

A FIFTEEN YEAR PERSPECTIVE ON THE DESIGN AND PERFORMANCE OF THE SNS ACCELERATOR

S. M. Cousineau, Oak Ridge National Laboratory, Oak Ridge, TN

Abstract

Construction of the Spallation Neutron Source (SNS) accelerator began approximately fifteen years ago. Since this time, the accelerator has broken new technological ground with the operation of the world's first superconducting H⁻ linac, the first liquid mercury target, and 1.4 MW of beam power. This talk will reflect on the issues and concerns that drove key decisions during the design phase, and will consider those decisions in the context of the actual performance of the accelerator. Noteworthy successes will be highlighted and lessons-learned will be discussed. Finally, a look forward toward the challenges associated with a higher power future at SNS will be presented.

INTRODUCTION

The SNS accelerator was designed as a short pulse, high power proton driver for neutron production. The top levels goals of the accelerator are to provide 1.4 MW of proton beam power with 90% reliability. The final proton beam exiting the accelerator is composed of 1 us pulses of 1 GeV protons operating at a repetition rate of 60 Hz. A third goal, related to the reliability metric, is to maintain beam loss levels to the order of < 1 W/m, corresponding to approximately 100 mrem/hr residual radiation at 30 cm distance, throughout the accelerator in order to allow for routine, hands on maintenance of system components.

To obtain the short pulse structure of the beam, an H⁻ linac and accumulator ring combination was chosen. From there, the design decisions for the subsystems were driven by the high power, high reliability, low loss goals stated previously. These decisions and their impact on accelerator performance will be discussed in the forthcoming sections, following a brief summary of the accelerator performance to date.

PERFORMANCE METRICS

The accelerator was commissioned beginning in 2002, and the first neutrons were produced in 2006. The power ramp up took longer than planned due to difficulties in the target systems [1]. The accelerator was operated in production mode with 1.4 MW of beam power for the first time in the fall of 2015. After this production cycle resulted in a premature target failure, the beam power was reduced to 1.1 MW in a move to prioritize reliability for the neutron users. The power will be ramped back up in to 1.4 MW in a stepwise fashion over the course of the next few years in a controlled study of target cavitation damage versus beam power. Fig. 1 shows the beam power evolution of the SNS accelerator since the beginning of operations in 2006.

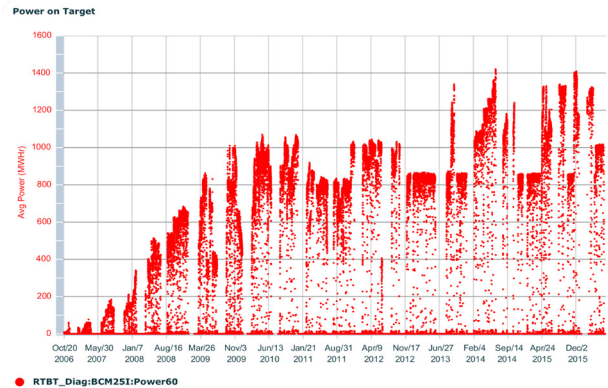


Figure 1: SNS beam power evolution.

The accelerator reliability metric is defined as the number of hours of delivered neutron production divided by the number of hours scheduled. Catastrophic equipment failures such as target failures that result in long downtimes have major impact on the reliability metric. The SNS has suffered seven premature target failures, as well as one equipment failure in the MEBT with comparable downtime. Target failures have progress from early life failures due to manufacturing details to failures due to cavitation damage at high power. Aside from these single event failures, the remainder of the accelerator systems are operating with very high reliability, as demonstrated in Fig. 2, which shows the historical reliability metric with and without the target and MEBT failures.

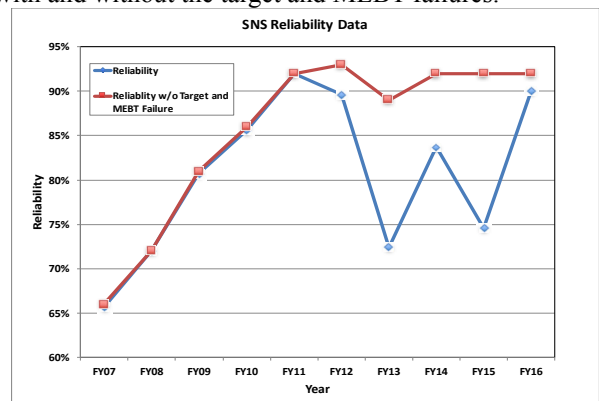


Figure 2: SNS accelerator reliability metric.

Finally, while reliability and beam power are the metrics relevant to the neutron user program, for the accelerator the level of residual radiation is top level consideration which determines how efficiently equipment maintenance can be performed. Table 1 below shows the activation levels throughout the accelerator following the 1.4 MW production run. With the exception of the ring injection area, which was always anticipated to be hot, the activation is well below the 1 W/m criterion.

Table 1: Activation 3-5 Hours after a 1.3 MW Production Beam Run

Region	Activation Level (mrem/hr @ 30 cm)
DTL	2 – 30
CCL	8 – 60
SCL	5 – 45
LEDP	90
HEBT General	< 5
HEBT DH25	90
Ring Injection	1000
Ring Extraction	80
Ring Collimation	90
Ring General	5-40

THE LINEAR ACCELERATOR

The linear accelerator is composed of a warm linac DTL and CCL combo to a beam energy 186 MeV, followed by a superconducting linac (SCL) with a medium beta section ($\beta=0.61$) and a high beta section ($\beta=0.61$) to 1 GeV. The superconducting technology was chosen over the warm linac technology for a number of reasons, including reduced cost of construction and operation, higher availability compared to a warm linac, high vacuum to beam-gas scattering, and a large bore aperture to reduce beam loss. Since it was the world's first H⁻ SCL, the performance expectations were somewhat unknown, and the choice was considered both high risk and high potential. Although by now it is clear that the SCL has been a success, the first decade of operation has offered a number of surprises, both good and bad.

Expectations vs. Realities – SCL Cavities

The design gradients for the SCL cavities were 10.2 MV/m for the medium beta cavities, and 15.8 MV/m for the high beta cavities. When the cavities were first powered up at their design repetition rate of 60 Hz, the gradients for the medium beta cavities were generally above the design values, but the high beta cavities were well below expectations, with some cavities unable to run at all. The most majority of problems were associated with electron activity in the cavities, which limited the gradients to below design values for 51 cavities [2]. Due to these issues, the SNS linac has not yet been run in production with the design beam energy of 1 GeV. Much progress has been made toward repairing defective cavities and improving field gradients through mechanisms such as plasma processing, and the production beam energy is now 958 MeV. Fig. 3 shows the cavity gradients versus design early in the operational cycle in 2007, and during the 1.4 MW run in 2015.

One performance aspect that became apparent during the process of removal and repairs of cavities is the exceptional flexibility and adaptability of the SCL. This is due to the combination of a spare cavity to provide energy reserve, and individually powered cavities that allow for

retuning an rephasing to obtain the same beam energy for the accumulator ring. The additional benefits of a digital LLRF system that allows beam blanking, a robust BPM system for time of flight measurements, and sophisticated applications software make it possible to tune up the entire SCL from scratch in a completely automated fashion in 40 minutes, and to retune after a cavity loss in only 20 seconds [3]. This level of expediency was never imagined during the design phase.




Figure 3: SCL cavity gradients in 2007 (red), in 2016 (blue), and for design (black).

Expectations vs Realities - Beam Dynamics in the Linac

Since the SNS SCL is the world's first superconducting H⁻ linac and there was not wisdom and experience to be garnered from elsewhere, simulations played a key role in setting the performance expectations [4]. From simulation work, it was initially thought that the SCL would be tuned up with design cavity phases to preserve the longitudinal match, and additionally to correlate the longitudinal and transverse phase advances. Finally, it was expected that the beam would be transversely matched throughout the entire linac, using matching algorithms and the intermittent dedicated matching sections in the linac. Between this model, and the very large bore in the SCL, simulations predicted a negligible amount of beam loss in the SCL region.

None of these expectations met reality. Regarding the cavity phases, the SCL was first configured to a constant focusing value (-18 degrees from the synchronous phase) and the final configuration resulted from tuning the cavity phases on beam loss until the lowest beam loss state was reached.

In the transverse plane, a significant and unanticipated beam loss phenomenon has driven the lattice settings far away from the initial design value. H⁻ intrabeam striping, which was not realized during the design stage, lead to significant beam loss in the SCL. The loss mechanism, which scales with beam density, and was confirmed through an experiment which compared beam loss levels for operating with protons and with H⁻ in the SCL [5]. To mitigate beam loss from this mechanism, the quadrupoles in the SCL have been reduced to approximately half of their design values. Any further reduction from these

values results in an increase in beam loss, which will be discussed shortly.

Prior to the reduction in transverse focusing, the activation in the SCL was increasing with beam power at an alarming rate. Fig. 4 shows the historic evolution of the activation in the SCL, and the reduction in slope after the defocusing was put in place. Clearly, if the design quadrupoles had been used for beam powers up to 1.4 MW, the average activation in the SCL would have been ~100 mrem/hr, which falls into the regime of a high radiation area that requires special work permit for maintenance.

It is worth noting that the original SNS linac design was a warm linac with approximately half the beam pipe aperture. If this design had been chosen instead of the large bore (76 mm diameter) SCL option, the SNS could not have been able to simultaneously achieve its high power goal while maintaining the < 1 W/m beam loss standard. Operation at high power would have resulted in high levels of residual radiation in the linac tunnel which would have complicated maintenance and cause degradation of system components.

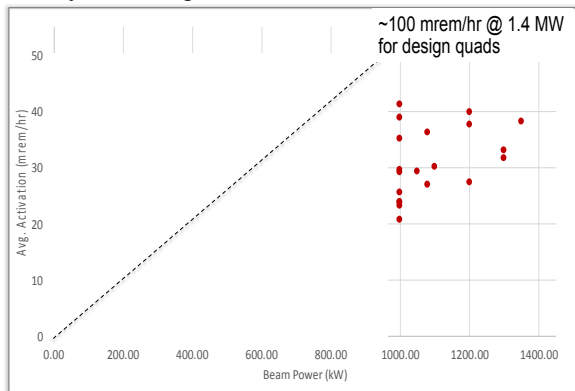


Figure 4: Average SCL activation versus beam power. The dashed line indicates the trend prior to reduced focusing.

Following along the trend of breaking from the design plans, the beam in the linear accelerator is not matched in any of the three planes, nor do the transverse phase advances match the longitudinal phase advances [3]. In practice, the beam envelope is arrived at through “monkey tuning” of the beam losses in the linac, and has significant beating in all planes.

Early efforts to transversely match the beam in the linac were thwarted hindered by an inability of the models to reproduced the measured RMS values. A multi-year campaign has now resulted in good agreement between model and measurement in an RMS sense in the SCL in both the longitudinal and transverse planes. The next steps in this effort are to achieve the same agreement in the warm linac, and then to apply control of the beam optics and matching using these tools.

Even after the defocusing of the optics to reduce intra-beam stripping, the transverse RMS beam size is still small compared to the diameter SCL bore aperture. In fact, the RMS is a factor of ~10 less than the bore size of

76 mm. Therefore, remaining beam loss in the SCL is likely due to extended beam halo. Presently, the source and character of the beam halo is not understood, but efforts are underway to resolve this problem.

Although the definition of “beam halo” is often disputed, SNS adopts the convention from the 2014 Workshop on Beam Halo Monitoring which defines beam halo as $10^{-4} - 10^{-6}$ of the beam intensity. Recently, a few diagnostics systems have been upgraded to detect beam at this level and efforts to utilize these diagnostics to understand and control beam halo in the SNS linac are underway [6].

Expectation vs. Realities – The MEBT Chopper System

One significant surprise which occurred outside of the SCL was the performance of the MEBT chopper system. The design of the SNS accelerator relied on two chopper systems – the LEBT chopper and the MEBT chopper – to provide the necessary microsecond beam structure for the accumulator ring. The LEBT performs the initial chopping, but leaves behind a 25 ns partially chopped beam tails. It was assumed during design that the partially chopped beam would fall outside of dynamic aperture and result in unacceptable beam loss in the linac and beam in gap in the ring. The fast MEBT chopper, with a rise time of 5 ns was designed remove most of these tails.

The inclusion of this fast chopper severely constrained the design of the MEBT. A MEBT without the fast chopper would have required four quadrupoles and one rebuncher to match the beam into the first DTL tank. On the other hand, the SNS MEBT required fourteen quadrupoles and four rebunchers to achieve the proper phase advance between chopper and antichopper.

While the MEBT chopper performed according to specifications, it did not have an appreciable impact on the beam loss in the linac. A small difference in the extraction region loss in the accumulator ring was observed with the MEBT chopper on, but since the losses were already low in this region, there was no significant benefit. In fall of 2014, the MEBT chopper target leaked and flooded the entire MEBT. The recovery resulted in a complete disassembly and reassembly of the MEBT, requiring four weeks of unscheduled downtime.

THE ACCUMULATOR RING

The SNS accumulator ring, which has accumulated up to 1.56×10^{14} ppp is the most intense proton ring in the world on a charge per pulse basis. In order to meet the activation goals, beam loss must be kept on the order of 10^{-4} of the beam intensity. The design of the ring was highly focused on controlling this beam loss and a variety of both large and small investments were made to assure that this goal was met. A detailed description of the design can be found in [7]. Overall, the ring has performed extremely well and in fact, and could accept a higher beam intensity while still maintaining reasonable loss levels. Nonetheless, ten years after the start of operations, it is interesting to the pose questions: What investments paid off? What investments didn't pay off? What was over-

looked during the design and what was the consequence? These are the questions that will be addressed in this discussion.

High Payoff Investments

Probably the largest pay off investment in the ring was the large beam aperture, which varies between 10 – 16 cm in the collimation region, to 20 – 30 cm for the beam pipes. The large aperture allows the beam to be injection painted into a large area which subsequently reduces the space charge tune shift and associated effects such as resonance crossing. In practice, the entire aperture is used during the accumulation cycle, which minimizes the foil traversals and resulting beam loss in the injection region. Another high payoff investment was the dual plane injection painting system, which allows independent painting in each plane according to an arbitrary user defined waveform. The aim of the injection painting is two-fold: First, to optimize the beam distribution for the target requirements, and second, to reduce the foil traversals and resulting injection beam loss. The impact of the latter goal can be easily demonstrating by accumulating nearly identical beam distributions with and without injection painting, and comparing the resulting beam loss, as shown in Fig. 5. The beam loss for the case with no injection painting is as much as five times more than the case with painting.




Figure 5: Top - Transverse profiles for painted (blue) and not painted (red) beams. Bottom - Injection area beam loss monitor signals for the above profiles.

The third and final large payoff investment is the ring collimation system. The collimation system is a two stage system with a set of primary scrapers and three secondary absorbers. The acceptance of the collimation system is ~300 pi mm mrad, significantly below the 480 pi acceptance of the remainder of the ring, ensuring that most particles far outside of the beam core intercept the collimation system instead of the beam pipe. While the primary scraper system has never been used during production, mainly because it is not needed, the secondary absorbers

are critical and are considered to be the main reason that the remainder of the ring has very low activation.

Medium and Low Payoff Investments

Based on the experience of predecessor machines such as the Protons Storage Ring (PSR) at Los Alamos, one of the major concerns during the design of the SNS accumulator ring was beam loss due to collective effects such as space charge and instability. There was a significant concern over the possibility of an intensity-limiting electron proton (e-P) instability. As such, a number of pre-emptive measures were taken to prevent the instability for 1.4 MW operations. Most notably, the entire SNS vacuum chamber was TiN coated to reduce the secondary emission yield for electrons, and a 2nd harmonic RF stations was included in the baseline RF buncher design to allow control of the longitudinal beam distribution. In addition, both clearing electrodes and clearing solenoids were installed, and a instability feedback system was developed.

So far, there has not been any e-P instabilities during production style SNS beam operations. Trace levels of e-P activity are sometimes observed near the end of the beam accumulation, but it does not result in any observable beam loss. It is not possible to know if this success is due to the TiN coating or not. To date there has been no need for use of either the clearing electrodes or the suppression solenoids, and in fact the solenoids have never been powered up. The e-P instability has been observed during dedicated physics studies where the machine is configured specifically to excite the instability, and during these times the second harmonic has been shown to be a strong knob in extinguishing the instability. The feedback system, under development for the entire beam power ramp up, is now able to reliably damp the instability, as described in other works in these proceedings [8]. Along with the second harmonic RF, this represents a robust suite safety feature against the possibility of future e-P instabilities at higher beam powers.

While the e-P instability was the primary concern in the arena of collective effects, there was also significant attention dedicated to avoiding space charge induced beam resonances. The baseline lattice tunes of (6.23, 6.20) were carefully chosen to be in a resonance free zone, and another back up working point (6.4, 6.3) was carefully studied. In order to ensure that higher order resonances were avoided for all both tune sets, a set of four of sextupole families were installed in the ring to provide chromaticity correction and tune spread reduction, and two octupole corrector families for compensation of higher order resonances. Normal and skew sextupoles were also installed.

To date, neither the sextupoles nor the octupoles, nor any of the correctors of this order, have ever been used during a production beam run. While the main sextupoles are routinely used during beam physics studies, they have not been shown to reduce beam loss in the ring. Although it is possible that these magnets may become necessary for higher power beam operations in the future, for the time being, they are not nearly as critical as originally

anticipated and a significant cost savings could have been realized by excluding them from the baseline design.

Unanticipated Challenges

While the ring has operated mostly as planned from the beginning, a few unforeseen issues have arisen, mainly in the injection region.

The first issue results from a design change whose consequences were not fully appreciated, and a lack of sufficient modelling of the trajectories in the injection region. The result was that it was not possible to obtain good injection into the ring while simultaneously providing clean transport of the waste beam to the injection dump. Rectifying this problem required several modifications to the injection region in the first few years after turn on [9]. These modifications included, but were not limited to: A change in size and width of both the primary and secondary foil, and increase in the injection dump beamline aperture, an increase in the injection dump line septum magnet gap, and the installation of a new C-magnet in the injection dump line.

A final unexpected challenge in the injection region is associated with the convoy electrons, i.e., the electrons which are stripped from the H⁻. For the SNS 1.4 MW beam power, these electrons constitute 1.6 kW of beam power and hence need to be properly handled. The foil is located in a magnetic field such that the electrons spiral downward along the field lines and intercept an electron catcher at the bottom of the vacuum chamber. The catcher is designed to capture the electrons and prevent reflection and out-scatter. However, due to a combination of fabrication errors and due to the modifications in the injection region, the electron catcher is not nor has ever been in the correct position. As a result, a significant number of electrons are reflected back toward the foil and intercept the foil mounting bracket and the surrounding beam pipe. This has caused damaged to the brackets leading to progressive changes in both the bracket geometry and material to mitigate the damage [10]. In addition, the catcher itself is suffering significant damage from the electrons and will need to be redesigned in the future.

FUTURE CHALLENGES

The SNS facility is planning for a beam power upgrade from the current baseline 1.4 MW to 2.8 MW to accommodate a second target station [11]. To achieve the new beam power, the beam energy will be increase from 1.0 GeV to 1.3 GeV, and the ion beam current will have to be increase from ~35 mA to ~50 mA, which is challenging. The new parameters result in approximately the same space charge tune shift in the accumulator ring, such that space charge effects are not a major concern. However, due to its highly nonlinear nature and notoriously unpredictable behaviour, the e-P instability could still be an issue. In addition, the beam power increase will result in a significant increase in the foil temperatures ~ 300 K [9]. Because the current foil temperature is not known, and because the sublimation rate versus temperature curve has a large error bar, there is some concern over the sublima-

tion rate of foils at the higher power. An effort is underway to measure the current foil temperatures to more accurately predict the sublimation rate for 2.8 MW of beam power.

ACKNOWLEDGEMENT

ORNL is managed by UT-Battelle, LLC, under contract DE-AC05-00OR22725 for the U.S. Department of Energy. This research was supported by the DOE Office of Science, Basic Energy Science, Scientific User Facilities.

REFERENCES

- [1] G. Dodson, "SNS Commissioning and Operations, the first 10 Years. An Overview of the Components Status after High Intensity Beam Operating Experience", presented at HB2016, Malmö, Sweden, July 2016, paper THAM3X01, this conference.
- [2] S. Kim, "SNS Superconducting Linac Operational Experience and Upgrade Path", in *Proc. LINAC08*, Victoria, USA, 2008, paper MO103, pp. 11.
- [3] A. Shishlo, "Model Benchmark With Experiment at the SNS Linac", presented at HB2016, Malmö, Sweden, July 2016, paper WEPM2Y01, this conference.
- [4] S. Nath *et al.*, "Longitudinal Beam-dynamics of the SNS SRF-Linac", in *Proc EPAC02*, Paris, 2002, paper THPLE024, pp. 1031.
- [5] A. Shishlo, A. Aleksandrov, V. Lebedev, and M. Plum, *Phys. Rev Lett*, 108, 114801 (2012).
- [6] A. Aleksandrov, "Path to Beam Loss Reduction in the SNS Linac Using Measurements, Simulation and Collimation", presented at HB2016, Malmö, Sweden, July 2016, paper THAM5Y01, this conference.
- [7] S. Henderson *et al*, *Nucl. Inst. and Methods*, 763, 2014.
- [8] N. Evans, "Broadband Feedback System for Instability Damping in the SNS Ring", presented a HB2016, Malmö, Sweden, July 2016, paper TUPM1X01, this conference.
- [9] M. Plum, "SNS Injection and Extraction Systems--Issues and Solutions", in *Proc. HB2008*, Nashville, 2008, paper WGC04, pp. 268.
- [10] M. Plum, "H⁻ Charge Exchange Injection Issues at High Power", presented at HB2016, Malmö, Sweden, July 2016, paper TUPM6X01, this conference.
- [11] Technical Design Report Second Target Station, ORNL/TM-2015/24, 2015.