

Three-dipole kicker injection scheme for the Advanced Light Source upgrade accumulator ring

M. Ehrlichman^{✉,*}, T. Hellert, S. C. Leemann, G. Penn[✉], C. Steier[✉], C. Sun, M. Venturini, and D. Wang[✉]

Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA

 (Received 16 June 2021; accepted 29 November 2021; published 17 December 2021)

The Advanced Light Source Upgrade will implement on axis single-train swap-out injection employing an accumulator between the booster and storage rings. The accumulator ring (AR) design is a twelve period triple-bend achromat that will be installed along the inner circumference of the storage-ring tunnel. A nonconventional injection scheme will be utilized for top-off off axis injection from the booster into the AR meant to accommodate a large ~ 300 nm emittance beam into a vacuum-chamber with a limiting horizontal aperture radius as small as 8 mm. The scheme incorporates three dipole kickers distributed over three sectors, with two kickers perturbing the stored beam and the third affecting both the stored and the injected beam trajectories. This paper describes this “3DK” injection scheme and how it fits the AR’s particular requirements. We describe the design and optimization process, and how we evaluated its fitness as a solution for booster-to-accumulator ring injection.

DOI: [10.1103/PhysRevAccelBeams.24.120702](https://doi.org/10.1103/PhysRevAccelBeams.24.120702)

I. INTRODUCTION

The Advanced Light Source-Upgrade (ALS-U) is the upgrade of the Advanced Light Source to a fourth generation diffraction-limited soft x-ray light source [1]. Achieving ALS-U’s high-brightness goal requires strong focusing-magnet gradients; the strong gradients necessitate strong chromatic sextupoles and these will shrink the dynamic aperture of the storage ring (SR) to about 0.5 mm radius once lattice imperfections are taken into account.

To inject into such a small dynamic aperture, ALS-U will implement on axis single-train swap-out injection utilizing an accumulator ring (AR) housed along the inner wall of the SR tunnel. See Fig. 1 for an overview of the ALS-U accelerator complex. A small 2 nm-rad natural emittance, much smaller than the approximately 300 nm-rad emittance of the beam delivered by the existing booster, is required to inject into the SR with near-100% efficiency and sufficient margin. In addition to its function as a damping ring, the AR is intended to act as a beam-charge recycler in between swap-outs. The SR average current is 500 mA, distributed evenly over a 284-bunch beam consisting of eleven trains of 26 (or 25) bunches each. The AR is designed to carry a single train at a time; this is

swapped with one of the SR trains once every about half-minute and replenished with top-off injection from the booster before the next swap-out.

The AR design has to fulfill two competing demands on the vacuum-chamber aperture: it should be wide enough to accept the large emittance beam from the booster, with a goal of $\gtrsim 95\%$ injection efficiency, but as narrow as possible to minimize the magnets’ aperture and thus their weight and volume so that the AR can fit in the same tunnel as the SR. Consideration of these demands has guided the choice of injection scheme that recognizes that the AR can tolerate a significant injection transient. Thus, an injection cycle leaves both the stored and injected beams with significant transients, both contained in the dynamic aperture of the machine. Such a technique is also referred to as aperture sharing. Fully exploiting the latitude offered by the latter observation, we have developed a simple and nearly 100% efficient injection scheme, coined “3DK”, that utilizes three dipole kickers distributed over three sectors in combination with a pair of thick and thin pulsed septa.

In short, the first two dipole kickers (“prekickers”) kick only the stored beam and place it on a trajectory that partly compensates the main injection kick placed further downstream, which kicks both the stored and injected particles onto stable (but not closed) trajectories. The features of this scheme include (i) the reduced straight length of the AR is accommodated by distributing the kickers among three straights, (ii) near-100% injection efficiency, (iii) a reduced septum aperture that eases vacuum and magnet engineering in the injection straight, (iv) flexibility to trade between residual oscillations of the stored and injected beams, and

*mike@lbl.gov

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




FIG. 1. The Advanced Light Source Upgrade complex. The three dipole kickers and injection septum of the 3DK booster-to-accumulator injection scheme are highlighted. Every 1.4 s up to four bunches are generated and accelerated to 2 GeV through the existing linac and booster and then injected through the BTA into the new AR to fill or top-off bunches in a single 26- (or 25-) bunch train. Every ~ 30 s this train is swapped with one of the eleven trains circulating in the Storage Ring (SR).

(v) reliance on only conventional power supply and kicker technology.

An important aspect of our study is the demonstration that potential drawbacks resulting from the combination of relatively large stored-beam oscillations and collective effects are manageable.

Other schemes that were considered and ultimately discarded included a nonlinear kicker (NLK), a conventional four-kicker single-straight closed-bump injection scheme, as well as a two-dipole kicker variant of the 3DK scheme. These will be briefly discussed in this paper as well.

Section II is a brief introduction to the ALS-U AR. The layout and design of the 3DK injection scheme is described in Sec. III. For completeness, Sec. IV briefly describes the design of the booster-to-AR transfer line. The robustness of the injection scheme against lattice imperfections is discussed in Sec. V. Impedance and collective effects, as well as multibunch feedback issues are addressed in Sec. VII. Alternatives to the 3DK injection scheme are discussed in the Appendix.

II. ALS-U ACCUMULATOR RING

The AR is an intermediary that needs to have a large enough acceptance to capture near-100% of the relatively

large booster beam, and small enough emittance to inject with near-100% efficiency into the SR. To that end, the existing ALS twelve-sector triple-bend acromat layout meets these needs. Such a layout is well understood, allowing R&D efforts to be focused on the SR and transfer lines. With that, the ALS-U AR is essentially a slightly smaller version of the existing ALS twelve period triple-bend acromat. The AR's basic parameters are compared with those of the existing ALS in Table I. The optics and layout through one arc are shown in Fig. 2. To adapt the ALS layout to the AR layout, the length of the straight sections was shrunk from 9.386 to 8.762 m and the arcs were shrunk from 7.014 to 6.415 m by reducing the magnet spacing. Nominal bend, quadrupole, and sextupole magnet lengths are the same between the ALS and AR. These adjustments bring the total circumference down from 196.805 to 182.122 m which makes the AR fit neatly along the inner tunnel wall. The AR utilizes a nonswept gradient dipole design (picture a geometry similar to a partially opened book), as does the ALS. The beam energy is increased from 1.9 to 2.0 GeV to match that of the ALS-U SR. The AR emittance of 1.8 nm is comparable to that of the ALS.

To preserve egress, the AR will be mounted close to the tunnel ceiling, about 2 m above the floor and about 0.6 m

TABLE I. Parameter list of the ALS-U AR and existing ALS.

	AR	ALS ^a
Beam energy	2.0 GeV	1.9 GeV
Circumference	182.122 m	196.8 m
Tune x/y	16.221/8.328	16.165/9.25
Natural chromaticity x/y	-43/-36	-46.8/-39.6
Momentum compaction	1.04×10^{-3}	0.9×10^{-3}
Emittance	1.8 nm	1.8 nm
Dispersion in straight	11.6 cm	15.0 cm
Charge per bunch	1.15 nC	1.1 nC
Energy spread	8.5×10^{-4}	9.6×10^{-4}
Energy loss per turn	269 keV	228 keV
Damping time $x/y/z$	6.2/8.5/5.2 ms	7.7/8.9/5.0 ms
Harmonic number	304	328
Main rf frequency	500.417 MHz	499.654 MHz
Main rf voltage	1.0 MV	1.2 MV
Synchrotron tune	4.9×10^{-3}	5.4×10^{-3}

^aEmittance and energy loss per turn reported for the ALS are for the lattice without superbends and without insertion devices.

above the plane of the SR. The AR magnet stands are a combination of floor and wall attachments. The elevated installation and limited space inside the tunnel places a premium on reducing the size and weight of the magnets, and so the vacuum chamber dimensions were shrunk, bringing the pole tips closer to beam, achieving the same gradients with less bulky magnets. The arc and straight chambers are round and 14.2 mm in radius. The dipole chambers are elliptical with a 20 mm axis in the horizontal and 7.28 mm axis in the vertical. The limiting aperture is set by the injection septum, which has an 8 mm outward aperture.

The AR-to-SR injection swap is expected to occur approximately every 30 s. A long beam lifetime in the AR is not necessary and so dynamic aperture and injection efficiency are prioritized when optimizing the AR. The AR

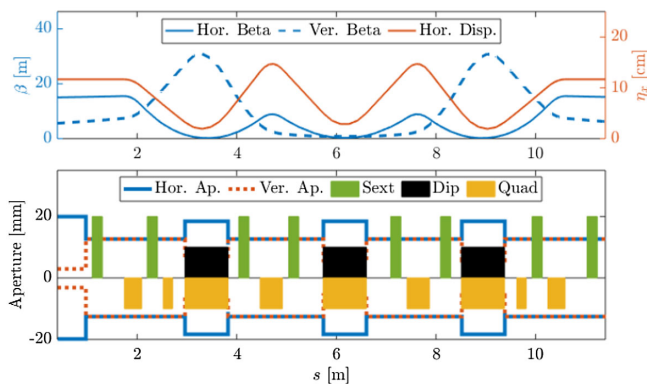


FIG. 2. Lattice and magnet distribution of an arc in the ALS-U AR. The ring is composed of 12 such arcs. Shown are the beta and the dispersion functions (top), the aperture model and the distribution of magnets (bottom).




FIG. 3. The layout diagram of the BTA thick and thin septa in sector 1. The separation between the injected beam and unknicked stored beam is 14 mm and the thickness of the thin septum blade is 1 mm.

dynamic aperture thus extends to the physical apertures. Simulations taking imperfections and correction algorithms into account bring the expected usable dynamic aperture close to linear acceptance.

Transfer lines intersect the AR along three consecutive straights. Sector 12 is the takeoff of the accumulator-to-storage ring (ATS) transfer line; the straight section has to accommodate the dipole kicker and pulsed thin septum required for extraction from the AR into the ATS in addition to the first pre-kicker for AR injection from the booster. Sector 1 houses the pair of booster-to-accumulator (BTA) injection septa, which are depicted in Fig. 3; immediately upstream in the same straight section is the second pre-kicker for AR injection from the booster. The sector 2 straight contains the landing point of the storage ring-to-accumulator (STA) transfer line. It has to accommodate the STA pulsed thin septum, as well as both the dipole kicker for injection from the SR and the main kicker for AR injection from the booster. The close proximity of the three transfer lines and injection or extraction components posed a design challenge, requiring careful consideration of the layout and close interaction between the beam physics and engineering groups.

The arrival of the injected bunches from the booster is timed to replenish the bunches of the depleted train in a top-off manner. In normal ALS-U operations we expect that a train circulating in the SR will lose about 10% of its charge before being swapped out. The beam coming out of the booster is 4 pulses long, with a separation of 3 empty buckets between each pulse. The total charge in the 4 pulses is approximately 0.5 nC, or about 10% of the charge that needs to be replenished, implying that about 10 injection shots will be needed to replenish the circulating train before swapping it back to the SR.

On-axis injection from the SR into the AR is more straightforward than injection from the booster; this and extraction from the AR into the SR are not in the scope of this paper.

III. 3DK INJECTION SCHEME

The AR layout along with injection elements and beam trajectories during an injection cycle are depicted in Fig. 4. At the start of a BTA injection cycle, the stored bunch train is given an outward “prekick” at the upstream end of the sector 12 straight. A second outward prekick is applied at the upstream end of the next straight, which is sector 1. The injection septum is at the downstream end of sector 1. Injected particles make a large excursion through the sector 1 arc. The main injection kicker, located at the downstream end of sector 2, applies an outward kick to both the stored and injected particles, leaving both on trajectories that damp with very little particle loss. The strengths of the 3 kickers and injection septum are adjusted to control the injection transients and are shown in Table II. The kickers’ bumps are confined to one turn and their flat top spans the entire bunch train and are timed to the train head, irrespective of which buckets are being injected into.

The vacuum chamber between the injection point and the second downstream dipole is horizontally widened to accommodate the large excursion of the injected-beam trajectory but no special-aperture magnets are required.

The injected bunches arriving from the pulsed thin septum are offset 14 mm outward from the nominal chamber center, as depicted in Fig. 3. The transverse position of the septum was arrived at by considering both the limitation the septum places on dynamic aperture and the constraint it imposes on the beam-trajectory oscillation

amplitude during the post-injection transient. The distance of the injected beam centroid to the septum sheet was arrived at by considering particle loss on the septum sheet and the amplitude of the injected particles.

The starting point for determining the optimal septum angle is found by manually adjusting the angle of the injected particles to create a zero-crossing at the main injection kicker. The septum angle is then refined along with the strengths of the three kickers by applying a Levenburg-Marquardt minimizer to the amplitudes of the stored and injected transients in action or angle coordinates. The optimum is that which minimizes the action of the centroids while preserving minimum distances from the limiting physical apertures and tolerating no losses out to 3σ of stored beam. Design and optimization of the injection scheme was done using TAO, which is based of the Bmad accelerator code [2].

The parameter space is explored in Table II, which presents the thin septum and kicker strengths for different injection modes, as well as the maximum amplitude and rms of the horizontal injection transient for the stored and injected beams. The injection losses and stored beam survival percentages were calculated using a 9000 particle Gaussian-weighted distribution extending to $4\text{-}\sigma$ in all three planes, tracked for 300 turns. The “on axis injection mode,” as was done at MAX IV [3], is expected to be useful during early commissioning. It will place the injected beam on axis, but the stored beam, if present, will be left with a large

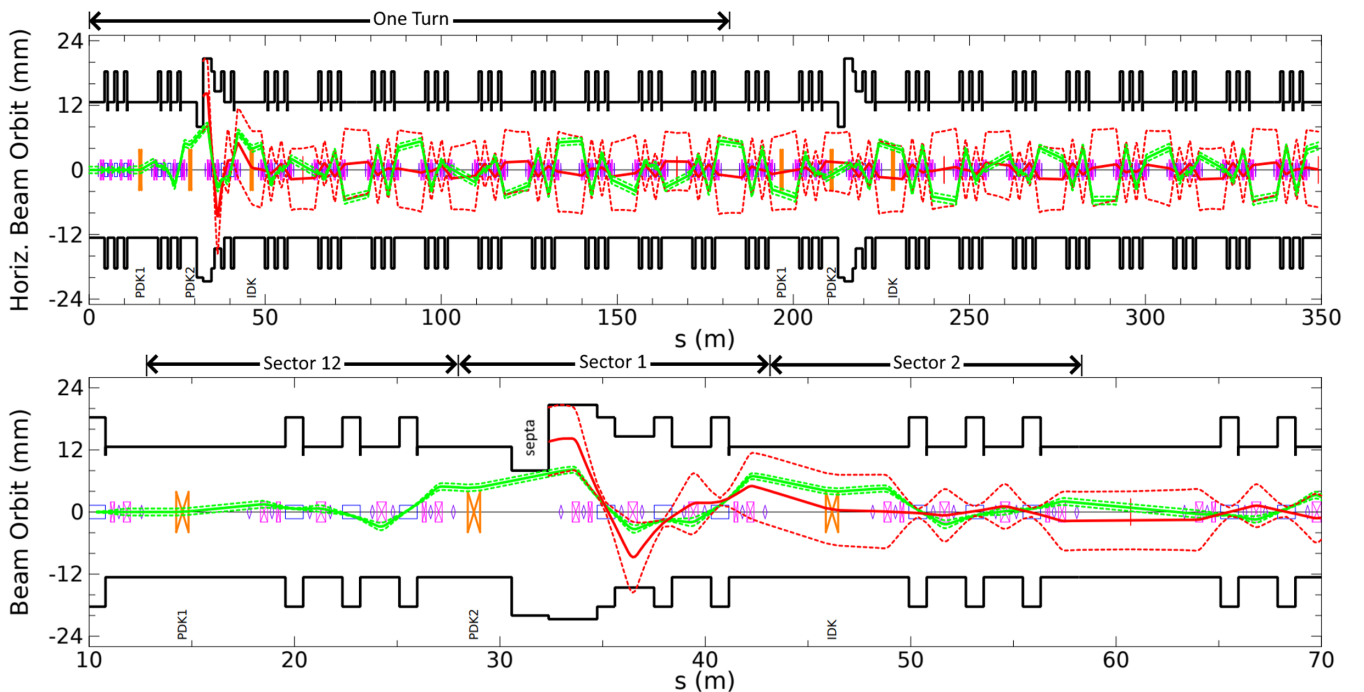


FIG. 4. Trajectories of the stored (green) and injected (red) beams during an injection cycle for the “balanced” Mode A kicker strength settings. Solid black profile indicates horizontal physical aperture. Dashed lines delineate 3σ of the beam size. Kickers are labeled and represented by orange bow-ties. The first pre-kicker is located at the start of sector 12. The second pre-kicker and injection septum are located in sector 1. The main injection kicker in sector 2 kicks both the stored and injected particles.

TABLE II. Injection modes achieved by adjusting kickers and septum strength. Mode A achieves a high-injection efficiency with a small stored beam transient. Mode B reduces the stored beam transient to mitigate short range wake fields, at the expense of greater losses on the BTA-side septum sheet and due to injected particles exceeding the acceptance. Mode C recovers high-injection efficiency by shrinking (mismatching) the horizontal β of the injected beam. Short-range wake fields were not studied for the on axis injection and closed stored-beam bump modes.

		On axis injection	Closed stored-beam bump	Mode A	Mode B	Mode C
Injection element strengths	Thin septum (mrad)	35.82	35.85	35.85	35.72	35.67
	Prekicker 1 (mrad)	0.365	0.489	0.383	0.417	0.417
	Prekicker 2 (mrad)	1.09	0.759	1.04	0.933	0.933
	Injection kicker (mrad)	1.21	0.613	1.11	0.987	0.987
Post-injectioncycle residuals	Stored beam rms (mm)	3.54	0.0	3.00	2.27	2.31
	Stored beam maximum (mm)	6.32	0.0	5.39	4.18	4.18
	Injected beam rms (mm)	0.0	4.58	0.898	1.24	1.22
	Injected beam maximum (mm)	0.0	8.16	1.59	2.21	2.17
Optics at septum exit	β_x (m)	14.82	14.82	14.82	14.82	11.33
	α_x	-0.139	-0.139	-0.139	-0.139	-0.275
	BTA-side septum losses (%)	0.76	0.76	0.80	1.05	0.13
	Capture losses (%)					
	No wakes	0.21	77.7	0.30	1.77	1.30
	With wakes	16.4	6.8	6.9
	With wake mitigations	0.7	1.1	1.0
Stored beam survival (%)	99.9	99.9	99.9	99.9	99.9	

oscillation amplitude that limits imperfection tolerance and generates large short-range wake fields. The “closed stored-beam orbit bump mode” eliminates the stored beam injection transient, but achieves only 22.3% injection efficiency. It is included here to bound the parameter space. The remaining three modes indicate different strategies to optimize performance for small but nonzero residual amplitudes of the stored and injected beams. Mode A, or “balanced mode,” represents a compromise solution between the two previous modes and maintains near-100% injection efficiency while reducing the amplitude of the stored beam oscillations with an injected beam that is matched to the SR. For this mode, the transient of the stored beam is allowed to be larger than that of the injected beam as it has a much smaller beam size than the injected particles. The trajectories in Fig. 4 are based on Mode A.

The parameters of the injection elements are further refined after considering short-range transverse wake fields. Because of the nonzero injection transient, the stored beam generates wakes strong enough to sweep some of the injected particles into the vacuum chamber. This is described in more detail in Sec. VI. When taking short range wakes into account, it is found that a smaller perturbation to the stored beam improves the overall injection efficiency, even though the injected bunch has a larger initial offset and suffers slightly larger losses initially.

Mode B brings the injected beam closer to the septum as it enters from the BTA thus requiring a weaker main injection kick and allowing for a smaller stored beam transient at the expense of slightly reduced injection

efficiency. The injected beam losses to the machine acceptance increase from 0.30% to 1.77%, while the losses due to wake fields decrease from 16.4% to 6.8%. BTA-side septum sheet losses also increase slightly from 0.80% to 1.05%.

Mode C has the same kicker settings as Mode B but decreases β_x to better fit the injected beam into the AR acceptance available after taking the septum sheet into account, similar to the technique applied in Ref. [4]. This differs from Modes A and B where the beam distributions are matched to the ring. Reducing β_x shrinks the injected beam horizontally, reducing losses along the septum sheet. After adjusting the septum strength to recenter the injected beam within the available machine acceptance, the centroid is brought closer to the septum, thereby reducing the injection transient. BTA-side septum sheet losses are reduced to 0.13%, and losses to the acceptance are reduced to 1.30%.

Figure 5 shows the DA of the latter three modes viewed from the septum exit during an injection cycle, along with the injected and stored particle distributions, and also the thin septum. Both the inward and outward profiles of the thin septum in this figure were determined by particle tracking. Comparing Modes B and A, the acceptance appears to move since the acceptance itself relative to the septum is affected by the choice of main injection kick.

In producing the nominal operational mode and its two variants, we conclude that the 3DK scheme has the flexibility to accommodate perturbations and offers various optimization opportunities.

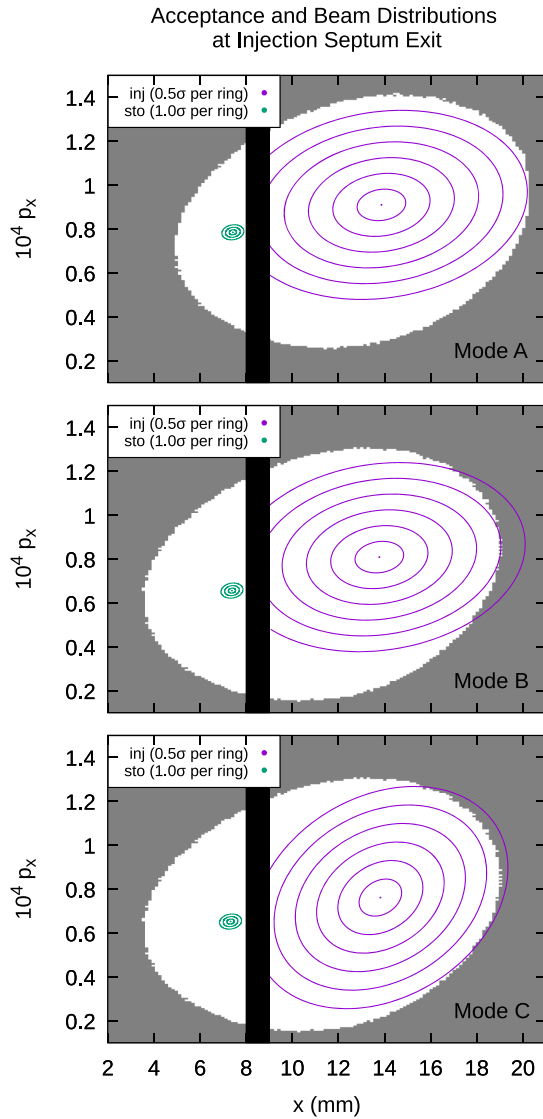


FIG. 5. Injected and stored beam profiles out to 3σ at septum exit during an injection cycle for the three injection modes. The white region is the acceptance in $x-x'$ space at the septum exit. The septum is drawn in black and is 1 mm thick with its inward edge 8 mm from the stored beam reference trajectory. The septum drawing was obtained by tracking. The parameters of the three modes are detailed in Table II.

IV. BOOSTER-TO-ACCUMULATOR (BTA) LAYOUT DESIGN AND OPTICS MATCHING

Transport from the booster to the AR is through the Booster-to-Accumulator (BTA) transfer line. The BTA utilizes about 2/3 of the existing booster-to-ALS storage ring (BTS) transfer line. For early commissioning a slow bending switch will be installed for on-demand steering of the beam into the new BTA branch while ALS can continue normal operation.

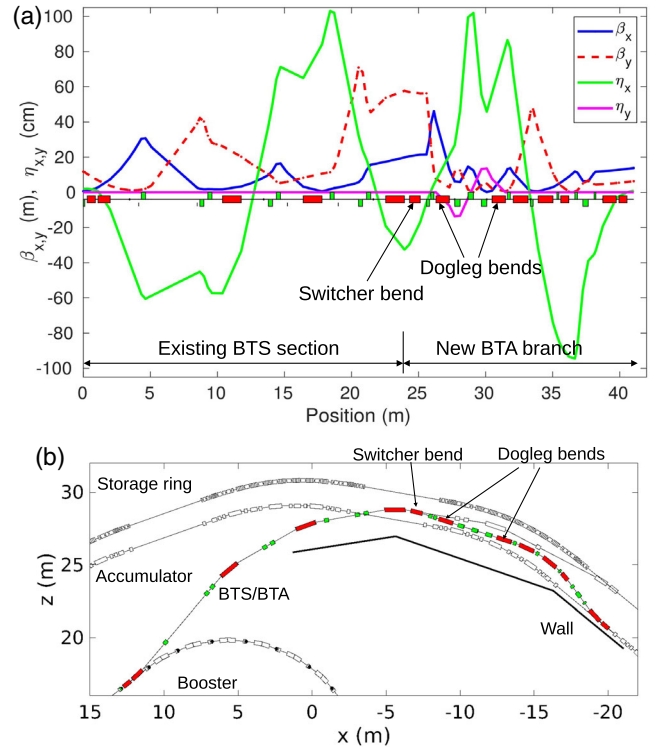


FIG. 6. Twiss functions (a) and layout (b) of the BTA transfer line.

To accommodate the different elevation between booster and AR (see Fig. 1), the BTA includes a vertical two-bend achromatic dogleg [5], where the achromatic condition is met with two quadrupoles between the two dogleg bends to generate the required vertical π phase advance. The BTA terminates into a pair of pulsed thin and thick septa providing a combined 12 deg (horizontal) bending. Since the space is limited, thick and thin septa are located at the far end of the injection straight (straight 1) as shown in Fig. 3 to increase the distance from the last bend of the BTA to the septa and lower the incoming angle into the septa, therefore reducing the required kick angles from septa.

The main design difficulty posed by the BTA is the avoidance of interference between the BTA and other systems' components in the tight and congested injection region while meeting the layout and lattice functions' matching boundary conditions. Linear-optics matching was performed with constraints on beam sizes along the transfer line to minimize vacuum chamber and magnets' aperture. The decoupling of horizontal and vertical planes provided by the dogleg simplifies the matching problem but we found it useful to deploy a Multiobjective Genetic Algorithm [6] to optimize the design. The optics functions and beam-size envelopes of the baseline lattice are shown in Fig. 6. The desired Twiss functions at the end of the transfer line are obtained and the beam sizes are below 4 mm throughout, which satisfies the design requirements.

V. INJECTION EFFICIENCY ANALYSIS AND ROBUSTNESS AGAINST LATTICE ERRORS

The injection efficiency is evaluated using AR and BTA lattices obtained after undergoing a process of simulated commissioning [7].

Optimizing the injection efficiency in the presence of errors requires small adjustments to the nominal kick angles, with the exact value depending on the error realization. In simulations and eventually in operation, the optimization can be accomplished by running parameter scans in a neighborhood of the kick-angle nominal settings.

While one would have the freedom to adjust all three DKs independently, in the simulations we used a simplified setup in which both pre-DKs are varied by the same amount, thus resulting in a two-dimensional optimization problem. We found that this setup gave satisfying results while significantly reducing the computation time compared to a full three-parameter scan.

In a preliminary study, we took a first cut at the parameter space by delimiting the region that exhibits no or minimal losses in both injected and stored beam in the absence of errors: see the top images in Fig. 7. In the simulations, the injected and stored bunches are represented with 1000 particles and tracked for 1000 turns.

When errors are added to the lattice, the best combination of kicker settings may vary slightly depending on the specific random error realization and correction. For each of 100 random lattice-error realizations, the kicker settings are scanned in the vicinity of their nominal, ideal values.

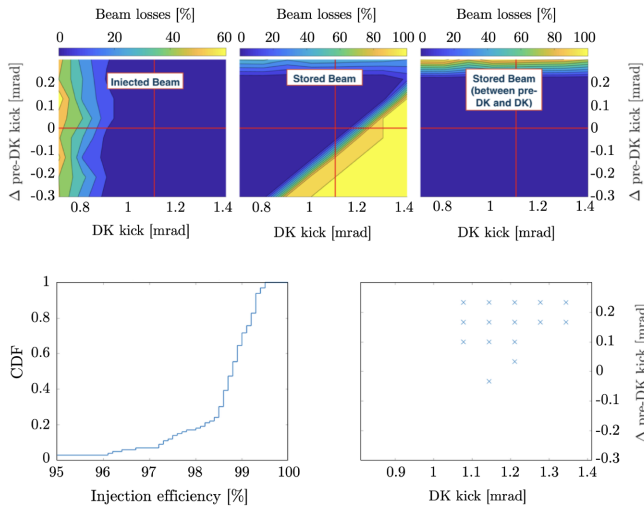


FIG. 7. Top: particle losses in the injected (left) and stored (center and right) beam vs kickers' angles for the ideal lattice. The pre-kickers' kick (vertical axis) is reported relative to Mode A settings (red lines). Bottom: for each one of 100 lattice-error realizations, a DK scan is carried out on a grid to identify the setting yielding the highest injection efficiency. The ensemble of the best DK settings is reported in the right figure. The left figure is the injection efficiency cumulative density function (CDF).

TABLE III. Imperfections assigned to injection elements in the commissioning simulations. The error distributions are Gaussian and truncated at 2σ . The stability errors are relative to the set point, while the roll error is absolute.

Error type	Distribution width
Septum stability	2.5×10^{-4}
Dipole kicker stability	1×10^{-3}
Dipole kicker roll	4 mrad

For each error realization, the kickers' settings that induce no stored-beam particle losses were selected and among these the one kickers' setting that gives the highest injection efficiency was identified. The set of those best 100 settings (one per lattice) is shown in the bottom-right image of Fig. 7 (only 15 distinct data points appear, as a single kickers' setting will in general represent the optimum for several lattice-error realizations.) The bottom-left picture in Fig. 7 is the cumulative density function of this set, showing that about half of the lattices have injection efficiency larger than 98.5%.

We then proceeded by refining the simulations to include a more complete set of errors and perturbations. Specifically, to the AR lattice errors we added: septa and kickers shot-to-shot variations (Table III), nonuniformity of the kicker-pulse temporal profile, and septum leakage fields.

The nonuniformity of the kicker pulse is simulated by assigning a linear slope to the pulse profile. The slope is defined as the difference between the kicks on the first and last bunch of the bunch train relative to the nominal kick. The simulation of the leakage fields is done by calculating the kick map associated with the field accounting for both the temporal and spatial field extension, to be applied at every turn before the field has died off, Fig. 8. A result from these more complete simulations is shown in Fig. 9 demonstrating the sensitivity of injection efficiency to the sloping of the kickers' pulse. As before, for each lattice-




FIG. 8. Temporal and spatial profiles of the thin-septum leakage field as included in the injection efficiency studies. The left plot shows the integrated kick at various horizontal coordinates as a function of turns. Conversely, the right plot shows the integrated kick at various times as a function of a particle horizontal coordinate.

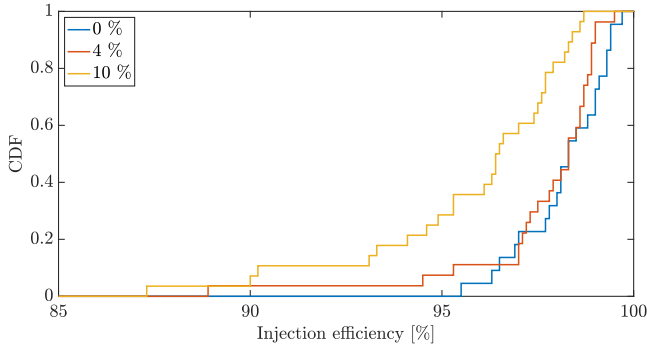


FIG. 9. Sensitivity of the injection efficiency to the nonuniformity of the injection kickers' pulse profile. The three data sets correspond to a 0%, 4%, and 10% sloping of the pulse profile. The simulation is conducted with a complete set of errors including shot-to-shot variation of the kick and septa strengths, as well as septum leakage fields.

error realization, the injection efficiency was determined after performing a two-dimensional scan of the kicker's amplitude to identify the optimum. Among other things, from these results we concluded that a 4% slope would be acceptable.

VI. SINGLE BUNCH (SHORT-RANGE WAKE FIELDS)

During injection both the electrons from the booster and the stored bunch onto which they are added undergo several millimeter transverse oscillations and this has the potential to induce potentially harmful transverse collective effects. The investigation of these effects is the topic of this section. These investigations were conducted using the ELEGANT accelerator code [8].

The short-range transverse wake fields have been calculated for each vacuum chamber element using a detailed model of its geometry. Longitudinal wakes are not included in these results, however, only the horizontal wakes have been observed to have a significant impact on the injection process. The wakes are applied close to the location of the impedance, although simulations are in good agreement with simpler calculations that use beta-function weighting to group together the impact of wake fields in every sector or even at a single location in the ring. The threshold for the Transverse Mode-Coupling Instability is 5.8 nC, about 8 times the nominal charge for a single bunch. Figure 10 shows the beta-weighted wakefields from different components in both horizontal and vertical planes, along with their total. Wake fields are calculated for a Gaussian beam with rms bunch length of 1 mm and serve as pseudo-Green functions in the following beam dynamics simulations.

The resulting transverse oscillations, especially of the stored bunch (because it contains about 90% of the charge), induce short-range wake fields. Despite the fact that the

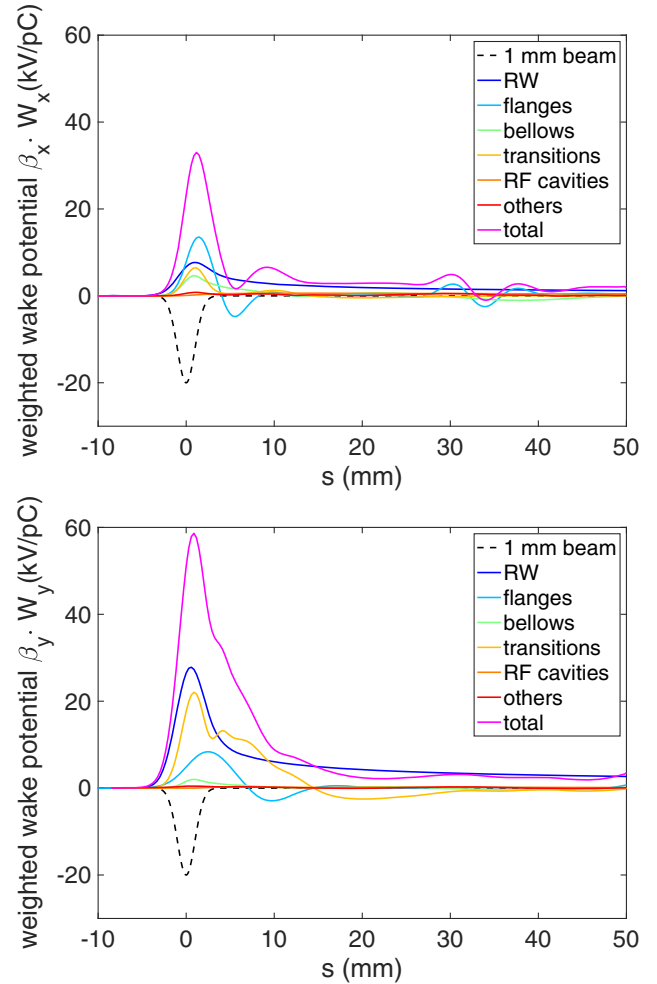


FIG. 10. Transverse short-range wakefields in the AR from individual components and their combined total. The wake potentials are calculated for a Gaussian beam with rms bunch length of 1 mm and serve as pseudo-Green functions in beam dynamics simulations. Left: horizontal plane, right: vertical plane.

total bunch charge is expected to be far from the threshold for instabilities driven by short-range wakes, the wake fields do build up in intensity over a time scale of order 1 ms (roughly 1600 passes around the ring). Simulations show that the wake fields are strong enough to confine the stored bunch and delay phase decoherence, as shown in Fig. 11. At the same time, the more diffuse injected electrons increase in both oscillation amplitude and transverse width.

The 3DK “balanced” configuration (Mode A) was chosen by optimizing for injection efficiency but without considering the impact of wakes. Less than 1% beam losses are predicted in this case. However, when the expected wake fields are included, the injection efficiency drops to 83.6%. As shown in Fig. 12, most of the losses occur between 0.5 and 1 ms after injection, roughly between 1000 and 2000 passes around the ring.

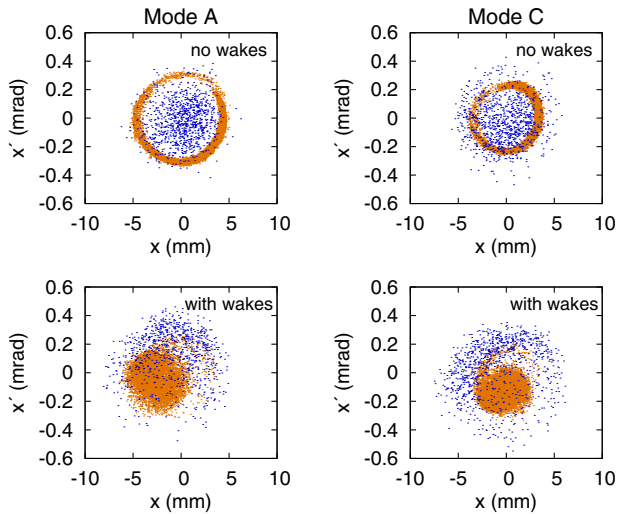


FIG. 11. Horizontal phase space of the stored (orange) and injected (blue) bunches after 2000 turns, without wakes (above) and with nominal wakes (below). The left figures show the results for Mode A, the right for Mode C.

An important aspect of the 3DK scheme is that it is flexible enough to allow for additional adjustments to accommodate the impact of wake fields. An alternate set of all three kicker magnet amplitudes, referred to as “Mode B”, was found that reduced the displacement of the trajectory of the stored beam by 20%, while increasing losses by a modest amount when wake fields are neglected. These additional losses occur in the first 50 passes around the ring. The capture efficiency remains above 99%. For the nominal wake fields the capture efficiency is 93.2%,

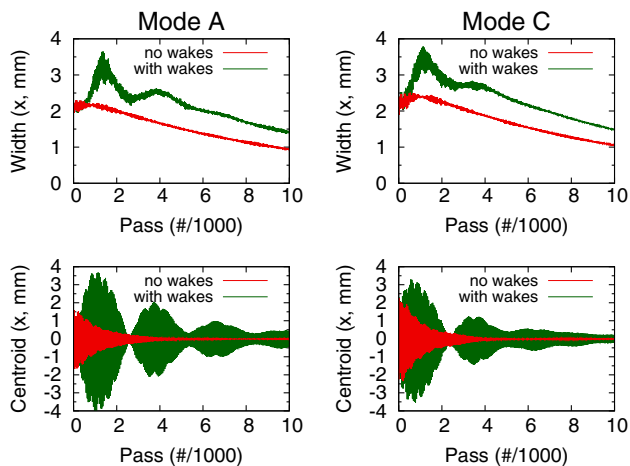


FIG. 12. Comparison of the evolution of the size (top) and centroid motion (bottom) of the injection bunch for the cases without wake fields (red) and with nominal wake fields (green). The left figures show results for the 3DK settings referred to as Mode A (“balanced”), and right figures show Mode C, which is tuned for better performance in the presence of wake fields. The corresponding injection efficiencies for the two cases are 83.6% and 93.1%.

a significant improvement compared to 83.6% for Mode A. In none of these cases are there losses in the stored bunch.

Mode C is similar to Mode B except that the septum field is weaker and the horizontal beta function at the end of the septum is lower than the matched value so as to reduce transfer-line losses on the injection septum. These adjustments also reduce the injected beam transient slightly, though when wake fields are taken into account the capture efficiency is comparable to that of Mode B. The performance of Modes B and C are much more similar to each other than they are to Mode A, and comparisons will focus on the differences between Modes A and C.

Figure 12 shows the evolution of the size and horizontal offset of the injected bunch when short-range wake fields are either included or ignored. Results for Mode A are shown on the left-side plots, and for Mode C on the right-side plots. In Mode A, the width of the injected bunch steadily decays without wake fields, while with the nominal wakes the width of the bunch grows until it peaks after 1800 turns around the ring with an rms of 3.5 mm, up from 2.2 mm. The centroid motion rapidly damps from 1.5 mm without wakes, but with the nominal wakes it increases to 3.5 mm amplitude, followed by envelope oscillations which slowly decay over thousands of passes around the ring. For Mode C, there is similar behavior in the width of the injected bunch but the amplitude of the centroid motion starts out at 2 mm and only grows to 3.0 mm. The centroid motion again has continued envelope oscillations in the presence of wakes, but at a lower amplitude than for Mode A.

For the nominal wake fields, the goal of at least 95% injection efficiency is not met even for Mode C. Thus, other strategies to mitigate the impact of short-range wake fields have been examined. The nominal chromaticity of 1 in both planes can easily be increased and will drive more rapid phase decoherence of the bunches. Both horizontal and vertical chromaticities contribute to this effect at high amplitudes. A small temporal offset of the injected bunch relative to the stored bunch directly reduces the wake field forces experienced by the injected electrons, as the injected particles experience the peak of the wake field for only brief periods during their synchrotron oscillations.

Increasing the chromaticity from 1 to 1.2 increases the capture efficiency to 93.5% for Mode A and 97.7% for Mode C. Similarly, keeping the chromaticity at unity and delaying the injected bunch by 100 ps relative to the stored bunch leads to a capture efficiency of 98.5% for Mode A and 98.9% for Mode C. Applying both changes brings the capture efficiency above 99% in both modes.

To exercise a measure of caution, we considered the possibility that wake fields would be twice as high as expected (“enhanced wakes”). In this case, the capture efficiency drops to 64.8% for Mode A and 78.3% for Mode C without additional mitigations. The evolution of the capture efficiency as a function of number of passes




FIG. 13. Capture efficiency as a function of number of passes around the ring after injection, comparing Mode A, where the kicker settings are optimized without including wake field effects, to Mode C, which uses an alternative set of kicker strengths which is more tolerant to wake fields and also changes the injection beta function. Results are shown with and without wake fields, for enhanced wake fields, and for additional mitigations.

after injection is shown in Fig. 13 for both Modes A and C. The wake field levels are varied among no wakes, nominal wakes, and enhanced wakes. In addition, some examples include using a combined mitigation of setting both chromaticities to 1.2 and shifting the injected bunch by 100 ps. For the enhanced wakes, this brings the capture efficiency to 97.8% for Mode A and 98.3% for Mode C. Creating a temporal offset between the stored and injected bunches is particularly effective at reducing the impact of wake fields. For Mode C with the above mitigations, the capture efficiency is 97.2% when the strength of the wake fields is multiplied by a factor of three. The corresponding efficiency for Mode A is 95.9%.

VII. MULTIBUNCH STABILITY AND INTERFERENCE WITH TFB SYSTEM

The transverse bunch-by-bunch feedback system (TFB) will detect the transient of the stored beam and apply kicks to damp it. Unfortunately, those kicks will antidamp the injected particles since the stored beam and injected particles are out of phase. The injected particles carry relatively little charge and will offset only slightly the center-of-charge that the feedback system detects. To prevent the loss of injected particles as the TFB damps the injection transient, the TFB will be masked for those buckets into which charge is injected. A single injection shot delivers charge into four buckets, spaced four buckets apart, for example, buckets 10, 14, 18, and 22. Masking the buckets significantly diminishes the ability of the TFB to damp multibunch instabilities and so we investigate here the injection transient and TFB interaction using a tracking simulation.

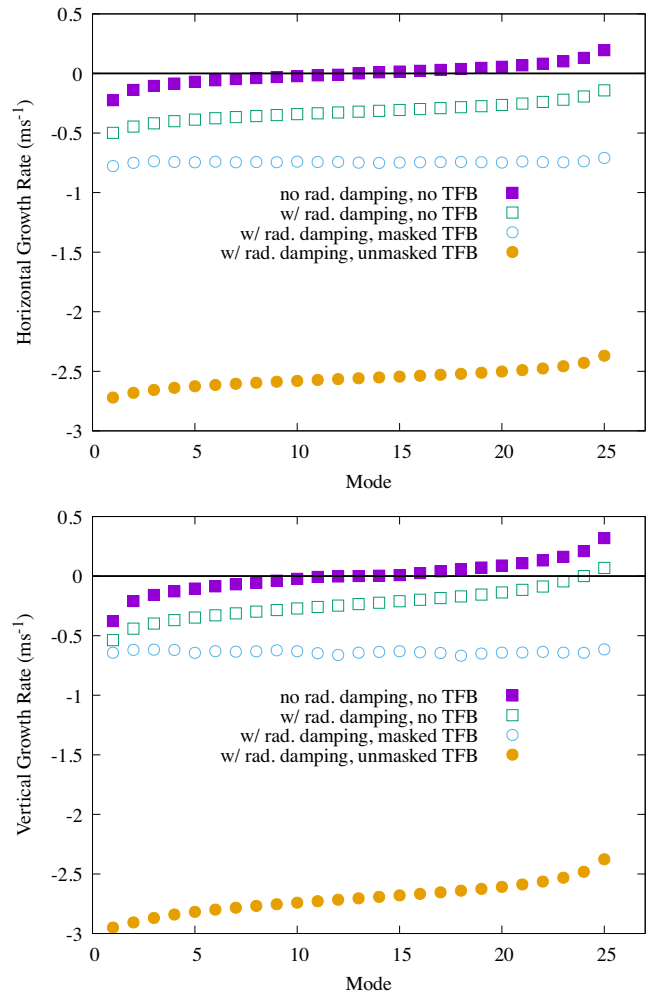


FIG. 14. Per-mode growth rates of a 25 bunch train in the accumulator ring determined by multibunch macroparticle tracking with resistive wall wake fields. During steady-state, all RW modes are safely damped with a total growth rate of about -2.5 ms^{-1} . During the injection transient, the effectiveness of the TFB is significantly diminished by the necessity of masking the four buckets into which charge is injected, though all modes remain damped.

Under nominal conditions, the growth rates of multibunch impedance instabilities in the AR are less than that of the radiation damping and multibunch feedback system. The AR highest order mode (HOM) wakes have been evaluated with theory and tracking simulations. The strongest cavity HOM growth rate is 0.03 ms^{-1} , and the radiation damping rate is 0.16 ms^{-1} . The growth rates of the resistive wall modes are shown in Fig. 14. Except for the two HOMs, all are below the radiation damping time, and the HOMs are well within the capabilities of the TFB.

Following an injection cycle, the stored beam is left with a large transverse offset and the four buckets into which charge is injected are masked from the multibunch feedback. Since the resistive wall modes are near or exceed the

radiation damping rate, a tracking simulation is used to evaluate their stability during the injection transient when the effectiveness of the TFB is diminished.

The simulation models the long-range resistive wall wakes at any position in the ring using a fit of 15 pseudomodes to the $1/\sqrt{t}$ dependence of the long-range wake fields and follows closely the technique described in Ref. [9]. The long-range resistive wall wake is,

$$W(t) = \frac{\sqrt{c}}{\pi b^3} \sqrt{\frac{Z_0}{\pi \sigma}} \frac{L}{\sqrt{t}}, \quad (1)$$

where b is the chamber dimension and σ its conductivity, c and Z_0 are the speed of light and impedance of free space, L is the length of the chamber and t is time since the passage of the source particle [10]. The basis of pseudomodes is

$$W_i(t) = \frac{A_i^2}{S_i^2} \exp\left(-\frac{d_i^2}{S_i^2} t\right), \quad (2)$$

where A_i , S_i , and d_i are fit parameters. Here, S_i is redundant but including it as a fit parameter has been found to help the fitter converge consistently and quickly with simply default parameters. The fitter is Mathematica's `NonlinearModelFit` [11]. The log of the basis is fit to $\log(1/\sqrt{t})$ sampled with a logarithmic distribution of t . The residuals of the fit are less than 0.1% from 1 ns to 44,000 turns, and reliably decay to zero beyond the ends of the fit. The same basis is used for all chambers but is scaled according to the local chamber material and dimension shown in Eq. (1). Resistive wall wakes are represented at 216 locations in the ring.

For each horizontal mode n , each bunch i in the train is seeded with an appropriate offset,

$$x_i = \sqrt{J\beta_x} \cos\left(\frac{\pi}{2} + \frac{2\pi ni}{N_{\text{bunches}}}\right) \quad (3)$$

$$x'_i = \sqrt{\frac{J}{\beta_x}} \left(\alpha_x \cos\left(\frac{\pi}{2} + \frac{2\pi ni}{N_{\text{bunches}}}\right) + \sin\left(\frac{\pi}{2} + \frac{2\pi ni}{N_{\text{bunches}}}\right) \right). \quad (4)$$

and tracked for 5000 turns.

The multibunch tracking simulation is developed using `Bmad` [2], which keeps track of the pseudomodes as they build up and influence the beam turn-by-turn. Radiation damping is implemented by a per-turn decrement to the action of each macroparticle. The multibunch feedback is simulated using a pickup, FIR filter, and kicker.

Each turn the normal mode coordinates of each bunch are recorded and a Fourier transform is taken along the bunch train. For the seeded mode, the growth rate is obtained from the fit of an exponential to the height of the spectral peak versus turn number. This is done similarly

for the vertical modes. The growth rates of each of the 25 modes are shown in Fig. 14. From Fig. 14, we conclude that while masking will significantly reduce the damping obtained from the multibunch feedback system, all resistive wall modes will remain stable by comfortable margin during the transient.

Following an injection cycle, the bunch train is offset with a large uniform transient. However, the study just described used small amplitudes and particular offsets given by Eqs. (3) and (4). To simulate absolute stability following an injection cycle, the beam is given a uniform 5 mm horizontal offset and tracked for 5000 turns. Without masking, after 1702 turns the bunch oscillating with the largest amplitude has an action that is 50% of the initial condition. With four bunches masked, 50% is achieved after 3266 turns. All 15 RW pseudomodes are reliably decaying in strength throughout the simulation.

VIII. CONCLUSION

The three-dipole kicker injection scheme is well fit to the particular constraints and allowances of the ALS-U AR. It reliably injects the relatively large beam coming out of the booster with nearly 100% efficiency using conventional technology. It meets the AR's tight space constraints by distributing the injection kickers across multiple straights. It fits the large beam coming in from the booster into the relatively small vacuum chamber by allowing a transient on the stored beam so as to shrink the transient of the injected beam.

The injection transient has consequences for long- and short-range wake fields, but several techniques have been prepared to effectively mitigate these consequences. The scheme is robust in the presence of lattice and pulsed element imperfections. Both the wake field issues and imperfection tolerances benefit from the ability of the scheme to tune the amplitudes of the stored and injected transients.

In the Appendix, the merits of three alternative injection schemes are compared to those of the 3DK scheme.

ACKNOWLEDGMENTS

The authors would like to thank David Sagan of Cornell University for supporting this work through the development of `Tao` and `Bmad`. This work was supported by the Director, Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

APPENDIX: ALTERNATE INJECTION OPTIONS

1. Nonlinear kicker

Unlike conventional dipole-kicker magnets, a nonlinear kicker (NLK) generates a transverse nonlinear

magnetic-field profile with a maximum located off axis at the injected-beam arrival point and zero in the center. The injected beam is kicked while the stored beam at the center is minimally perturbed. In addition, a single NLK device is sufficient for injection, thus easing issues of space and synchronization. Because of these reasons, several light source facilities have adopted NLK in alternative to more conventional closed-orbit bump schemes [12–15].

Although injection perturbations to the stored beam are not a major concern in the AR (the stored beam does not serve user experiments), NLK injection represents an attractive option because of its compactness and simplicity, while a relaxed tolerance to stored-beam perturbation can be exploited to optimize the design in ways that are not suitable for light-source applications.

The NLK design configuration we considered consisted of eight main conductors symmetrically positioned around the vacuum chamber in the direction of the beam as shown in Fig. 15. The NLK was located in the same lattice position as the main kicker in the 3DK scheme. The conductor placement was optimized using Multiobjective Genetic Algorithm methods [6] in start-to-end macroparticle injection simulations; for a detailed description of the optimization setup and results see Ref. [16]. With an optimized NLK design, injection tracking studies in the presence of lattice errors and including orbit or optics correction

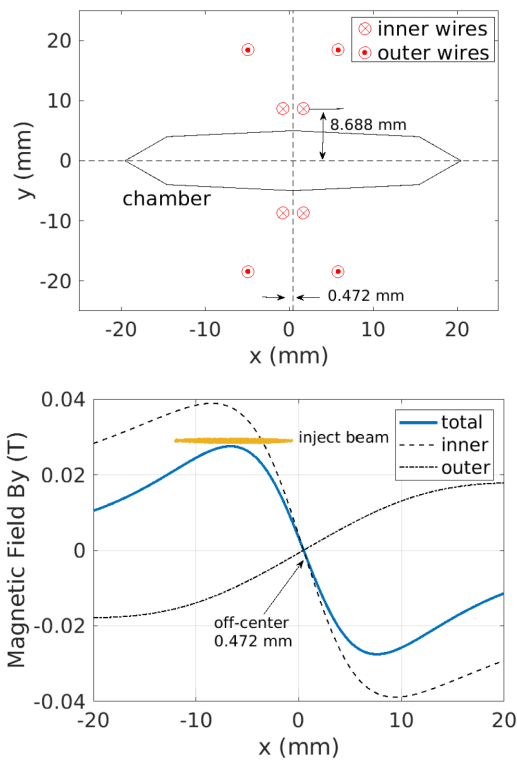


FIG. 15. NLK-design concept. Top: main conductors’ transverse positions (the return conductors are not shown). Bottom: magnetic-field profile in the mid plane.

showed injection efficiency above 95% (not including transverse wakefield effects).

While these simulations were encouraging, an NLK solution for the AR injection was eventually abandoned after the R&D effort initiated at ALS [16] encountered technical difficulty with the prototyping of the ceramic vacuum chamber.

2. Conventional closed orbit bump 4DK injection in one sector

The existing ALS utilizes a conventional four-kicker orbit-bump injection scheme with all four kickers in the same straight section. A similar scheme was explored for the AR with all four injection kickers placed in Sector 2. Potential advantages of this scheme are (i) it is conventional, (ii) residual oscillations of the stored beam following the injection cycle are eliminated, and (iii) it eliminates the need to carefully control the phase advance between the pre-kicker and main injection kicker.

The straights in the AR are 0.624 m shorter than those in the ALS. This shorter space drives up the strength required from the dipole kickers. Additionally, because the orbit bump is closed for the stored beam, the injected particles oscillate with larger amplitude. To avoid loss of injected particles, the injection septum aperture would need to be increased to the nominal vacuum chamber aperture of about 14 mm. This complicates engineering in the injection straight and further increases the necessary dipole kick to about 10 mrad. Such a large dipole kick requires some combination of kicker R&D, a multiturn ramp, or placing the four kickers on a single power supply string. A multiturn ramp would constrain the machine tune, thus negating one potential advantage of this scheme. Placing the four kickers on a single string would make it difficult to trade off between stored beam and injected beam oscillation amplitudes, which is a disadvantage compared to the 3DK scheme. An additional drawback is that the septa would need to be moved upstream (toward the center of the straight) to make room for the four kickers. This has the consequence of forcing a stronger bending requirement in the thick septum.




FIG. 16. Injection efficiency for the conventional 4DK scheme evaluated over a population of 100 lattice-error realizations including orbit and optics corrections. Results are for the septum positioned at 8 mm (left) and 13 mm (right) from the vacuum chamber center. Tracking simulations were done with 10k particles over 10k machine turns.

Even with a 14 mm septum aperture, a conventional four kicker design delivers less than 90% injection efficiency, as shown in Fig. 16. In contrast, the 3DK scheme approaches 100% injection efficiency. The 3DK scheme also compares well in requiring only an 8 mm septum aperture, and its kicker technology requirements are modest: three 0.6 m long ferrite-loaded dipole kickers delivering ~ 1 mrad kick.

3. Two dipole kicker (2DK)

It is possible to implement an injection scheme similar to 3DK while using only two dipole kickers. Coined “2DK” this scheme saves costs by requiring one less kicker. With 3DK, the two pre-kickers allow adjustment of the position and angle of the stored beam at the main injection kicker. With 2DK, only one pre-kicker is employed and its location in the SR is chosen to have a horizontal phase advance from the pre-kicker to the main injection kicker such that the stored beam has a near-zero crossing at the main injection kicker. The pre-kicker is placed in sector 7, which is six sectors upstream of the main injection kicker.

2DK is competitive as far as injection efficiency goes, but the lack of flexibility limits the ability to accommodate difficult-to-predict consequences of imperfections and wake fields. It also constrains the ring tune, as the pre-kicker to main injection kicker phase advance must be precisely set. Generally, the phase advance along a single straight is too little to allow for significant adjustment. It was judged that the versatility of the 3DK scheme justified the cost.

-
- [1] C. Steier *et al.*, in *Proceedings of 9th International Particle Accelerator Conference (IPAC'18), Vancouver, BC, Canada, 2018* (JACoW Publishing, Geneva, Switzerland, 2018), pp. 4134–4137.
- [2] D. Sagan, Computational accelerator physics, *Nucl. Instrum. Methods Phys. Res., Sect. A* **558**, 356 (2006).
- [3] S. Leemann, *Nucl. Instrum. Methods Phys. Res., Sect. A* **693**, 117 (2012).
- [4] A. Streun, SLS booster-to-ring transferline optics for optimum injection efficiency, Technical Report No. SLS-TME-TA-2002-0193, Paul Scherrer Institut, Villigen PSI, Switzerland, 2005.
- [5] C. Sun *et al.*, in *Proceedings of 10th International Particle Accelerator Conference (IPAC'19), Melbourne, Australia, 2019* (JACoW Publishing, Geneva, Switzerland, 2019), pp. 2756–2759.
- [6] K. Deb, A. Pratap, S. Agarwal, and T. Meyerivan, *IEEE Trans. Evol. Comput.* **6**, 182 (2002).
- [7] T. Hellert, P. Amstutz, C. Steier, and M. Venturini, *Phys. Rev. Accel. Beams* **22**, 100702 (2019).
- [8] M. Borland, elegant: A flexible SDDS-compliant code for accelerator simulation, Technical Report No. LS-287, Advanced Photon Source, Argonne, Chicago, 2000.
- [9] S. Wang, M. Billing, S. Poprocki, D. L. Rubin, and D. Sagan, in *Proceedings of International Particle Accelerator Conference (IPAC'17), Copenhagen, Denmark, 2017* (JACoW, Geneva, Switzerland, 2017), pp. 3204–3207.
- [10] A. W. Chao, *Physics of Collective Beam Instabilities in High-Energy Accelerators* (John Wiley & Sons, Inc., New York, 1993).
- [11] W. R. Inc., “Mathematica, Version 12.1,” champaign, IL, 2021.
- [12] S. C. Leemann, *Phys. Rev. ST Accel. Beams* **15**, 050705 (2012).
- [13] H. Rast, T. Atkinson, M. Dirsat, O. Dressler, and P. Kuske, *Conf. Proc. C* **110904**, 3394 (2011), <https://accelconf.web.cern.ch/IPAC2011/papers/thpo024.pdf>.
- [14] L. Liu, X. R. Resende, A. R. D. Rodrigues, and F. H. de Sá, in *Proceedings of International Particle Accelerator Conference (IPAC'16), Busan, Korea, 2016* (JACoW, Geneva, 2016), 10.18429/JACoW-IPAC2016-THPMR011.
- [15] P. Alexandre, R. B. E. Fekih, A. Letrésor, S. Thorau, J. da Silva Castro, F. Bouvet, J. Breunlin, Åke Andersson, and P. F. Tavares, *Nucl. Instrum. Methods Phys. Res., Sect. A* **986**, 164739 (2021).
- [16] C. Sun, P. Amstutz, T. Hellert, S. C. Leemann, C. Steier, C. Swenson, and M. Venturini, *Phys. Rev. Accel. Beams* **23**, 010702 (2020).