USING NORMALIZING FLOWS IN NORMAL-FORM SPACE TO MODEL INTRA-BEAM STRIPPING*

S. Kakkar†, L. Lee, University of Tennessee at Knoxville, Knoxville, United States
N. Evans, A. Hoover, V. Morozov, Oak Ridge National Laboratory, Oak Ridge, United States

Abstract

Intra-beam stripping (IBSt) is a critical beam-loss mechanism in high-intensity H⁻ linacs and presents a significant limitation to increasing beam power. This work presents a computational framework to evaluate IBSt-induced beam losses along the Spallation Neutron Source (SNS) linac for arbitrary bunch distributions. The calculation is based on evaluating the IBSt loss integral using a probability density function (PDF) trained on discrete-particle bunch distributions with normalizing flows. The input distribution is transformed to scaled normal-form coordinates, which improves both normalizing-flow training and Monte Carlo (MC) sampling. The method is benchmarked against simplified analytically solvable Gaussian bunches and then applied to canonical-angular-momentum-dominated (CAM-dominated) beam distributions, which contain strong inter-plane correlations. The results show that the normal-form-coordinate approach improves MC sampling stability and enables efficient IBSt loss calculations for correlated beam distributions.

INTRODUCTION

High-intensity H^- linacs require careful control of beam loss to support reliable operation and future power upgrades. At the SNS, IBSt is one mechanism that can contribute to beam loss when particles within the bunch interact and strip the weakly bound electron from H^- ions. Accurate prediction of IBSt loss requires knowledge of the six-dimensional beam phase-space distribution, since the loss rate depends on both the relative velocity of interacting particles and their spatial overlap. For realistic or correlated beams, this distribution is generally not known analytically, making direct evaluation of the IBSt loss integral challenging.

Previous work developed a Monte Carlo framework for evaluating IBSt losses using analytic Gaussian distributions and machine-learned PDFs trained on PyORBIT tracked phase-space data [1, 2]. In this work, we extend that framework by reformulating both normalizing-flow training and MC sampling in scaled normal-form coordinates. The lab coordinates are transformed to normal-form coordinates, where they are uncorrelated in the rms sense, and then scaled so that each of the transformed coordinates has a dimensionless unit rms spread. This improves the overlap between the sampling distribution and the high-density region of the bunch PDF.

To demonstrate the efficiency of the updated algorithm, we apply it to the challenging case of canonical-angular-momentum-dominated (CAM-dominated) beams. These beams contain strong inter-plane correlations, such as x-y' and y-x', which may make conventional MC sampling in the lab frame inefficient. Therefore, they provide a useful test of the normal-form-coordinate method presented in this paper.

IBST LOSS RATE FORMALISM

We begin with an analytic expression for the IBSt loss rate specified in the bunch rest frame [3]. In this frame, the stripping cross section σ_H describing the interaction probability depends on the magnitude of the relative velocity $\mathbf{u} = \mathbf{v}_1 - \mathbf{v}_2$ between the two interacting particles. The total loss rate can be written as

$$\frac{dN}{dt_{\text{rest}}} = \frac{N^2}{2} \int |\mathbf{u}|\sigma_H(|\mathbf{u}|) f(\mathbf{v}_1, \mathbf{r}_1) f(\mathbf{v}_2, \mathbf{r}_2) \times \delta(\mathbf{r}_1 - \mathbf{r}_2)\, d\Gamma_1\, d\Gamma_2 \tag{1}$$

where f is the phase-space probability density, which is a function of the six phase-space coordinates. The delta function accounts for the fact that IBSt occurs locally, so only particle pairs that overlap at the same position in the bunch contribute to the loss rate.

We convert beam loss per unit time in the rest frame to beam loss per unit length along the linac in the lab frame by means of the following substitutions $dt_{\text{lab}} = \gamma\, dt_{\text{rest}}$ and $ds = \beta c\, dt_{\text{lab}}$ arriving at

$$\frac{dN}{ds} = \frac{1}{\gamma \beta c} \frac{dN}{dt_{\text{rest}}} \tag{2}$$

In the results below, we report the normalized particle number loss rate $(1/N)dN/ds$ as a function of the longitudinal position s along the linac.

NORMAL-FORM-COORDINATE MONTE CARLO METHOD

Coordinate Transformation

In the previous implementation [2], Monte Carlo integration points were sampled directly in the scaled lab-coordinates. This becomes inefficient if the beam distribution is stretched, or strongly correlated. In such cases, a fixed sampling region can include many points where the beam density is low.

The new calculation uses a sequence of coordinate transformations. The integrand functions in Eq. (1) are expressed in terms of the rest-frame coordinates. They are first converted to the lab-frame coordinates. We next apply a transformation to the normal-form coordinates and finally scale them

* This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE).
† nlns1@ornl.gov

so that the resulting distribution has an identity covariance matrix.

The six-dimensional particle coordinate vector can be written as

$$\mathbf{X} = (x, x', y, y', z, \delta)^T, \tag{3}$$

where δ denotes the relative momentum deviation. The lab coordinates are related to the scaled normal-form coordinates by

$$\mathbf{X} = A\mathbf{U}. \tag{4}$$

where $\mathbf{U}$ denotes the scaled normal-form coordinate vector used for sampling, while A maps these coordinates back to the lab frame. $\mathbf{U}$ is also scaled using the beam covariance matrix in the normal-form coordinates, so that the resulting covariance matrix of $\mathbf{U}$ is a unit matrix,

$$\Sigma_U = I. \tag{5}$$

Neural Spline Flow (NSF) training and MC integration occur in this scaled normal-form coordinate space, keeping the learned PDF and MC sampling consistent. The coordinate transformations are accounted for in the integral evaluation using their respective Jacobians.

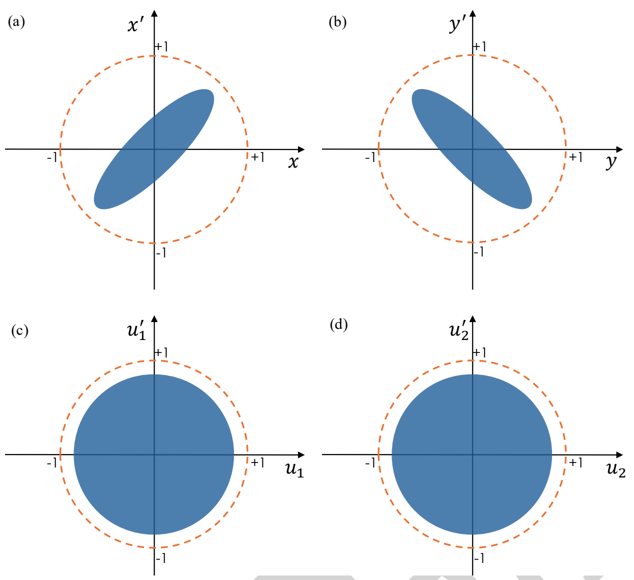

Figure 1: Schematic comparison of the sampling strategy. The blue regions indicate the beam distribution while the orange dashed circles indicate the sampling region. Panels (a) and (b) show sampling in the scaled lab-coordinates. Panels (c) and (d) show sampling in scaled normal-form-coordinates.

Figure 1 illustrates that the sampling efficiency depends on the coordinate choice despite the same underlying physics. It is convenient to sample the non-uniform MC evaluation coordinates using a simple uncorrelated 9D Gaussian distribution. When the 2D beam projections are strongly correlated in the lab coordinates, such as x-x', y-y', x-y', etc., they may not be efficiently covered by the fixed sampling region. In scaled normal-form coordinates, the same distribution is placed on an approximately unit scale, improving overlap between the sampling region and the high-density part of the bunch. This reduces low-weight samples and makes the sampling procedure more efficient.

Treatment of Spatial Delta Function

The spatial delta function in the IBSt integral constrains the two interacting particles to be at the same position in the lab frame:

$$(x_1, y_1, z_1) = (x_2, y_2, z_2). \tag{6}$$

After transforming to scaled normal-form coordinates, this condition cannot be imposed by simply setting certain components of $\mathbf{U}_1$ and $\mathbf{U}_2$ equal. The equality must still hold after mapping back to the lab frame.

If A_r denotes the three rows of A that map the six scaled normal-form coordinates to the lab spatial coordinates x, y, z, then the position vector is

$$\mathbf{r} = A_r\mathbf{U}. \tag{7}$$

The spatial overlap condition therefore becomes

$$A_r\mathbf{U}_1 = A_r\mathbf{U}_2. \tag{8}$$

Equation (8) provides three constraints on the phase-space coordinates of the two particles. In the Monte Carlo implementation, nine independent scaled normal-form variables are sampled, while the remaining three variables are determined from the spatial overlap condition. This treatment preserves the local nature of the IBSt interaction while allowing the MC sampling to be performed in scaled normal-form coordinates. This step significantly improves the efficiency of IBSt loss calculation for a wide class of arbitrary distributions.

CAM-DOMINATED BEAM APPLICATION

We apply the scaled normal-form-coordinate method to CAM-dominated beam distributions. These beams are a useful test case because their transverse phase space contains strong inter-plane correlations, such as x-y' and y-x' [4]. These correlations make the distribution more difficult to sample efficiently using only scaled lab-coordinates.

For fixed projected beam parameters, the level of the canonical angular momentum is characterized by the eigen-emittance ratio

$$\Delta = \frac{\epsilon_2}{\epsilon_1} \tag{9}$$

where ϵ_1 and ϵ_2 are the transverse eigenemittances. Smaller values of Δ mean that one transverse eigenmode dominates the distribution. In the limit of $\Delta \to 0$, the beam's canonical angular momentum approaches maximum with the distribution consisting of a single circular mode. As $\Delta \to 1$, the beam approaches a conventional planar transverse distribution. Therefore, scanning Δ provides a way to study how this degree of the bunch's intrinsic correlation influences the IBSt loss.

For comparison, we also calculate the IBSt loss rate for a matched planar waterbag case. The waterbag distribution is constructed with the same projected beam parameters as the CAM-dominated beam, so that the comparison mainly reflects the effect of internal phase-space correlations rather than changes in the beam size or projected emittances.

RESULTS AND DISCUSSION

Sampling Efficiency

Figure 2 shows the effect of the coordinate choice on the loss calculation along the SNS linac. The scaled lab-coordinate calculation exhibits stronger point-to-point fluctuations, whereas the scaled normal-form calculation gives a more stable and precise estimate for the same number of MC samples. This improvement comes from sampling in coordinates that better follow the correlated structure of the beam distribution.

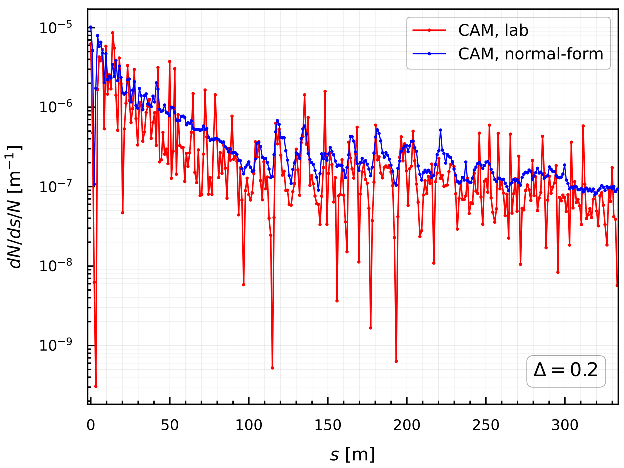

Figure 2: Normalized IBSt loss per unit length rate along the SNS linac for a CAM-dominated beam at $\Delta = 0.2$. The red curve is calculated using scaled lab-coordinate sampling, and the blue curve is calculated using scaled normal-form coordinate sampling.

The two calculations follow the same overall loss trend, indicating that the normal-form transformation does not change the underlying loss model. Instead, it improves the overlap between the sampling distribution and the high-density region of the bunch, reducing point-to-point fluctuations compared with lab-coordinate sampling.

Loss Dependence on Eigenemittance Ratio

To study how the correlated beam structure affects the IBSt loss, we vary the eigenemittance ratio $\Delta = \epsilon_2/\epsilon_1$, at a fixed linac location using local beam parameters. In Fig. 3, the red curve shows that the scaled lab-coordinate sampling gives a noisier estimate of the CAM loss rate. This is expected because the CAM beam contains stronger internal correlations, especially when one transverse eigenmode dominates at small Δ. In this case, many samples do not fall in the high-density region of phase space, as illustrated in Fig. 1. The blue curve shows that calculation using scaled normal-form coordinates is more robust.

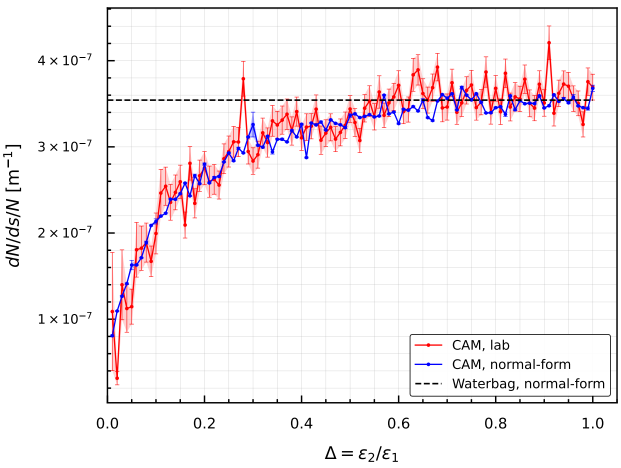

Figure 3: Normalized IBSt loss as a function of eigenemittance ratio $\Delta = \epsilon_2/\epsilon_1$ for CAM-dominated beams, evaluated at a fixed linac location. The red and blue curves compare scaled lab-coordinate sampling and scaled normal-form coordinate sampling, respectively. The dashed black line shows the matched waterbag reference, and shaded bands indicate $\pm 1\sigma$ Monte Carlo uncertainty.

The normalized IBSt loss for CAM-dominated beams increases with Δ. As the two eigenmode contributions become more comparable, the result approaches the matched waterbag reference. For smaller Δ, the CAM-dominated beams give lower loss than the matched waterbag case.

CONCLUSION

We developed a new algorithm to accurately calculate IBSt losses for arbitrary bunch distributions. We used an existing normalizing-flow toolkit to learn the beam PDF from discrete-particle tracking data. The resulting neural network was then used in the evaluation of the IBSt loss integral. The integration was performed in scaled normal-form phase space to optimize the training and MC sampling. We demonstrated that the new scaled normal-form method produces a more stable MC estimate with reduced point-to-point fluctuations compared with direct lab phase-space sampling. As an additional application, we scanned Δ and found that the CAM-dominated beam has substantially lower loss than the matched waterbag reference. This result suggests that the rotating, correlated structure can reduce IBSt losses.

REFERENCES

[1] A. Shishlo, S. Cousineau, J. Holmes, and T. Gorlov, "The particle accelerator simulation code pyorbit", *Procedia Comput. Sci.*, vol. 51, pp. 1272–1281, 2015. doi:10.1016/j.procs.2015.05.312

[2] S. Kakkar *et al.*, "Beam loss modeling and mitigation due to intra-beam stripping", in *Proc. NAPAC'25*, Sacramento, CA, USA, Aug. 2025, pp. 971–974. doi:10.18429/JACoW-NAPAC2025-THP014

[3] N. Solyak, A. V. Aleksandrov, V. A. Lebedev, A. P. Shishlo, and J.-F. Ostiguy, “Intrabeam Stripping in H- Linacs”, in *Proc. LINAC'10*, Tsukuba, Japan, Sep. 2010, pp. 929–931. doi:10.48550/arXiv.1207.5492

[4] K.-J. Kim, “Round-to-flat transformation of angular-momentum-dominated beams”, *Phys. Rev. Spec. Top. Accel. Beams*, vol. 6, no. 10, p. 104002, Oct. 2003. doi:10.1103/PhysRevSTAB.6.104002