

Fault-tolerance analysis for an injector of continuous-wave free-electron lasers

Zhen Zhang¹,* Yuantao Ding, and Feng Zhou

SLAC National Accelerator Laboratory, Menlo Park, California 94025, USA



(Received 26 June 2025; accepted 3 September 2025; published 26 September 2025)

High-repetition-rate free-electron laser (FEL) facilities such as LCLS-II rely on continuous-wave (cw) injectors to deliver high-quality electron beams to multiple undulator lines simultaneously. Ensuring injector reliability under nonideal conditions is critical for sustained user operations. In this work, we investigate the fault tolerance of a representative cw photoinjector by analyzing its performance under various degraded scenarios, including reduced gun energy, radio-frequency (rf) cavity failures, solenoid field constraints, and early-stage energy gain requirements. Using multiobjective optimization with a genetic algorithm framework (Xopt), we assess achievable beam quality—focusing on transverse emittance and bunch length—across a range of fault conditions. Our results reveal that while certain failures (e.g., buncher off) significantly impact emittance, coordinated adjustments to laser parameters, rf phases, and cavity gradients can still partially recover performance. The injector demonstrates resilience to a variety of perturbations, confirming its robustness for high-demand cw FEL operations.

DOI: [10.1103/zz4w-tvxt](https://doi.org/10.1103/zz4w-tvxt)

I. INTRODUCTION

Continuous-wave (cw) free-electron laser (FEL) facilities, such as LCLS-II [1], SHINE [2], and the planned cw upgrade of the European XFEL [3], place stringent demands on the injector to deliver high-brightness, low-emittance electron beams with exceptional stability and reliability requirements [4–6]. Candidate electron sources for cw FELs include very-high-frequency (VHF) normal-conducting (NC) radio-frequency (rf) guns [7], superconducting rf (SRF) guns [8], direct-current (dc) guns [9], and hybrid dc—SRF designs [10].

In large-scale cw FEL facilities, it is common to operate multiple undulator lines simultaneously, serving a diverse range of user experiments. As a result, maintaining continuous injector operation is critical not only to maximize machine uptime but also to ensure that various scientific users can make full use of beam time. However, in realistic operating environments, deviations from ideal settings are inevitable due to technical limitations or hardware faults. For example, for VHF-gun-based photoinjectors, the rf gun may suffer gradient degradation due to hardware limitations or operational instabilities; the solenoid strength may be constrained for emittance optimization in order to suppress

the dark current transmission, and individual SRF cavities may experience complete/partial failures. These disruptions can potentially compromise beam quality and disrupt FEL operation. Therefore, it is essential to assess whether the injector can continue to support FEL lasing, even with reduced performance or partial functionality. This study investigates the robustness and fault tolerance of a VHF-gun-based cw injector under various nonideal conditions, with a particular focus on whether acceptable beam quality can still be achieved to support FEL lasing, especially for soft x-ray or other operational modes that are less demanding in terms of emittance. By systematically analyzing the effects of key component constraints and outages through detailed simulations, we identify the operational margins and performance trade-offs necessary to maintain reliable injector operation. These results provide critical guidance for machine protection, fault response strategies, and operational flexibility in next-generation FEL facilities.

This work was originally motivated by operational experiences during the LCLS-II injector commissioning. Several failure or constraint conditions studied in this paper were directly encountered during early operation. For example, the gun gradient was reduced to mitigate dark current, and the strength of the first solenoid was limited for the same reason. An unexpected field emission in the fourth cavity of the injector cryomodule reduced its maximum energy gain to approximately 5 MeV, while the fifth cavity operated below 10 MeV—resulting in a total energy shortfall of about 15 MeV compared to the design values. In this scenario, additional energy had to be extracted from the first two cavities to maintain overall injector performance. Despite

*Contact author: zzhang@slac.stanford.edu

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

these challenges, we successfully delivered electron beams with reasonably good emittance and achieved satisfying FEL output for users.

The experimental results of recent LCLS-II injector operation will provide a critical benchmark for the studies discussed here. As a separate paper is currently in preparation to present the experimental progress, this work is intentionally limited to simulation-based analysis for several practical reasons. First, the machine configurations during the relevant fault scenarios varied in beam charge, laser profile, and other settings, making systematic comparisons within a unified simulation framework impractical. Second, limited machine development time—driven primarily by the urgent demand for user experiments—precluded dedicated experimental studies of each fault scenario in the short term. Nonetheless, such experimental studies remain essential for long-term operational planning at facilities such as LCLS-II, and once enough data is collected in the future, a full benchmark will be carried out. In contrast, simulation-based exploration enables consistent evaluation of all fault scenarios starting from the same baseline configuration, thereby providing meaningful and comparable results. The simulation tools employed here are well benchmarked, and the insights and mitigation strategies derived are expected to provide valuable guidance for both commissioning and operational decision making.

The paper is organized as follows. We begin by introducing the LCLS-II injector layout and presenting a baseline solution obtained through multiobjective optimization using a genetic algorithm. This optimized configuration serves as a reference for evaluating performance under various degraded conditions. We then examine the injector's response to several realistic fault scenarios, including gun energy reduction, rf cavity failures, solenoid field constraints, and energy gain requirements in the first few accelerating cavities. For each case, dedicated optimizations are performed to identify strategies that mitigate

performance degradation. The resulting trends highlight the injector's flexibility and resilience, offering practical insights into maintaining beam quality under operational constraints.

II. REPRESENTATIVE INJECTOR LAYOUTS FOR CW FELS

cw FEL facilities typically adopt a superconducting linac fed by a high-brightness injector optimized for stable operation and low emittance. In this study, we focus on the injector layout of the LCLS-II facility [4] as a representative example, shown in Fig. 1. The design features a VHF NC APEX-like rf gun [7] as the electron source, followed by a NC buncher with two cavities for velocity bunching [11] and two solenoids for beam envelope control and emittance compensation [12,13]. The beam is then accelerated by a cryomodule containing eight 1.3-GHz superconducting 9-cell cavities [14], bringing the energy to approximately 100 MeV. Figure 1 shows the layout of the LCLS-II injector and the field profiles of the rf cavities and solenoids used in the simulation. For clarity, we refer to the eight cryomodule cavities as CAV1–CAV8.

Our simulations are based on the physical layout of the LCLS-II injector, incorporating the positions of all solenoids and rf structures. However, the optimization is not restricted to replicate the nominal operating field strengths. Instead, both the magnetic and rf field parameters—as well as the temporal and transverse properties of the drive laser—are treated as tunable variables within a multiobjective optimization that aims to simultaneously minimize transverse emittance and bunch length at the injector exit. While our simulations are based solely on the LCLS-II injector, it is worth noting that other cw FEL facilities, such as SHINE, use a similar layout but with one key difference [5]: SHINE includes an additional single superconducting cavity placed between the buncher and the cryomodule. This extra cavity introduces a longer drift space that facilitates further




FIG. 1. Top: layout of the LCLS-II cw injector, including the VHF gun, buncher cavity, two solenoids (SOL1 and SOL2), and a cryomodule containing eight superconducting rf cavities (CAV1–CAV8). Bottom: axial magnetic (blue line) and electric (green line) field profiles of the solenoids and rf cavities used in the LCLS-II injector simulation.

TABLE I. Parameters used in the multiobjective optimization of the LCLS-II injector layout. The table lists fixed values, discrete sets of values, or parameter ranges considered during the optimization.

Parameter	Value/range	Unit
Beam charge	100	pC
Temporal profile	Gaussian shape	...
Laser rms pulse length	[3, 10]	ps
Transverse profile	1- σ -cut Gaussian	...
Laser spot size (σ)	[0.1, 1]	mm
Intrinsic (thermal) emittance	1.0	$\mu\text{m}/\text{mm}$
Gun on-crest energy gain	750/650/550	keV
Gun gradient	20/17.33/14.67	MV/m
Solenoid SOL1	[0.03, 0.08]	T
Buncher on-crest energy gain	220	keV
Buncher gradient	2	MV/m
Buncher phase ^a	[-120, 20]	°
Solenoid SOL2	[0.01, 0.05]	T
CAV1 gradient ^b	[0, 30]	MV/m
CAV1 phase	[-120, 20]	°
CAV2 gradient	[0, 30]	MV/m
CAV2 phase	[-120, 20]	°
CAV3–8 phase	0	°
CAV3–5 gradient	24.84	MV/m
CAV6–8 gradient	30.54	MV/m

^arf cavity phase is the relative phase to on-crest acceleration, which is defined as 0°.

^bCavity gradients refer to the peak fields used in ASTRA simulations. The corresponding average gradient is obtained by dividing the peak gradient by 1.905.

emittance compensation and allows for the inclusion of a third solenoid for additional envelope control.

Table I summarizes the beamline and laser parameters involved, which include the rf phases and gradients of the gun, buncher, and accelerating cavities; the strengths of the solenoid magnets; and the laser pulse duration and spot size. All simulations in this study are performed for a 100-pC bunch charge, corresponding to the nominal operating point of LCLS-II. While flattop temporal laser profiles are often favored in design studies for their potential to improve beam quality, our study here adopts a Gaussian-shaped temporal profile to reflect the actual operating conditions at LCLS-II [15]. The drive laser is modeled with a truncated Gaussian transverse distribution (cut at 1σ). Parameters are either varied over a defined range or selected from discrete sets, while certain values are held fixed to reduce dimensionality and focus on the optimization. For all simulations presented in the subsequent sections, the intrinsic (thermal) emittance is fixed at 1.0 $\mu\text{m}/\text{mm}$ rms [16]. Note all the emittance values in the following optimization results refer to the 100% root-mean-square (rms) emittance.

Specifically, while we examine the impact of different gun gradients on beam quality by considering three discrete gun energy values (750, 650, and 550 keV), the gun gradient is

not treated as a free variable during the optimization for each case. Instead, it is preset to match the desired on-crest energy gain—20, 17.33, and 14.67 MV/m, respectively. The gun phase is kept close to the on-crest value, as the laser pulse duration is much shorter than the rf period, making the beam largely insensitive to small phase deviations. The buncher gradient is also fixed at 2 MV/m, resulting in an on-crest energy gain of approximately 220 keV. In contrast, the buncher phase is treated as a free variable to control compression through velocity bunching. For the cryomodule cavities, the gradients listed in Table I and all the following related figures represent the peak accelerating field gradients used in the simulations. The corresponding average gradients can be obtained by dividing the peak values by 1.905. Considering the effective electrical length of each cavity (1.038 m), these gradients translate directly into the total effective accelerating voltage per cavity. The phases of the downstream cavities (CAV3–CAV8) are fixed on crest, and their gradients are not varied in the optimization, as they have minimal influence on bunch length and emittance. This strategic balance of fixed and tunable parameters enables efficient and physically meaningful exploration of the injector’s performance space.

The beam dynamics of the injector, including space-charge effects, are simulated using ASTRA [17]. During injector optimizations, we employ 10^4 macroparticles per run, which provides reliable estimates of beam emittance and bunch length while keeping the optimization process computationally practical. For example, in the baseline configuration at a 1 mm bunch length, increasing the number of macroparticles from 10^4 to 10^5 and refining the space-charge grid changes the normalized emittance only slightly—from 0.32 to 0.33 μm —while the bunch length remains essentially unchanged.

The parameter optimizations in this study are performed using the Xopt framework [18], which incorporates the widely used NSGA-II multiobjective genetic algorithm [19]. NSGA-II is an evolutionary algorithm designed to efficiently explore high-dimensional search spaces, simultaneously handling multiple objectives and constraints to produce a set of Pareto-optimal solutions—configurations where improving one objective necessarily degrades another. This approach has been successfully applied to accelerator optimization problems in numerous previous works [20–26]. In our implementation, each optimization typically uses a population size of 256 per generation, selected to balance the number of decision variables with available computational resources, and is run for at least 500 generations to ensure convergence. While the stochastic nature of the algorithm and the limited computational resources mean that absolute global optima cannot be guaranteed, the solutions obtained represent well converged, high-performance configurations that are sufficient to reveal the performance trade-offs and feasible operating points for the various fault scenarios considered here.

III. IMPACT OF GUN ENERGY DEGRADATION ON INJECTOR PERFORMANCE

In this section, we present the optimization results for varying gun energies (gradient) of 750 keV (20 MV/m), 650 keV (17.33 MV/m), and 550 keV (14.67 MV/m), focusing on how gun gradient degradation affects injector performance. Each of these gun energy values was optimized to evaluate the achievable beam quality as the gun energy decreases.

Figure 2 illustrates the Pareto fronts for the three gun energies, showing the trade-off between normalized transverse emittance and rms bunch length at the injector exit. As expected, higher gun energies lead to lower emittance for the same bunch length, demonstrating the reduced impact of space-charge forces during early beam acceleration. The 750 keV case yields the best performance, achieving $0.32 \mu\text{m}$ emittances with rms bunch lengths under 1 mm. In contrast, the 550 keV case shows considerable degradation on the emittance ($0.45 \mu\text{m}$) at the same bunch length. These findings underscore the importance of gun energy in shaping the injector's performance and highlight the advantage of maintaining a high gun voltage to achieve ultralow-emittance beams.

An additional feature of interest appears in the Pareto front for both the 750 and 650 keV gun energy cases, where the trade-off curve exhibits a distinct kink, separating it into two visibly different segments. This nonsmooth transition likely reflects a shift in the dominant compression mechanism or optimization strategy. On the left side of the kink, the solutions favor lower emittance with relatively longer bunch lengths, suggesting moderate compression at the buncher and a more conservative use of downstream cavities. On the other side, solutions achieve shorter bunch lengths at the cost of increased emittance, indicating a stronger role of the buncher in the compression process. This behavior may result from the optimizer navigating a




FIG. 2. Pareto fronts showing the trade-off between normalized transverse emittance and rms bunch length at the injector exit for three gun energies: 750, 650, and 550 keV.

nonconvex performance landscape with multiple local optima or switching between competing beam dynamics regimes. The kink thus serves as a signature of the injector's operational flexibility and highlights the complex trade-offs that emerge when tuning rf and laser parameters to maintain performance.

Figure 3 presents the optimal settings for gun energies of 750, 650, and 550 keV, with the bunch length fixed at 1 mm to select points from the Pareto front. As the gun energy decreases, the optimization reveals a consistent trend toward larger laser spot sizes, which help mitigate emittance growth by reducing space-charge effects in the low-energy regime. Concurrently, the strength of the first solenoid (SOL1) is also reduced to avoid excessive focusing and preserve beam quality. At 650 keV, the optimal laser pulse length is notably shorter than at both 750 and 550 keV, likely reflecting two distinct optimization strategies associated with the Pareto kink discussed earlier.




FIG. 3. Comparison of optimal injector settings at different gun energies (750, 650, and 550 keV) for the LCLS-II injector. All parameters are selected from the Pareto front of the multi-objective optimization, corresponding to a target bunch length of 1 mm.




FIG. 4. Beam evolution and field profiles for the optimized baseline injector configuration at a gun energy of 750 keV. Top: transverse beam size and rms bunch length; middle: normalized transverse emittance and beam energy; bottom: rf and solenoid field profiles along the beamline. This case serves as the reference for subsequent fault-tolerance studies.

Despite changes in gun energy, the phase of CAV1 remains close to on crest, while CAV2 operates near the zero-crossing phase. This phasing strategy indicates that moderate longitudinal compression occurs in and after the buncher, followed by further compression in CAV2.

Figure 4 shows the beam evolution and field profiles for the optimized configuration at a gun energy of 750 keV, which serves as the reference case for the fault-tolerance studies presented in the following sections. The top panel illustrates the transverse beam size and rms bunch length along the injector. The beam size is well controlled throughout the beamline, with a minimum occurring after CAV3 where beam energy is accelerated to 18 MeV. Two-stage longitudinal compression—moderate compression after the buncher followed by main compression in CAV2—yields a final bunch length of approximately 1 mm. The middle panel shows the normalized transverse emittance and beam energy evolution, with the emittance preserved below $0.32 \mu\text{m}$ and the beam accelerated to nearly 100 MeV. The bottom panel displays the rf and solenoid field profiles that govern the beam dynamics. This optimized configuration serves as the nominal operating point against which degraded scenarios are compared in the fault analysis.

IV. INJECTOR PERFORMANCE UNDER KEY RF COMPONENT FAILURES

To evaluate the injector's fault tolerance, we investigate its optimized performance under individual rf component failures. Specifically, we consider three critical scenarios but one at a time: the buncher cavity turned off, CAV1 off, and CAV2 off. These represent realistic failure modes in cw operation, as each component plays a key role in initial compression and energy gain. In the cases of CAV1 and CAV2 failure, the downstream cavity CAV3 is included in the optimization, with its phase and gradient varied within the same parameter ranges of the cavity CAV2 or CAV1.

Figure 5 shows the Pareto fronts for each of these three fault scenarios in comparison with the baseline configuration. Among them, the buncher-off case exhibits the most severe performance degradation, with the normalized transverse emittance at 1 mm bunch length increasing from 0.32 to $0.60 \mu\text{m}$. In contrast, the CAV2-off case shows minimal degradation, with the emittance rising modestly to approximately $0.40 \mu\text{m}$. The CAV1-off case yields an intermediate impact, with the emittance reaching roughly $0.45 \mu\text{m}$. These trends can be understood by examining the optimized settings in Fig. 6 and the beam evolution in Fig. 7.

When the buncher is disabled, the beam undergoes a long low-energy drift from the gun to CAV1, during which significant longitudinal expansion occurs due to space-charge forces. Compression is subsequently applied in both CAV1 and CAV2, requiring strongly off-crest rf phases and a larger initial laser spot size to manage emittance growth. Despite these compensations, the final emittance increases substantially, reflecting the difficulty of achieving effective compression at low energy.

In contrast, when CAV2 is turned off, the required compression is shifted upstream to the buncher and CAV1. This configuration avoids long low-energy drifts




FIG. 5. Pareto fronts showing the trade-off between normalized transverse emittance and rms bunch length at the injector exit for three rf component failures: buncher off, CAV1 off, and CAV2 off.




FIG. 6. Comparison of optimal injector settings at different rf component failures (buncher off, CAV1 off, and CAV2 off) for the LCLS-II injector. All parameters are selected from the Pareto front of the multiobjective optimization, corresponding to a target bunch length of 1 mm.

and reduces the need for aggressive downstream compression, resulting in only a modest emittance increase. In the case of CAV1 failure, the beam drifts from the buncher to CAV2, which becomes the first active cavity. The buncher phase is optimized to impart an energy chirp that partially offsets space-charge-induced distortions during the drift. The beam is then chirped and compressed in CAV2 and reaches the final target bunch length in CAV3.

These results highlight an important design principle: in scenarios involving cavity failures, the first active cavity encountered by the beam in the CM01 module should operate with a relatively weak gradient (as seen with CAV1 in the buncher-off case and CAV2 in the CAV1-off case). This approach enables the introduction of a controlled energy chirp without rapidly increasing the beam energy, preserving flexibility in the longitudinal phase space for downstream compression. Such adaptive strategies are essential for maintaining beam quality in the presence of rf component degradation. We emphasize that this




FIG. 7. Evolution of normalized transverse emittance and rms bunch length along the beamline for three rf component failure scenarios. Top: buncher off; middle: CAV1 off; bottom: CAV2 off. In each case, the configuration is selected from the corresponding Pareto front with a target bunch length of 1 mm.

conclusion is specific to VHF-gun-based injectors with the LCLS-II layout; however, similar fault-tolerance analyses can be carried out for other gun types or injector layouts.

V. INJECTOR PERFORMANCE WITH SOLENOID STRENGTH LIMITS FROM DARK CURRENT CONSTRAINTS

In cw photoinjectors, effective dark current management is essential for maintaining safe and stable operation, particularly when the rf gun produces a significant amount of dark current. Under such conditions, we observed from the practical operation of the LCLS-II injector that the strength of the first solenoid (SOL1) is subject to a practical lower limit to ensure that the dark current is intercepted by the collimator located just downstream of the buncher. A sufficiently strong magnetic field is adopted to overfocus the dark current into the collimator, preventing it from reaching high-energy sections such as the laser heater or downstream cryomodules, where it can cause unwanted activation or radiation damage. This constraint, however, is not necessary when the dark current level is low, in which




FIG. 8. Pareto fronts showing the trade-off between normalized transverse emittance and rms bunch length at the injector exit for cases where the minimum SOL1 strength is enforced to be 0.8%, 1.8%, 2.4%, and 3.7% above the optimal value used in the baseline configuration.

case SOL1 may be set to lower values to favor the best emittance compensation. Although the current dark current level has been significantly reduced—e.g., through over-insertion of the cathode plug [27]—it remains valuable to examine the limitations imposed by solenoid strength in scenarios where dark current suppression is a concern.

To assess the injector's performance under such constraints, we perform multiobjective optimizations while enforcing increasingly strict lower bounds on SOL1: 0.8%, 1.8%, 2.4%, and 3.7% above the optimal value used in the baseline configuration. The resulting Pareto fronts are shown in Fig. 8. At a fixed bunch length of 1 mm, the emittance penalty is small for the 0.8% case, indicating that the system can compensate for this mild constraint by adjusting other parameters such as the laser settings, the second solenoid, or rf phases. However, at shorter bunch lengths (< 0.6 mm), the same constraint produces a noticeable emittance increase, as reflected in the broader spread of the Pareto front. As the SOL1 constraint tightens further, the emittance at 1 mm bunch length increases to 0.40 μm (1.8%), 0.44 μm (2.4%), and 0.49 μm (3.7%), compared to 0.32 μm in the baseline.

Figure 9 presents the beam evolution and field profiles using the 2.4% stronger SOL1 configuration as an example. The elevated SOL1 field causes a strong initial focusing effect just after the gun, which significantly reduces the beam size early in the injector but increases it at the entrance of the cryomodule. Since the standing-wave rf structures (the buncher and cryomodule cavities) also exert transverse focusing, the downstream components—including the buncher phase, the phases of the cryomodule cavities, and the SOL2 field—must be carefully tuned to avoid overfocusing, especially in the low-energy region. In this case, excessive focusing occurs in CAV3, where the beam energy has reached approximately 15 MeV. Compression is distributed across two regions: a weak




FIG. 9. Beam evolution and field profiles for the optimized injector configuration with the minimum SOL1 strength constrained to be 2.4% higher than the optimal value in the baseline case. Top: transverse beam size and rms bunch length; middle: normalized transverse emittance and beam energy; bottom: rf and solenoid field profiles along the beamline.

initial compression after the buncher, followed by the main compression in CAV2, with the bunch length remaining nearly unchanged across CAV1. The strong transverse focusing due to 2.4% stronger SOL1 strength enhances space-charge-induced energy chirp, which rapidly cancels the rf-induced energy chirp and causes compression to stall shortly after buncher cavity. As a result of the optimization for this scenario, the CAV1 is set near on-crest phase to accelerate the beam to about 6.5 MeV. To enable further compression, the CAV2 phase is shifted to -98.7° , placing it at a strong chirping and deeper deceleration phase. As a result, the beam energy decreases from 6.5 to about 5 MeV across CAV2, enabling renewed compression and allowing emittance compensation to complete before the beam is frozen by the high-energy gain in CAV3. Despite the solenoid constraint, these coordinated adjustments preserve reasonable beam quality at the target bunch length, demonstrating the injector's tuning flexibility even under restricted SOL1 conditions.

VI. INJECTOR PERFORMANCE WITH ENHANCED ENERGY GAIN IN EARLY CAVITIES

In this section, we investigate the injector performance under the constraint of increased energy gain in the first two




FIG. 10. Pareto fronts showing the trade-off between normalized transverse emittance and rms bunch length at the injector exit for cases where the combined energy gain from CAV1 and CAV2 is constrained to exceed 12, 16, 20, and 28 MeV. The baseline configuration has a total energy gain of 5.1 MeV from these two cavities.

accelerating cavities, CAV1 and CAV2. This scenario is motivated by the commissioning experience at the LCLS-II injector, where field emission in the injector cryomodule cavities (mainly in CAV4) led to gradient degradation and reduced energy gain in these cavities (CAV3 to CAV5). To compensate for such losses and still achieve the required beam energy for the laser heater and subsequent beamline transport, it may become necessary to extract more energy earlier in the injector (CAV1 and CAV2). In the baseline configuration, CAV1 operates at low gradient and CAV2 is set near the zero-crossing phase—both choices optimized primarily for emittance preservation. While this setup is effective when downstream cavities can provide sufficient acceleration, it may no longer be viable under degraded conditions. Here, we explore how enforcing higher energy gain in CAV1 and CAV2 impacts the beam quality and assess the trade-offs required to maintain injector performance under these operational constraints.

Figure 10 presents the Pareto fronts for injector configurations where the combined energy gain from CAV1 and CAV2 is constrained to exceed 12, 16, 20, and 28 MeV, respectively. In the baseline configuration, the total energy gain from these two cavities is only 3.4 MeV. As expected, increasing the early-cavity energy gain generally degrades beam quality, leading to higher transverse emittance for a given bunch length on the Pareto fronts. However, this effect is not severe: at a fixed bunch length of 1 mm, the normalized emittance increases from 0.32 μm in the baseline case to 0.39 μm for 12 MeV, 0.43 μm for 20 MeV, and 0.47 μm for 28 MeV. Note that 28 MeV is close to the maximum achievable energy gain from CAV1 and CAV2 (32.7 MeV) in our simulations. These results reveal that the first two cavities provide a substantial energy headroom of roughly 25 MeV, with only modest degradation in beam quality.




FIG. 11. Comparison of optimal injector settings for different energy gain requirements from CAV1 and CAV2 (12, 16, 20, and 28 MeV) in the LCLS-II injector. All parameters are selected from the Pareto front of the multiobjective optimization, corresponding to a target bunch length of 1 mm.

Figure 11 shows the optimized parameter settings for these different energy gain requirements, while Fig. 12 presents the corresponding bunch length and beam kinetic energy evolution along the injector. In the baseline configuration, CAV1 and CAV2 together provide only 3.4 MeV, with most of the compression occurring in CAV2. As the required energy gain increases, strong compression in CAV2 becomes less feasible, shifting the compression burden upstream to the buncher and CAV1—at the cost of increased emittance. Compression in the buncher is influenced by both the space-charge-induced chirp before the cavity and its rf phase. Accordingly, increases in early energy gain are accompanied by adjustments in laser pulse duration and buncher phase, with CAV1 also shifted off crest to provide additional compression in the 12, 20, and 28 MeV cases.

Interestingly, for the 16 MeV case, the optimizer identifies a solution where most of the compression occurs in the buncher rather than in CAV1 or CAV2, resulting in an evolution pattern distinct from the 12 and 20 MeV cases.




FIG. 12. Bunch length and beam kinetic energy evolution along the beamline (up to 6 m) for five cases: baseline and configurations requiring total CAV1 and CAV2 energy gains of 12, 16, 20, and 28 MeV.

A natural idea is to check whether higher energy requirements can be obtained starting from this 16 MeV solution by simply increasing CAV2’s amplitude. Indeed, if CAV2’s amplitude is raised to 28.9 MV/m, giving a total of 20 MeV from CAV1 and CAV2, the emittance is $0.44 \mu\text{m}$ —very close to the optimized 20 MeV case ($0.43 \mu\text{m}$). Though this configuration is not selected by the optimizer, the small difference from the optimum is well within simulation uncertainty. This indicates that for a given energy gain requirement, multiple configurations can yield comparably good emittance. The optimized solutions reported in the Pareto fronts and Figs. 11 and 12 therefore represent references and estimates of feasible performance rather than unique “best” configurations. In practice, the choice of working point would also be guided by additional factors such as beam stability, halo formation, and practical constraints on the drive laser parameters.

VII. INJECTOR OPTIMIZATION UNDER COMBINED CONSTRAINTS

To evaluate the injector’s performance under multiple simultaneous operational limitations, we conduct a combined-constraint optimization incorporating three fault conditions: reduced gun energy to 650 keV, a 1.8% increase in SOL1 strength to enforce dark current containment, and a minimum total energy gain of 14 MeV from CAV1 and CAV2 to account for potential downstream cavities




FIG. 13. Pareto fronts showing the trade-off between normalized transverse emittance and rms bunch length at the injector exit for combined constraints.

(CAV3–CAV8) gradient degradation. These conditions represent a realistic commissioning or fault scenario, where several hardware limitations may coexist.

Figure 13 presents the Pareto front for this combined scenario, compared with the baseline result. Across the full bunch length range, the combined-constraint curve shows a consistent increase in normalized transverse emittance relative to the baseline, reflecting the compounded impact of all three constraints. At a target bunch length of 1 mm, the optimized emittance rises from $0.32 \mu\text{m}$ (baseline) to approximately $0.58 \mu\text{m}$ under combined limitations. The overall compression capability is also reduced, as indicated by the less steep front at shorter bunch lengths. This result highlights the injector’s tuning margin under realistic multifault conditions and demonstrates the importance of coordinated adjustments across multiple components to preserve beam quality.

VIII. DISCUSSION AND SUMMARY

This study presents a comprehensive analysis of the fault tolerance of the LCLS-II injector by investigating its performance under a range of nonideal operational conditions. Using multiobjective optimization with a genetic algorithm, we explored how the injector responds to degradations in key components—including reduced gun energy, rf cavity failures, solenoid field constraints, and required early-stage energy gain—while aiming to preserve beam quality, particularly transverse emittance and bunch length.

Our baseline optimization establishes a reference configuration achieving a normalized emittance below $0.32 \mu\text{m}$ at 1 mm bunch length for a gun energy of 750 keV. Sensitivity to gun energy degradation was evident: as the gun energy decreased to 650 and 550 keV, the achievable emittance increased due to stronger space-charge effects at lower initial beam energy. Cavity failure scenarios revealed that turning off individual components—such as the buncher, CAV1, or

CAV2—can be compensated through careful reconfiguration of compression stages, though with varying levels of emittance penalty. Notably, the buncher-off case had the most detrimental impact, while the CAV2-off configuration remained relatively close to baseline performance.

Dark current management introduced practical lower limits on SOL1 strength. We found that modest increases in SOL1 (e.g., +0.8%) could be compensated with negligible emittance penalty at the 1 mm bunch length, but larger increases led to progressively higher emittance due to enhanced transverse focusing and disrupted emittance compensation. Likewise, enforcing minimum energy gain in CAV1 and CAV2 altered the rf phasing strategy and required a shift in compression distribution. The optimization results indicate that the first two cavities offer an energy headroom of approximately 25 MeV, with only moderate impact on beam quality.

Finally, we examined a combined-constraint scenario involving reduced gun energy (650 keV), increased SOL1 strength (+1.8%), and a minimum CAV1 + CAV2 energy gain of 14 MeV. The resulting Pareto front showed a measurable performance degradation that provides guidance about the injector's ability to support the FEL program under multiple simultaneous limitations with reduced beam quality.

Overall, these results highlight the robustness and flexibility of the VHF-gun-based cw photoinjector design when subjected to realistic fault and operational constraints. By identifying the compensation strategies and performance trade-offs associated with different failure modes, this study provides a valuable foundation for both commissioning support and future operational decision making in high-repetition-rate FEL facilities driven by VHF guns.

ACKNOWLEDGMENTS

This work was supported by U.S. Department of Energy Contract No. DE-AC02-76SF00515.

DATA AVAILABILITY

The data that support the findings of this study are openly available in Zenodo [28].

- [1] T. Raubenheimer *et al.*, LCLS-II final design report, Technical Report No. LCLSII-1.1-DR-0251-R0, SLAC National Accelerator Laboratory, 2015.
- [2] Z. Zhu, Z. Zhao, D. Wang, Z. Liu, R. Li, L. Yin, and Z. H. Yang, SCLF: An 8-GeV CW SCRF linac-based x-ray FEL facility in Shanghai, in *Proceedings of the FEL2017, Santa Fe, NM* (JACoW, Geneva, Switzerland, 2017), pp. 20–25.
- [3] D. Kostin and J. Sekutowicz, Progress towards continuous wave operation of the SRF linac at DESY, in *Superconductivity and Particle Accelerators 2018* (SPIE, Washington, 2019), Vol. 11054, pp. 34–42, [10.1117/12.2524952](https://doi.org/10.1117/12.2524952).
- [4] F. Zhou, D. H. Dowell, P. Emma, J. Schmerge, C. Mitchell, and F. Sannibale, LCLS-II injector physics design and beam tuning, in *Proceedings of the IPAC-2017* (JACoW, Geneva, Switzerland, 2017).
- [5] Z. Jiang, H. Qian, M. Zhang, Q. Gu, D. Gu, Z. Wang, X. Li, X. Wu, Z. Liu, X. Li *et al.*, Design and status of SHINE injector, in *Proceedings of the 15th International Particle Accelerator Conference (IPAC-2024)* (JACoW, Geneva, Switzerland, 2024), pp. 164–166.
- [6] D. Bazyl, Y. Chen, M. Dohlus, and T. Limberg, CW operation of the European XFEL: SC-gun injector optimization, S2E calculations and SASE performance, [arXiv:2111.01756](https://arxiv.org/abs/2111.01756).
- [7] F. Sannibale, D. Filippetto, C. F. Papadopoulos, J. Staples, R. Wells, B. Bailey, K. Baptiste, J. Corlett, C. Cork, S. De Santis *et al.*, Advanced photoinjector experiment photogun commissioning results, *Phys. Rev. ST Accel. Beams* **15**, 103501 (2012).
- [8] E. Vogel, J. Sekutowicz, S. Barbanotti, I. Hartl, K. Jensch, D. Klinke, D. Kostin, W.-D. Moeller, M. Schmoekel, S. Sievers *et al.*, SRF gun development at DESY, in *Proceedings of the 29th Linear Accelerator Conference (LINAC-2018)* (JACoW, Geneva, Switzerland, 2018), pp. 105–108.
- [9] C. Gulliford, A. Bartnik, I. Bazarov, L. Cultrera, J. Dobbins, B. Dunham, F. Gonzalez, S. Karkare, H. Lee, H. Li *et al.*, Demonstration of low emittance in the Cornell energy recovery linac injector prototype, *Phys. Rev. ST Accel. Beams* **16**, 073401 (2013).
- [10] Y. Q. Liu, M. Chen, S. Huang, L. Lin, K. X. Liu, S. W. Quan, F. Wang, S. Zhao *et al.*, Engineering design of low-emittance DC-SRF photocathode injector, in *Proceedings of the 39th Free Electron Laser Conference (FEL2019)* (JACoW, Geneva, Switzerland, 2019).
- [11] L. Serafini, M. Ferrario, C. Pellegrini, M. Cornacchia, S. Chattopadhyay, and I. Lindau, Velocity bunching in photoinjectors, *AIP Conf. Proc.* **581**, 87 (2001).
- [12] B. E. Carlsten, New photoelectric injector design for the Los Alamos National Laboratory XUV FEL accelerator, *Nucl. Instrum. Methods Phys. Res., Sect. A* **285**, 313 (1989).
- [13] L. Serafini and J. B. Rosenzweig, Envelope analysis of intense relativistic quasilaminar beams in rf photoinjectors: A theory of emittance compensation, *Phys. Rev. E* **55**, 7565 (1997).
- [14] T. Peterson, T. Arkan, C. Ginsburg, Y. He, J. Kaluzny, M. McGee, and Y. Orlov, LCLS-II 1.3 GHz cryomodule design—modified TESLA-style cryomodule for cw operation, in *Proceedings of the SRF'15* (JACoW, Geneva, Switzerland, 2015), pp. 1417–1421.
- [15] F. Zhou, C. Adolphsen, D. Dowell, and R. Xiang, Overview of cw electron guns and LCLS-II RF gun performance, *Front. Phys.* **11**, 1150809 (2023).
- [16] D. H. Dowell and J. F. Schmerge, Quantum efficiency and thermal emittance of metal photocathodes, *Phys. Rev. ST Accel. Beams* **12**, 074201 (2009).
- [17] K. Floettmann, ASTRA-A space charge tracking algorithm, DESY Technical Report, 1997, <http://www.desy.de/~mpyflo>.
- [18] R. Roussel, A. Edelen, A. Bartnik, and C. Mayes, Xopt: A simplified framework for optimization of accelerator problems using advanced algorithms, in *Proceedings of the*

- IPAC-2023* (JACoW, Geneva, Switzerland, 2023), Vol. 14, pp. 4796–4799.
- [19] K. Deb, A. Pratap, S. Agarwal, and T. A. M. T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II, *IEEE Trans. Evol. Comput.* **6**, 182 (2002).
 - [20] I. V. Bazarov and C. K. Sinclair, Multivariate optimization of a high brightness dc gun photoinjector, *Phys. Rev. ST Accel. Beams* **8**, 034202 (2005).
 - [21] R. Hajima and R. Nagai, Multiparameter optimization of an ERL injector, *Nucl. Instrum. Methods Phys. Res., Sect. A* **557**, 103 (2006).
 - [22] A. Hofler, B. Terzić, M. Kramer, A. Zvezdin, V. Morozov, Y. Roblin, F. Lin, and C. Jarvis, Innovative applications of genetic algorithms to problems in accelerator physics, *Phys. Rev. ST Accel. Beams* **16**, 010101 (2013).
 - [23] C. F. Papadopoulos, D. Filippettopresenter, F. Sannibale, P. Emma, T. Raubenheimer, J. F. Schmerge, L. Wang, and F. Zhou, RF injector beam dynamics optimization for LCLS-II, in *Proceedings of the 5th International Particle Accelerator Conference (IPAC'14), Dresden, Germany* (JACoW, Geneva, Switzerland, 2014), pp. 1974–1976, <https://proceedings.jacow.org/IPAC2014/doi/JACoW-IPAC2014-WEPRO015.html>.
 - [24] H. Chen, L.-M. Zheng, B. Gao, Z.-Z. Li, Y.-C. Du, R.-K. Li, W.-H. Huang, C.-X. Tang, D. Gu, Z. Qi *et al.*, Beam dynamics optimization of very-high-frequency gun photoinjector, *Nucl. Sci. Tech.* **33**, 116 (2022).
 - [25] Z. Zhu, D. Gu, J. Yan, Z. Wang, H. Yang, M. Zhang, H. Deng, and Q. Gu, Inhibition of current-spike formation based on longitudinal phase space manipulation for high-repetition-rate x-ray FEL, *Nucl. Instrum. Methods Phys. Res., Sect. A* **1026**, 166172 (2022).
 - [26] N. Neveu, T. H. Chang, P. Franz, S. Hudson, and J. Larson, Comparison of multiobjective optimization methods for the LCLS-II photoinjector, *Comput. Phys. Commun.* **283**, 108566 (2023).
 - [27] X.-H. Wang, G. Shu, H.-J. Qian, X.-D. Li, Z.-P. Liu, Z.-G. Jiang, H. Meng, C.-C. Xing, Q. Zhou, and H.-X. Deng, Experimental demonstration of dark current mitigation by an over-inserted plug in a normal conducting very-high-frequency gun, *Phys. Rev. Accel. Beams* **28**, 043401 (2025).
 - [28] Z. Zhang, Data for “Fault tolerance analysis for an injector of continuous-wave free-electron lasers”, Zenodo, 2025, [10.5281/zenodo.17057251](https://zenodo.org/record/17057251).