

LONGITUDINAL BEAM-DYNAMICS OF THE SNS SRF-LINAC*

S. Nath, J. Billen, J. Stovall, H. Takeda, L. Young, LANL, Los Alamos, NM 87545, USA
 K. Crandall, Tech Source, Santa Fe, NM 87501, USA, and D. Jeon, SNS/ORNL, Oak Ridge, USA

Abstract

The spallation neutron source (SNS) linac is composed of both normal- and super-conducting RF (SRF) structures. The normal-conducting section (up to 185 MeV) consists of a Low-Energy Beam Transfer (LEBT) line downstream of the H⁺ ion source leading to a 2.5 - MeV RFQ, a Medium-Energy Beam Transfer (MEBT) line, a 402.5-MHz DTL, followed by a 805-MHz coupled-cavity linac (CCL). The SRF structure accelerates the beam from a nominal energy of 185 MeV to 1000 MeV. The SRF section consists of two sections: a low beta ($\beta_g = 0.61$) and a high beta ($\beta_g = 0.81$). In this paper, we discuss a longitudinal beam dynamics feature of the SRF linac.

1 INTRODUCTION

An earlier paper [1] describes the design essentials of the SNS linac, which injects beam into an accumulator ring. It is designed to deliver 1.4 MW of beam at 1 GeV with room for upgrade. At ~87 MeV, the beam from the 402.5-MHz DTL structure enters the 805-MHz CCL; the CCL structure is followed by the SRF structure. The underlying design principles for these two structures are quite different. Longitudinal tuning schemes for the two sections also are different. A simplified SRF-tuning scheme illustrates an interesting feature of beam behavior in the SRF section. It is more tolerant of errors in the cavity-phase settings than expected.

2 SRF LINAC

The SRF section of the linac has two groups of cavities characterized by two different cavity lengths or "geometric β " (β_g). The low- β (β_1) section with 33 6-cell cavities, accelerates the beam to 394 MeV. The high- β (β_2) section with 48 6-cell cavities takes the beam to a final energy of 1 GeV. The low- β section has 11 cryomodules containing 3 cavities each, and the high- β section has 12 cryomodules with 4 cavities each.

2.1 Design

The final energy of 1 GeV in the SRF linac design assumes that $E_{max} = 27.5$ MV/m and $E_{max} = 35.0$ MV/m for the peak-surface electric field of the two cavity types, respectively. Figure 1 shows the design values of E_{max} , E_0 , the average axial accelerating field, and E_0T , where T is the transit-time factor.

Given the cavity field, the number of cavities, and the final energy leaves only the cavity-phase rule to be established. The design phases, $\phi_{design} = -20.5^\circ$ and -19.5°

in the low- and high- β sections, respectively, meet the design requirements and provide a smooth match at the CCL: β_1 and β_1 : β_2 transitions. Error studies indicate that this choice provides adequate longitudinal acceptance.

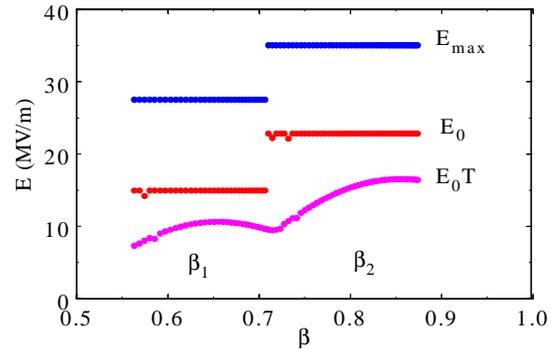


Figure 1. Electric field profiles in the SRF linac.

Unlike the CCL cavities, the SRF cavities are expected to have a spread in achievable accelerating gradients. Figure 2 shows the expected peak-surface electric field distribution for the two cavity types.

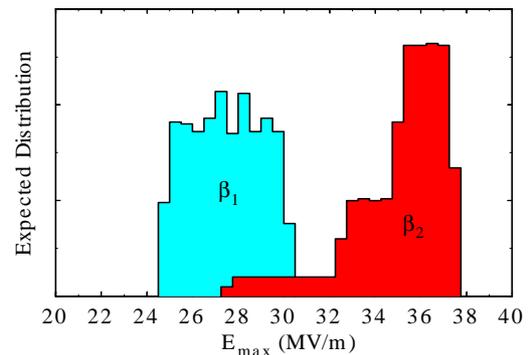


Figure 2. Expected peak surface electric field distribution in the SRF cavities.

The project schedule precludes any sorting of the cavities to optimize their location in the linac based on achievable fields. Thus, they will be installed in order of delivery. Because some of the cavities may not achieve the design field, the decision was made to set the field in each cavity to its maximum achievable value. This choice leaves several options for setting the phase in each cavity. One possibility would be to adjust the phase so that the energy gain in each cavity is the design value. Another possibility is to adjust the phase so that the longitudinal phase advances remain at the design values, which would preserve the matching along the structure. A third possibility would be to just set the phase of each cavity to its design value regardless of the accelerating field.

*Work supported by the Office of Energy Research, Basic Energy Science of the US Department of Energy.

2.2 Tuning Schemes

Our initial tuning strategy was to preserve the design longitudinal phase advance by adjusting the cavity phase corresponding to the operating field so that,

$$k_{l0}^2 \equiv \frac{E_{0,dsn} T(\beta) \sin \phi_{dsn}}{(\beta\gamma)^3} = \frac{E_{0,op} T(\beta) \sin \phi_{op}}{(\beta\gamma)^3} \quad (1)$$

Figure 3 shows the design cavity-accelerating field and the phase in the two sections. The scattered dots represent a random sampling of the cavity fields about the mean design value and the corresponding phase values consistent with the above prescription. The longitudinal phase advance per unit length k_{0l} remains at the design value across the entire SRF linac. The points outside the smooth curve are values used in matching the beam at the transitions.

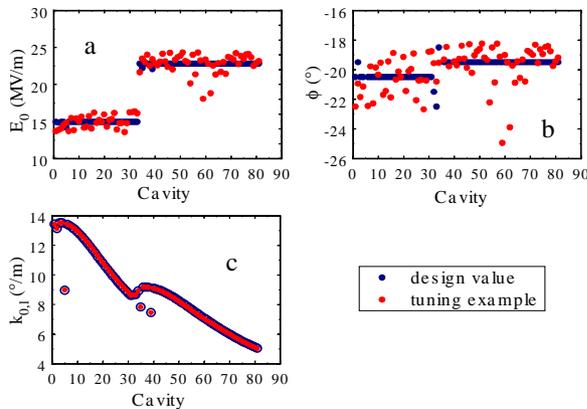


Figure 3. (a) Design accelerating field E_0 and a random sampling of expected values, (b) design phases consistent with keeping k_{0l} to design values, and (c) k_{0l} values for design and sample case.

This approach of setting up each cavity phase, $\phi_{operating}$ to preserve the longitudinal phase advance k_{0l} requires prior knowledge of the maximum accelerating field, E_0 . For each cavity, in sequence, the following steps are envisioned to set the phase and amplitude. With the cavity under test tuned to the designed frequency, we use the drifting beam to excite the cavity. Comparison of the measured excitation with the calculated value provides a calibration of the cavity pick-up loops and the LLRF (low level RF) system's phase and amplitude [2]. Since the beam-excited fields are decelerating, the calibration point corresponds to $\phi = -180^\circ$. We now set the cavity E_0 to its maximum achievable accelerating field using the calibrated loop. The corresponding cavity phase is then calculated and set. The next step is to turn on the RF to that cavity, accelerate the beam, and proceed to the next cavity. The procedure has to be repeated for each cavity in the linac.

This process of tuning is straightforward, but it requires stepping through the whole procedure when any of the cavities needs adjustment of the operating field. An alternative approach is to set the field of each cavity to its maximum achievable accelerating field E_0 and set the

phase of all the cavities to their design values i.e., $\phi_{design} = -20.5^\circ$ and -19.5° in the medium- and high- β sections, respectively. Thus, no attempt is made to keep the k_{0l} values the same as designed. Figure 4 illustrates this approach. The scattered dots in 4a represent a variation of E_0 within $\pm 10\%$ for β_1 - and $+10\%$ to -25% for β_2 -cavities, respectively. The blue scattered points in figure 4b correspond to the values if k_{0l} were to follow the design values. Instead, we kept the phases to the design values represented by the solid lines in 4b. The effect on k_{0l} is shown in figure 4c. The scattered points in 4c represent a sample of resultant k_{0l} values used for simulation. This alternative approach saves no significant time because the operating phase must still be determined by exciting each cavity in turn with the beam. Calibrating the cavity loop and calculating the proper phase can be automated and thus may not take more time.

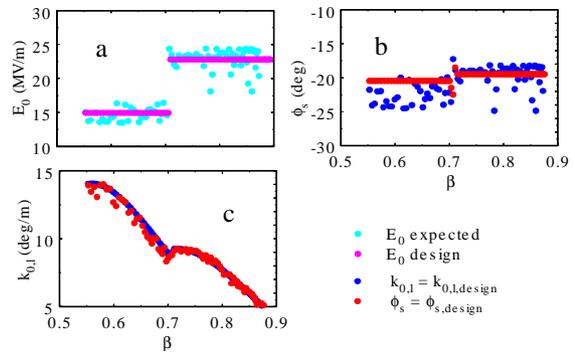


Figure 4. Design accelerating field E_0 and cavity phases without phase adjustment.

2.3 Simulated Beam Performance

In figure 5, we show the relative probability histograms for the final output energy. In figure 5a, we compare the spreads between the cases when k_{0l} is maintained at its design value and when E_0 is varied randomly by $\pm 10\%$ and each cavity set to the design phase. The difference of the probability distribution is insignificant between the two cases. For both cases, the median value is ~ 1 GeV with a maximum spread of $\sim \pm 25$ MeV. The same holds true when the results are compared for cases with $E_0 \pm 10\%$ and $E_0 \pm 30\%$ as is shown in figure 5b. Ten thousand simulations were made in each case.

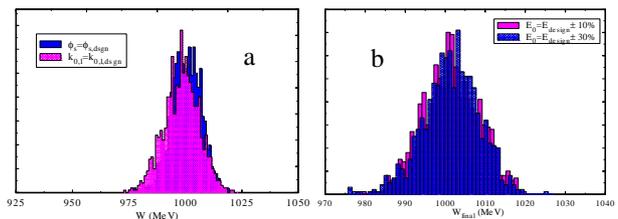


Figure 5. SRF linac output energy probability distribution for (a) $E_0 \pm 10\%$, and (b) $E_0 \pm 30\%$.

In order to study the effect of this alternate approach on beam dynamics, we started with a matched 6-D waterbag distribution of 10^5 macro-particles at the input to the SRF

section. We studied several cases where the accelerating field E_0 is allowed to vary randomly within certain tolerance limits. Results are shown here for only (a) $E_0 \pm 10\%$, and (b) $E_0 \pm 30\%$. Each cavity in the SRF section is independently powered by a separate klystron. This configuration enables us to phase the RF power to compensate for the deviation of output energy from each cavity, which affects the arrival time of the beam from one cavity to the next. In table 1, we compare the average emittance values for each case with that of the design case. For case (a), both transverse and longitudinal emittance values are virtually unchanged from the design. There is about 10% more transverse emittance growth compared to the design case when E_0 is allowed to vary randomly by up to $\pm 30\%$. There is, however, $\sim 35\%$ more increase in the longitudinal emittance value compared to the design case.

Table 1. Emittance values for studied cases.

$\langle \mathcal{E}_{99\%} \rangle$ (π cm-mrad)	Input	Design	$E_0 \pm 10\%$	$E_0 \pm 30\%$
$x-x'$	0.250	0.405	0.405	0.432
$y-y'$	0.231	0.253	0.254	0.269
ϕ -W	0.257	0.334	0.344	0.431

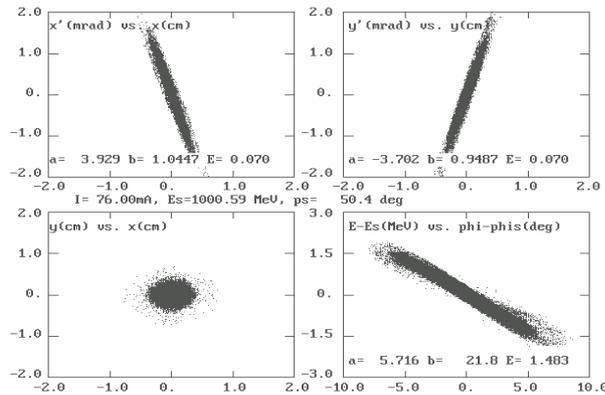


Figure 6. Transverse and longitudinal phase-space projections at the SRF linac end for the design case.

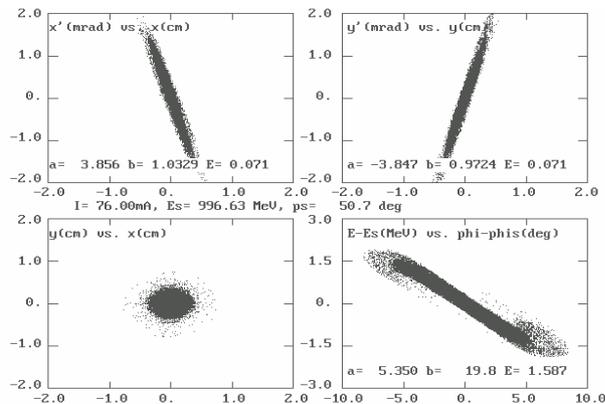


Figure 7. Transverse and longitudinal phase-space at the SRF linac end for case (a).

Figures 6 and 7 show the phase-space projections at the SRF linac end for the design case and for case (a). In agreement with the emittance-growth predictions shown

in table 1, the transverse phase-space projections are not visually distinguishable from that of the design case. In the longitudinal phase space, however, we see wings forming at the bunch ends. For case (b), (figure not shown) the bunch-shape distortion in the longitudinal phase-space becomes more prominent.

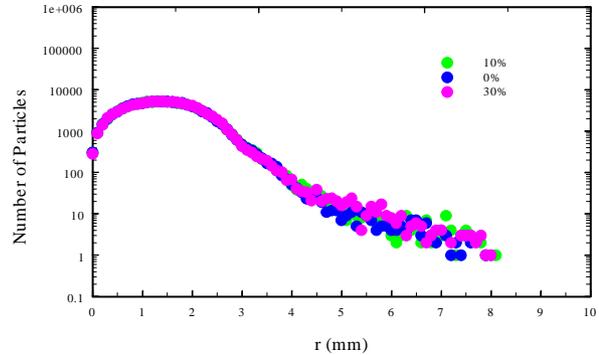


Figure 8. Radial particle distribution at the linac end.

In figure 8, we show the radial distribution of particles at the end of the linac. The values for both case (a) and (b) are comparable with that of the design case. In both cases, the maximum radial extent of the beam is about 8 mm which is almost the same as the design case.

3 CONCLUSION

The SRF linac having one klystron powering one cavity lends itself to a surprisingly tolerant design. Constant ϕ_s for accelerating electric field different from the design value by as much as $\pm 30\%$ represents significant deviation from the design longitudinal tune. Such deviation introduces distributed longitudinal mismatch along the entire SRF linac but appears to have minimal effect on the beam performance. For electric-field deviation within $\pm 10\%$, there is practically no effect either on the transverse or longitudinal phase-space. Deviations of E_0 as large as $\pm 30\%$ have a small effect on the transverse dynamics. This low sensitivity is, in part, a direct outcome of the nature of the SRF linac where each cavity's phase can be set to the desired value. In contrast, a normal-conducting CCL is a synchronous structure where each cavity is phase-locked to its neighbor and its phase is not adjustable. Hence, it does not lend itself to the same flexibility.

4 REFERENCES

[1] J. Stovall, J. H. Billen, S. Nath, H. Takeda, L. M. Young, D. Jeon, R. Shafer, and K. Crandall, "Expected Beam Performance of the SNS Linac," Proceedings of the 2001 Particle Accelerator Conference, p 446, June 18-22, 2001, Chicago, IL, USA.
 [2] L. M. Young, "A Procedure to Set Phase and Amplitude of the RF in the SNS Linac's Superconducting Cavities," Proceedings of the 2001 Particle Accelerator Conference, p 572, June 18-22, 2001, Chicago, IL, USA.