous study showing that defects in coatings are the dominant cause of damage from picosecond pulses (Laurence, 2016).

Since the operating laser waist incident on the vacuum window during operation is expected to be 1 mm, the average power that would be needed to reach the damage threshold would be 380 W. This is 9.5 times the maximum average power produced by the Carbide laser so vacuum window damage should not be a concern during operation.

6 Summary

Slicing via laser heater shaping on a picosecond level can replace femtosecond slicing via the slotted foil method and can be employed at the high repetition rates of LCLS-II. Of the various methods compared for femtosecond slicing, pulse stacking via a Michelson interferometer offers the simplest and most effective method for controlling the FEL pulse duration to as low as 10 fs with the standard 1 kA, 100 pC beam and 5 fs with a 700 A, 20 pC beam. Laser shaping with picosecond pulses will also be important for seeding the microbunching instability to form short, high current pulses for ESASE as well as head and tail horn suppression in lieu of collimators at high rate a (Cesar, et al., 2021). Damage studies of the transport through the vacuum window show that the Carbide laser can be used at its full power of 40 W and still be an order of magnitude from the vacuum window damage threshold.

7 Acknowledgements

We thank Sasha Gilevich for helpful discussions and help with the Carbide laser, and Nicole and Yuantao for the LCLS-II Gaussian beam injector simulation beams used in the simulations of this note.

8 References

Borland, M. (2000). *ELEGANT: A flexible SDDS-compliant code for accelerator simulation.* Argonne, IL: ANL. doi:10.2172/761286

Cesar, D., Anakru, A., Carbajo, S., Duris, J., Franz, P., Li, S., . . . Marinelli, A. (2021). Electron beam shaping via laser heater temporal shaping. *arXiv:2106.04022*. Retrieved from https://arxiv.org/abs/2106.04022

Ding, Y., Behrens, C., Coffee, R., Decker, F.-J., Emma, P., Field, C., . . . Turner, J. (2015). Generating femtosecond X-ray pulses using an emittance-spoiling foil in free-electron lasers. *Appl. Phys. Lett.*, 107, 191104. doi:10.1063/1.4935429

Gilevich, S., Alverson, S., Carbajo, S., Droste, S., Edstrom, S., Fry, A., . . . Zhou, F. (2020). The LCLS-II Photo-Injector Drive Laser System. *CLEO 2020 Proceedings*, 19995144. Retrieved from https://ieeexplore.ieee.org/abstract/document/9192625

Glebov, L. (2008). *Volume Bragg Gratings in PTR glass - New Optical Elements for Laser Design.* OSA. doi:10.1364/FIO.2008.SThA4

Greco, V., Molesini, G., & Quercioli, F. (1995). Accurate polarization interferometer. *Review of Scientific Instruments*, 66, 3729. doi:10.1063/1.1145429

Huang, Z., Brachmann, A., Decker, F.-J., Ding, Y., Dowell, D., Emma, P., . . . D. (2010). Measurements of the linac coherent light source laser heater and its impact on the x-ray free-electron laser performance. *Phys. Rev. ST Accel. Beams*, 13, 020703.

Jacobsson, B. (2008). *Spectral control of lasers and optical parametric oscillators with volume Bragg gratings.* Stockholm, Sweden: Department of Applied Physics, KTH – Royal Institute of Technology. Retrieved from https://www.diva-portal.org/smash/get/diva2:13465/FULLTEXT01.pdf

Laurence, T. (2016). *Picosecond Laser Interactions with Materials: Mechanisms, Material Lifetime, and Performance Optimization.* LLNL. Retrieved from https://ldrd-annual.llnl.gov/archives/ldrd-annual-2016/lasers/picosecond

(2014). *LCLS-II Conceptual Design Report.* SLAC.

Marinelli, A., Coffee, R., Vetter, S., Hering, P., West, G., Gilevich, S., . . . Z. (2016). Optical Shaping of X-Ray Free-Electron Lasers. *Phys. Rev. Lett.*, 116, 254801. doi:10.1103/PhysRevLett.116.254801

Neveu, N., & Ding, Y. (2021). *LCLS-II Standard Configuration with a Gaussian-Profile Injector Laser.* SLAC. Retrieved from https://portal.slac.stanford.edu/sites/lcls_public/lcls_ii/acc_phy/technotes_lib/LCLS-II-TN-21-01.pdf

OptiGrate. (2020). *Transmitting Bragg Gratings (TBG) Product Brief.* Retrieved from https://optigrate.com/product_briefs/TBG-1.pdf

Reiche, S. (1999). GENESIS 1.3: a fully 3D time-dependent FEL simulation code. *NIMA*, Volume 429, Issues 1–3, June 1999, Pages 243-248. doi:10.1016/S0168-9002(99)00114-X