

SUPERCONDUCTING LINAC FOR THE SNS

J. Stovall, S. Nath, H. Takeda, J. Billen, L. Young, M. Lynch, D. Rees (LANL), J. Galambos, D. Jeon, D. Raparia, J. Wei (SNS/ORNL), R. Sundelin (TJNAF), K. Crandall (TechSource), C. Pagani, P. Pierini (INFN)

1. ABSTRACT

The Spallation Neutron Source (SNS) linac is comprised of both normal and superconducting rf (SRF) accelerating structures. The SRF linac accelerates the beam from 186 to 1250 MeV through 117 elliptical, multi-cell niobium cavities. This paper describes the SRF linac architecture, physics design considerations, cavity commissioning, and the expected beam dynamics performance.

2. INTRODUCTION

The SNS linac uses four types of accelerating structures. A 402.5-MHz radio-frequency quadrupole (RFQ) linac accelerates an H^- ion beam from the injection energy of 65 keV to 2.5 MeV. A 6-tank drift-tube linac (DTL), also at 402.5 MHz, accelerates the beam to 86.5 MeV. The next section, to an energy of 186 MeV, is an 805-MHz coupled-cavity linac (CCL) of the side-coupled type. Each of its four rf modules contains 96 accelerating cells and is powered by a 5-MW klystron.

The major portion of the linac, which accelerates the beam to 1.25 GeV, is a 805-MHz SRF linac. This linac is designed to eventually deliver 2.65 MW of beam, having a peak current of 52 mA and a duty factor of 6%, for injection into the accumulator ring.

3. SRF ARCHITECTURE

The SRF linac has two sections distinguished by cavity length or “geometric β ”. The β_1 section with 33 6-cell cavities accelerates the beam to 382 MeV. The β_2 section will initially have 59 6-cell cavities and an interim output energy of 974 MeV. The civil construction includes space and utilities for future installation of 25 more β_2 cavities. In this final (1.25-GeV) configuration, there are 11 cryomodules with three β_1 cavities each, and 21 cryomodules with four β_2 cavities each.




Figure 1. Transit-time factors for 5 elliptical SRF cavities as a function of particle velocity.

The rf power system has 117 550-kW klystrons. We investigated several potential cost-saving schemes to power the 117 SRF cavities using a smaller number of very high-

power klystrons. However, we concluded that driving each cavity from a single klystron, each with its own control loop, was the best solution using current technology. The klystrons will be powered in groups of 11 or 12 per high voltage system, depending on the power required in each portion of the linac. Twenty transmitters provide electronics and interlocks for groups of 6 klystrons.

Because the SNS linac accelerates ions with $\beta < 1$, the beam requires longitudinal focusing. Thus, beam bunches traverse each acceleration gap at a particular rf phase. With Lorentz detuning forces expected to be 250 to 400 Hz, and microphonics expected to be up to ± 100 Hz, 25% of the klystron's power is reserved for field control. The balance of 412 kW is available to accelerate the beam. The power dissipated in the cavity walls is negligible, however the rf system is designed to dissipate a significant amount of reflected power.

4. SRF LINAC DESIGN

The cavity designs were initially driven by two constraints. Power coupler performance was limited to a maximum of 350 kW (later increased to >550 kW), and the design value for the cavity peak surface electric field was limited to 27.5 MV/m. Assuming one power coupler per cavity, these two constraints limited the β_2 cavities to 6 cells each. The β_1 cavities were limited to 6 cells by longitudinal beam-dynamics considerations related to phase slip at low energy. The 550-kW power limit corresponds well to SRF cavities that would be available commercially. We arrived at this value by scaling measurements of non-heat-treated TTF cavities to the SNS frequency.

The cavity shapes resulted from studies to optimize the transit-time factor, balance the peak surface magnetic and electric fields, raise the resonant mechanical frequencies, minimize the effects of Lorentz detuning and higher-order cavity modes.

After investigating a variety of SRF linac configurations, we found a rough cost optimum that used only two cavity shapes. The choice of $\beta_1 = 0.61$ was based on previous 2-cavity design studies and the fact that the injection energy fits well the 186-MeV output energy of the CCL. We investigated values for β_2 between 0.74 and 0.83.

The energy gain per cavity $\Delta W = E_0 T L \cos \phi$, where E_0 is the average axial electric field, T is the transit-time factor, L is the cavity length, and ϕ is the design phase. Three of these parameters increase with the cavity geometric β resulting in higher acceleration efficiency. First, E_0 is higher for a given peak surface electric field; second, T is higher over a broad energy range; and third, L is longer. Figure 1 shows the transit-time factors for some candidate cavities.

The transition energy between β_1 and β_2 sections has a broad optimum range. For a fixed-price machine, the optimum depends weakly on the final energy. Figure 2 shows the expected final energy for a 99-cavity linac as a function of transition point for three choices of β_2 . In this example, which

is similar to the initial configuration of the SNS linac, the optimum value is $\beta_2 = 0.81$ with a transition after cavity 33 of the β_1 section. For longer linacs this optimum range is broader.




Figure 2. Final energy for a 99-cavity SRF linac as a function of the number of β_1 cavities.

The high accelerating fields generate strong rf defocusing forces on the beam. Quadrupole doublet magnets between the cryostats provide the transverse focusing. Figure 3 shows the zero-current transverse real-estate phase advance k_{t0} for constant $G\ell$ product. Studies are underway to avoid crossing these two curves and to avoid a parametric resonance.




Figure 3. Transverse and longitudinal real-estate phase advance in the SRF linac.

Because the SRF linac is not a synchronous accelerator, in the sense that the cavity lengths do not keep step with the beam velocity, there is considerable flexibility in trading beam energy for beam current to achieve a desired beam power. The lower curve in Figure 4 is the power imparted to the beam in each cavity. The upper curve reflects the klystron power allocated to each cavity including the control margin.

Given values for β_1 and β_2 , and assuming optimized cavity designs, parameters $T(\beta)$, $E_0(E_{\text{peak}})$, and L are fixed. The only remaining design variable is the phase ϕ of the cavity field relative to the average arrival time of the beam at the center of each cell. We choose a phase law that assures smooth transitions from the CCL into the SRF linac and between the two SRF linac sections. The phase starts at a large negative value and then gradually ramps up, providing strong longitudinal focusing at low energy. The phase ramp stops at –

15°, and remains fixed over most of the linac, increasing the acceleration efficiency at high energy. The resulting zero-current phase advance per unit length (k_{t0}) appears in Figure 3.

We expect a 20% range in the cavities' maximum sustainable peak surface electric field. The installation schedule precludes sorting the cavities to optimize their location in the linac. They will be installed in the order of delivery. Rather than preserving the design energy gain in each cavity we chose to preserve the design longitudinal phase advance by adjusting the cavity phase to correspond to the operating field as defined by:

$$k_{t0}^2 \equiv \frac{E_{0,\text{dsn}} T(\beta) \sin \phi_{\text{dsn}}}{(\beta\gamma)^3} = \frac{E_{0,\text{op}} T(\beta) \sin \phi_{\text{op}}}{(\beta\gamma)^3}$$

The consequence of this choice is that the beam quality appears to be preserved for a wide variety of random samples of cavity performance.

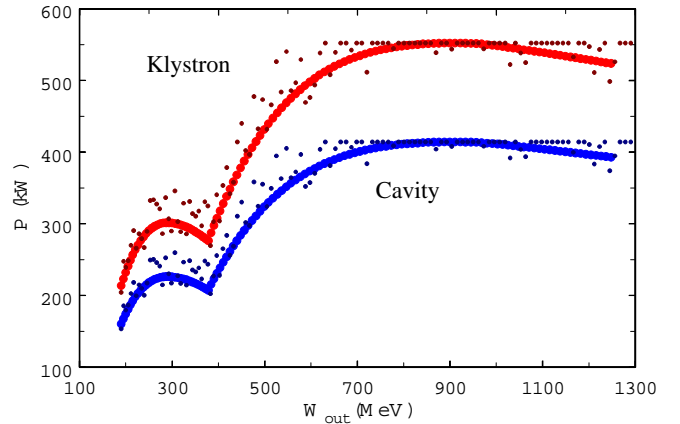


Figure 4. Power per SRF cavity versus beam energy.

We expect at least half the cavities to support a surface electric field higher than 27.5 MV/m. One might expect to increase the field in these cavities to further accelerate the beam. The scattered dots near the lower curve in Figure 4 represent a random sampling in cavity performance ranging from 5% below to +15% above the design field. The dots near the upper curve represent the klystron power that could be delivered to these cavities. For half of the linac (the most efficient half) the cavity fields are limited by the installed power and not by the quality of the cavities. Even so, in this example, the final full-current beam energy reaches 1.3 GeV.

5. SRF CAVITY COMMISSIONING

Before installing the cavities, we will have established during acceptance testing the approximate peak accelerating field E_{op} . After installation, drifting beam from the CCL will excite fields in each cavity allowing calibration of the cavity pickup loops and low-level rf (LLRF) system.

With all the SRF cavities unexcited and detuned we inject a 100- μ s-long pulsed beam into the SRF linac. The beam drifts through a portion of the linac to a low-power beam dump. This procedure requires that the transverse beam optics (quadrupole and steering magnets) have been tuned for maximum transmission.

For each cavity, in sequence, the following steps set the phase and amplitude. With the cavity under test tuned to the

design frequency, we allow the drifting beam to excite the cavity. We know the approximate beam energy, and we can measure accurately the integrated current. Therefore, the calculated cavity excitation provides a calibration of the cavity pickup loops and the LLRF system's phase and amplitude. Because the cavity fields rise linearly (approximately) with time for a constant-current beam pulse, the field at the end of the pulse is proportional to the total charge in the beam. We expect that the field derived from this measurement will be accurate to $\sim 1\%$, which is about five times better than needed for this technique.

The phase is very sensitive to the cavity tune and increases almost linearly with time during the excitation pulse. A 1-kHz error in the cavity frequency would introduce a phase error of $\sim 16^\circ$ at the end of a 100- μ s pulse, giving a clear indication that the cavity is mistuned. For cavity frequencies within 30 Hz of the design value, the phase changes less than 0.5° during the 100- μ s pulse. An advantage of measuring phase and amplitude in this manner is that during a short pulse (e.g., 100 μ s) the cavity excitation is nearly independent of the external Q. Excitation is also relatively insensitive to the resonant frequency of the cavity, which means detuning caused by microphonics is unimportant during the measurement. Because the rf fields are small during this measurement, Lorentz force detuning also is not an issue.

Because the beam-excited fields are decelerating, the calibration point corresponds to $\phi = -180^\circ$. We now excite the cavity at its design amplitude and phase to accelerate beam. These calibrated set points determine the beam energy at the exit of this cavity, which we use for performing the same procedure on the next cavity. To avoid accumulating errors, we measure the beam energy periodically using the time-of-flight method. We repeat the calibration procedure for each cavity until the entire linac has been tuned. The entire procedure must be repeated downstream of any significant change in the linac, such as removal or insertion of a cavity. Fortunately, the procedure lends itself to automation.

The calibration procedure has an expected accuracy of $\pm 5\%$ in amplitude and $\pm 2^\circ$ in phase, which results in an uncertainty of ± 20 MeV in the final beam energy. This uncertainty is a static error. Dynamic energy jitter depends on performance of the LLRF control system. Figure 5 shows a ± 1 - to ± 2 -MeV expected dynamic error.

6. SIMULATED PERFORMANCE

We used PARMILA for beam-dynamics simulations through the MEBT, normal conducting linac, CCL-to-SRF transition region, and SRF linac. The simulation starts with 10,000 macroparticles in a uniform 3-D distribution that has the RMS properties of the RFQ exit beam. The results shown here do not include errors. Figure 6 shows the beam's vertical spatial profile and the beam-energy profile from the RFQ exit to the end of the linac. In the SRF linac, the beam is well matched and occupies less than 25% of the clear bore.

Figure 7 shows the phase-space projections x - x' , y - y' and W - ϕ , and the spatial x - y projection at the exit of the linac. These calculations show that in the absence of errors the beam is stable, well behaved, and develops only about a ± 2 MeV energy spread. The emittance grows by 14% transversely and by 12% longitudinally through the SRF linac.




Figure 5. Expected energy jitter versus rf control tolerance.




Figure 6. Spatial and energy profiles of the SNS linac beam.




Figure 7. Beam distribution at 1.25 GeV.

7. CONCLUSION

We have designed a cost-optimized, simple SRF linac that has only two elliptical cavity types. The linac appears, from beam simulation studies, to be flexible and stable. We have developed techniques for tuning the linac, which will provide a high quality beam.

8. ACKNOWLEDGEMENT

This work is supported by the US Department of Energy.