Oak Ridge Spallation Neutron Source superconducting rf linac availability performance and demonstration of operation restoration with superconducting rf cavity off

V. S. Morozov⊙,[*] C. C. Peters⊙, and A. P. Shishlo⊙

Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831, USA

Recent advances in superconducting rf (SRF) technology allow for increasingly more reliable and higher power hadron linacs. They hold enormous potential for economic, scientific, medical, environmental, and national security fields. Moving this technology from one of a kind laboratory demonstration devices to industrial scale applications imposes strict requirements on the linac performance in terms of its stability, reliability, and availability. The Spallation Neutron Source (SNS) is currently the highest average power superconducting proton linac in the world that has been routinely operated since 2006. We present statistical data on the SNS SRF linac reliability obtained from its practical operational experience during 2020. We analyze the frequency and duration of SRF cavity trips and identify their causes. These data will show that SRF cavity trips are the most common source of single-point failure within the linac. In an industrial environment, where redundancy is a necessity to prevent interruption of facility operation, the linac beam can be quickly restored by redistributing the failed cavity function to downstream energy reserve SRF cavities. We present a practical demonstration of this approach. We intentionally turn off one of the cavities and recover the beam by readjusting the phases of downstream cavities to maintain the same linac beam energy output. We discuss limitations of this approach at SNS and how they can be overcome.

I. INTRODUCTION

High-power proton linac technology has been rapidly developing in recent years as several high-power linacs have been operating and are being built and many more are being designed and proposed around the world [1–10]. These linacs are currently used primarily as research devices. At the same time, they hold great promise for economic applications from medical isotope production to driving of subcritical reactors. For example, coupled with a subcritical reactor, they can provide solutions to reprocessing of spent nuclear fuel and clean energy generation.

One of the main challenges for applying this demonstration technology on the industrial scale is the requirement of nearly interruption-free operation due to economic considerations. In addition, accelerator-driven subcritical reactor (ADSR) applications [11–15] impose a technological constraint limiting the rate of beam trips that the reactor core can tolerate. Exceptionally high stability of the accelerator beam is necessary to prevent fatigue

development and eventual failure of the reactor components. For example, it is estimated that, for industrial scale power generation, the rate of beam interruptions for longer than 5 min should not exceed 3/year [11]. In this paper, we review the existing operational experience with the beam stability in the linac of the Spallation Neutron Source (SNS). We statistically analyze the causes of beam trips. We then present a proof-of-principle demonstration of retuning the linac "on the fly" in case of an SRF component failure and discuss its practical implementation.

We want to note that even before the SNS was built, an extensive amount of work has been done on the analysis and automated optimization of performance of ultrarelativistic electron linacs such as CEBAF [16–18]. However, in contrast to electron machines, the tuning process of a hadron linac must consider the beam velocity variation along the linac and the corresponding timing profile and velocity-dependent transit time factor for cavities. One must also pay close attention to space charge, longitudinal and transverse beam dynamics, and beam losses.

II. FACILITY OVERVIEW

Figure 1 shows an overall layout of the SNS facility [2]. Its main components are a 1 GeV H^- linac, a 248 m fixed-field accumulator ring with charge-exchange injection, a liquid mercury spallation target and connecting beam lines.
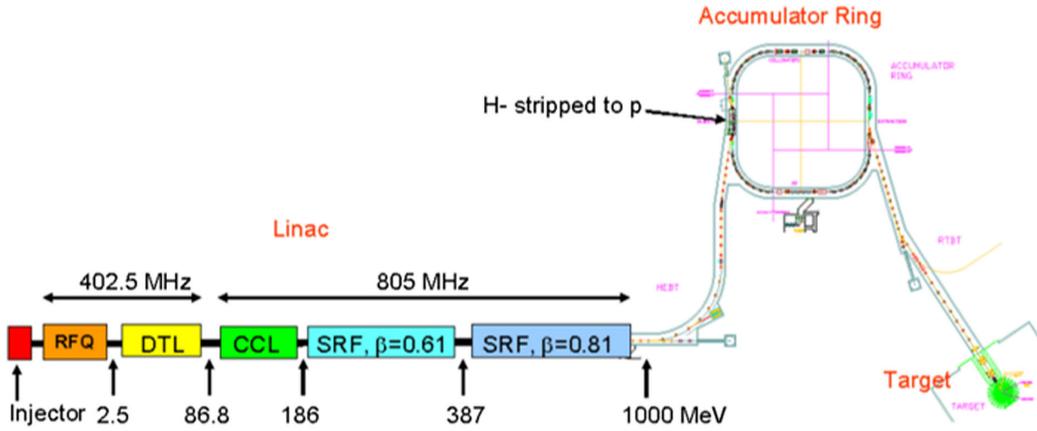
[*]morozovvs@ornl.gov

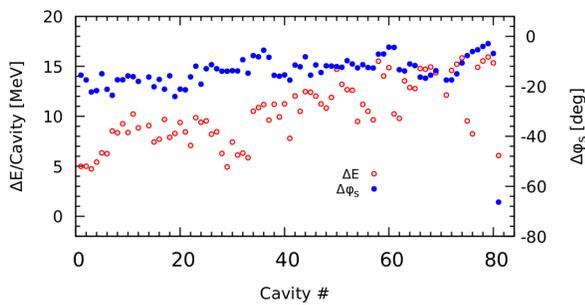FIG. 1.   Layout of the Spallation Neutron Source facility.

FIG. 2.   Typical distributions of the SCL cavity energy gains (left-hand-side vertical axis) and synchronous phases (right-hand-side vertical axis).

In this paper, we focus on the linac portion. The front end and warm linac accelerate to 186 MeV. It consists of an ion source, a 402.5 MHz radio-frequency quadrupole (RFQ), a medium-energy beam transport section (MEBT), a 402.5 MHz drift-tube linac (DTL), and an 805 MHz coupled cavity linac (CCL).

The warm linac is followed by the medium-β and high-β superconducting radio-frequency (SRF) cavity sections.

The medium-β and high-β SRF sections consist of 6-cell 805 MHz elliptical cavities optimized for relativistic β's of 0.61 and 0.81, respectively. The medium-β section is made up of 11 cryomodules each housing 3 cavities and accelerates the beam to 387 MeV. The high-β section consists of 12 cryomodules each containing 4 cavities and completes acceleration to 1 GeV. The average real-estate accelerating gradient of the superconducting linac (SCL) is 13 MV/m. Typical distributions of the SCL cavity energy gains and synchronous phases are shown in Fig. 2. The cavities are numbered in the ascending energy order. Cavities #13 and 70 were off due to hardware issues and are not included in Fig. 2. The main beam and linac operation parameters are summarized in Table I.

The SNS linac has been reliably delivering a MW-level proton beam since its commissioning in 2006. Figure 3 shows the time evolution of the accumulated energy and power on the target. As one can see, the production level of these energy parameters stabilized and have been staying constant over the last several years. The SNS typically operates in the production mode 5,000 hours per year. Additional 630 hours each year are devoted to machine development. The low-power data in Fig. 3 correspond primarily to times of target reliability issues. The SNS linac is currently the highest-average-power super-conducting

TABLE I.   SNS beam and linac parameters.

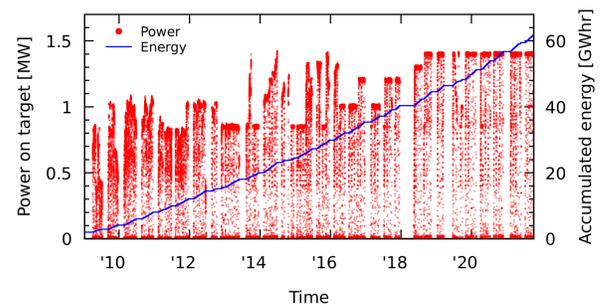| Parameter	Value
Top kinetic energy	1 GeV
Repetition rate	60 Hz
Pulse length	1 ms
Pulse current	26 mA
Average energy	1.4 MW
Bunch rate in pulse	402.5 MHz
SCL frequency	805 MHz
SCL cavities	
$\beta = 0.61$	11 cryomodules ×3 cavities
$\beta = 0.81$	12 cryomodules ×4 cavities

FIG. 3.   Time evolution of the accumulated energy (right-hand side vertical axis) and instant power on the target (left-hand side vertical axis).

linear proton accelerator in the world. The proton power upgrade (PPU) program is being implemented to advance its power even further. During the SNS operation, a wealth of experience and data have been accumulated on the practical aspects of running a high-power proton linac. Below in this paper, we statistically analyze and present some of those data.

III. STATISTICAL AVAILABILITY DATA

The SNS operated for about 5,000 hours during fiscal year 2020 (FY20). Figure 4 provides a weekly summary of the neutron beam availability for that time period. As one can see, the availability stayed above 90% every week and was on average about 95% for the entire year. The total downtime amounted to about 250 hours. The SCL accounted for about 100 hours of the downtime.

Figure 5 shows the breakdown of the SCL downtime by the category of its cause. The downtime is dominated by SCL cavity trips, which are responsible for almost 40% of the total downtime. These trips are hard to predict and prevent. Therefore, below we focus on analysis of this
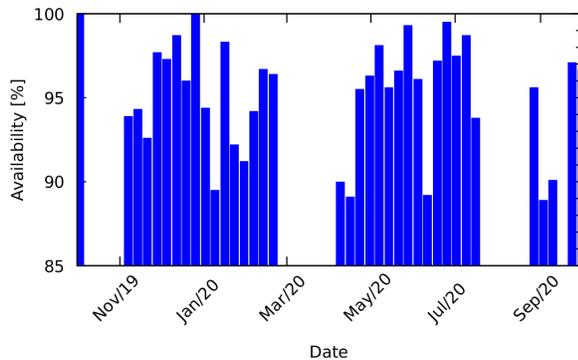

FIG. 4.   Weekly neutron beam availability at the SNS during FY20.
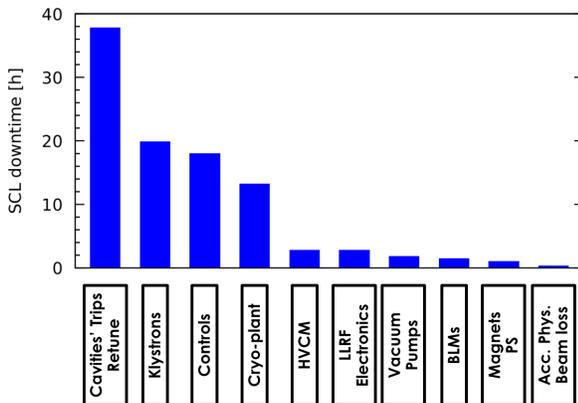

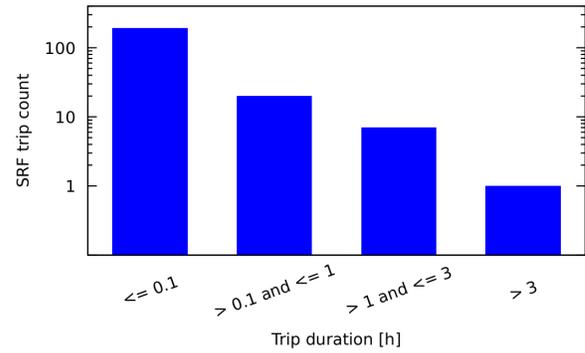FIG. 5.   Break down of the SCL downtime in FY20 by the subsystem of its origin.

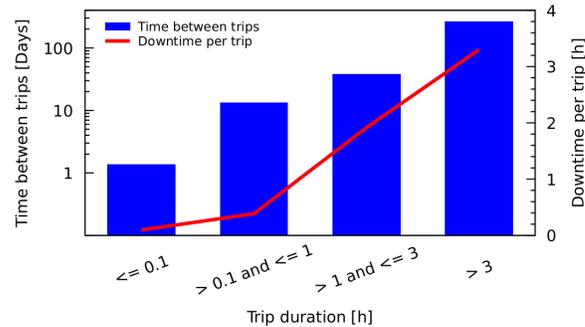FIG. 6.   Distribution of SRF trips in FY20 by their duration.

FIG. 7.   Average time between SRF trips (left-hand side vertical axis) and average downtime per trip (right-hand side vertical scale) in the indicated trip duration categories in FY20.

particular source of the downtime and approaches for recovering from the SRF trip consequences.

The impact of a trip can be judged by the amount of time it takes one to recover from it. The trips in FY20 are categorized by the duration of the recovery time in Fig. 6. As shown in the figure, the majority of the trips can be recovered from in a few minutes. One tenth of the trips can be handled in under an hour. Only about 5% of the trips require serious intervention taking over an hour and are usually related to a need to reduce an SRF cavity gradient and rephase downstream cavities to maintain energy
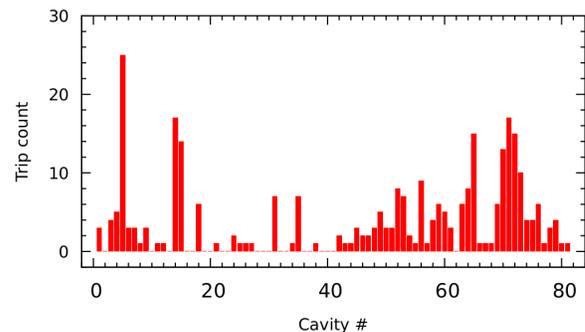

FIG. 8.   Numbers of trips of the individual SRF cavities in FY20.

output. Other statistical metrics of the SRF cavity reliability are the average time between the trips and the average downtime for each recovery time range of Fig. 6. These data are presented by Fig. 7. As one can see, short trips of less than a few minutes happen less than once in a day while long interruptions occur less than once in one hundred days.

Individual performance of SRF cavities varies from cavity to cavity. The FY20 trip statistics for each of the 81 SNS SRF cavities is shown in Fig. 8. The exact causes of the observed trip rate variation are not known. A machine learning approach may be appropriate for optimization of the cavity performance.

IV. TUNING AROUND A DISABLED CAVITY

In case of an SRF linac cavity trip in an industrial environment, one would like to minimize the amount of downtime caused by such an event to well below a minute or even a second. It is conceptually and technically possible. It has been suggested that, in case of a temporary cavity loss, operation can be quickly restored by redistributing the lost energy gain among the remaining cavities. Such an approach requires a sufficient hardware redundancy and an appropriate overhead in the accelerating voltage and power. Clearly, the greater the overhead the lower the chance of interruption of facility operation. This approach does not pose any conceptual issues. In this section, we discuss practical considerations and describe our experience with implementation of this approach at the SNS SRF linac.

For technical reasons such as transmission line and electronics delay, the phase setting of an SRF cavity in the control system of the SNS linac is offset with respect to the actual rf phase of the cavity. Thus, the control system is first calibrated against the rf phase by measuring the beam energy downstream of the cavity as a function of the cavity phase setting [19,20].

The cavities are calibrated one by one in the order of increasing beam energy from the start to the end of the linac. When a particular cavity is being calibrated, all of the cavities upstream of it are running at their production settings established earlier while acceleration by all of the downstream cavities is disabled by blanking their rf input pulses at the time of the beam passage. The calibration procedure uses a 1 μs pulse of 5 mA beam current at 1 Hz to avoid potential machine damage and prevent significant excitation of the blanked cavities. Thus, the beam energy remains constant downstream of the cavity that is being calibrated.

The downstream beam energy is determined using the time of flight (ToF) difference between two appropriately-selected downstream beam position and phase monitors (BPMs) [19,20]. There is a BPM after each SCL cryomodule along with several BPMs in the transport lines
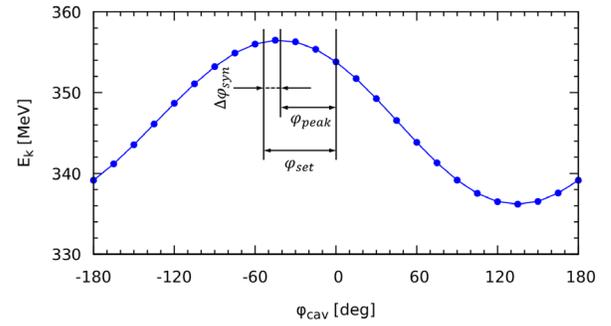

FIG. 9. Example of an SRF cavity phase calibration. The beam energy measured downstream of the cavity is plotted versus the phase setting of the input rf signal.

before and after the SCL. Different BPM pairs are used for calibration of different cavities. A BPM pair is selected for each cavity on the basis of minimizing the calibration error. The ToF technique provides a sufficiently accurate energy measurement in the SNS beam β range [21]. This method is not applicable to ultrarelativistic beams when the dispersion must be used but it is preferable at the SNS where the first dispersive section that could be used for such a measurement is in the HEBT. The HEBT is set for the full beam energy and does not provide an adequate energy acceptance for calibrating or optimizing the SRF cavities.

An example of the measured beam energy as a function of the cavity phase setting is shown in Fig. 9. The phase setting is varied from $-180°$ to $+180°$ in $15°$ steps. Such a dataset is obtained for each of the 81 cavities in the order described above. The data for each cavity are fit to a harmonic function to extract the location of the energy maximum as illustrated in Fig. 9. The synchronous phase $\Delta\varphi_{\mathrm{syn}}$ of the cavity is then set with respect to the maximum energy phase as shown in Fig. 9. The example data in Fig. 9 was taken for the 23rd cavity of the linac located at 45.3 m from the start of the linac. The energy was measured using the BPMs located at 54.3 and 137.0 m.

The entire procedure of phasing of all of the 81 SRF cavities takes about 40 min with $15°$ phase steps. This procedure can be sped up by taking wider steps in phase and would still provide an accurate calibration as can be seen from Fig. 9. The cavity synchronous phases are heuristically optimized to minimize the beam loss along the linac resulting in a synchronous phase setup shown in Fig. 2. The data collected during the calibration procedure are used to update the SCL model [22]. Note that the physics program at the SNS dictates very low, ultimately zero trip rates. The SNS SC cavities are operated with a sufficient safety margin so that trips are minimized. Therefore, rather than focusing directly on the trip rates, the SNS linac optimization is based on minimization of the beam loss along the linac, which in turn allows for more reliable SC cavity operation.
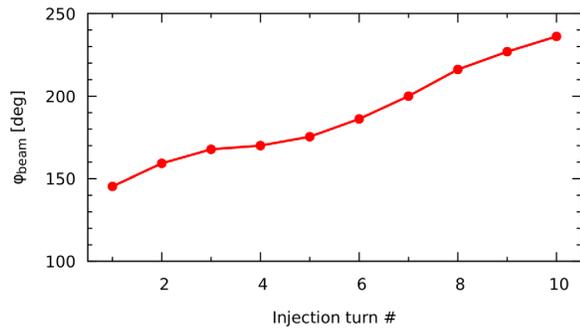
FIG. 10. Beam phase with respect to a reference rf signal measured in HEBT as a function of the injection turn number into the accumulator ring. The disabled SRF cavities are kept in resonance with the beam. Each point is the beam phase averaged over a time period corresponding to a single injection turn.

An unpowered cavity in a high-power beam line acts as a source of impedance. The front of a beam pulse passing through such a cavity excites its fundamental mode, which, in turn, adversely affects the tail of the pulse. This is a significant effect with the SNS production beam. Figure 10 shows the phase slew along the SNS 10 μs production beam pulse caused by 31 disabled cavities. The phase variation is related to change in the beam timing and therefore its energy. To prevent this effect from happening, the SNS SC linac employs the usual technique of detuning the disabled cavities sufficiently far from the resonant frequency. Due to a narrow width of the fundamental harmonic of the SNS SRF cavities, a frequency shift of about 20 kHz completely prevents its excitation by the beam.

In addition to measuring the bunch phase as shown in Fig. 10, the SCL energy deviation can be precisely determined by observing the orbit in the SNS accumulator ring. The dispersion in the ring results in perturbation of the closed orbit proportional to the offset in the SCL energy. Figure 11 compares the closed orbit distortions measured in the accumulator ring after injection with the frequency of the disabled 31 cavities offset by 7 and 20 kHz. The 20 kHz
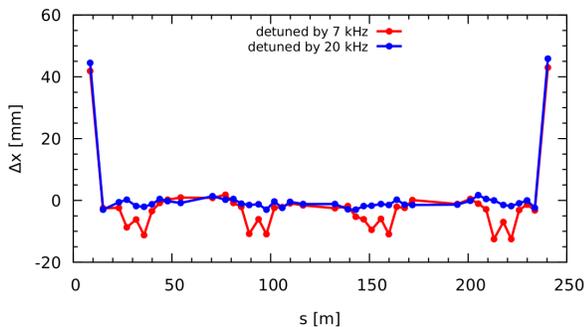

FIG. 11. Horizontal closed orbit distortions measured in the accumulator ring when the beam is injected with the frequency of the disabled SCL cavities shifted by 7 kHz (red) and 20 kHz (blue).
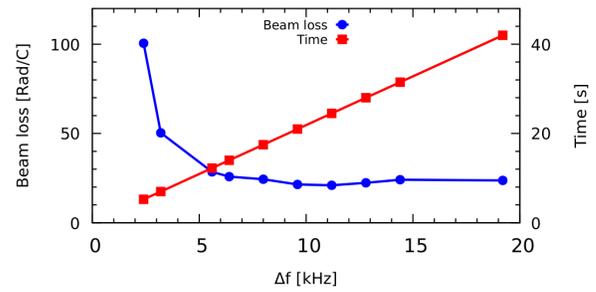

FIG. 12. Beam loss (shown in blue on the left-hand-side vertical axis) as a function of the frequency shift of the disabled SRF cavity. The time to move the cavity frequency (shown in red on the right-hand-side vertical axis) as a function of the amount of the frequency shift.

frequency offset restores an orbit alignment consistent with regular operation at the nominal energy when all SCL cavities are enabled.

Another parameter indicating whether the frequency shift is sufficient is the beam loss level. Figure 12 shows the change in the measured beam loss during the process of shifting the frequency. Such a coarse control of the cavity frequencies in the SNS linac is accomplished by mechanically moving tuning rods. There is no need at the SNS for quick changes in the cavity frequencies involving the tuning rods. Thus, for technical reasons, the rods move relatively slowly at 0.46 kHz per second. As shown in Fig. 12, the fastest the beam can be recovered safely based on beam loss is about 10 seconds. Similarly to what has been earlier established at CEBAF, this is the main technical limitation at the SNS on how quickly the beam can be restored to the original energy with one of the cavities off.

In practice, it is often faster to recovery a cavity after a trip than to retune the linac for running without it. In addition, the accelerating voltage overhead at the SNS linac is limited to a single SRF cavity, i.e., it is not possible to reach the full energy with more than one cavity off. Another caveat is that adjustment of the cavity voltages and phases to account for the missing cavity changes the beam loading of the active cavities. This change must be compensated by tuning the input rf power profile. In the SNS linac, the input power to each cavity is controlled by the adaptive feed forward (AFF) coefficients. The new AFF coefficients can currently only be determined experimentally with the full-power beam. In the future, we consider automating the AFF optimization procedure using machine learning. It should minimize the time this procedure takes. The aforementioned technical issues can be overcome with proper design and do not affect the conclusions of our conceptual demonstration described below.

The goal of our experiment was a conceptual demonstration of the idea that, in case of a cavity failure, linac operation can be promptly restored at the original energy by redistributing the lost energy gain among the remaining

cavities. The experimental procedure consisted of the following steps. (1) The linac SRF setup was documented with all cavities enabled. To simplify the later retuning procedure, the last SCL cavity is kept as a reserve. It is fully powered but, as shown in Fig. 2, has a large synchronous phase near the voltage zero crossing and provides little acceleration. (2) The input rf power to Cavity #41 (4th cavity of 13th cryomodule) was switched off and the cavity frequency was shifted by 20 kHz from the beam frequency. (3) The SCL model updated on the basis of the earlier cavity calibration procedure is used to adjust the phase settings of the rest of the cavities so that the SCL provides the original beam energy. In this process, the control system settings for the rf input phases are adjusted to keep the synchronous phases of all cavities nearly constant except for the reserved one. The synchronous phase of the reserved cavity is adjusted to compensate for the energy reduction due to the disabled cavity. (4) The AFF coefficients were optimized to bring the beam loss down to a level consistent with the original setup. (5) The rf phases and AFF coefficients of the new linac setup were documented. (6) The original linac state with all of the cavities enabled was restored. (7) The actual test was performed. A high-power beam was established, the test cavity was switched off and detuned, and the corresponding linac configuration prepared earlier was restored. It was then verified that the high-power beam was successfully recovered at the original energy. The time necessary to complete the entire process was dominated by the cavity tuner movement time.

Keeping the synchronous phase constant, as discussed above, results in nearly the same energy gains per cavity. However, since the input energies of the downstream cavities change, their energy gains shift slightly due to the energy-dependence of their transit time factors (TTFs). The last cavity that was playing the role of a hot spare, finally compensates the net energy change and restores the original beam energy.

The bunches arrive at each of the downstream cavities at a reduced energy and therefore at a time different from those in regular operation. In the described retuning
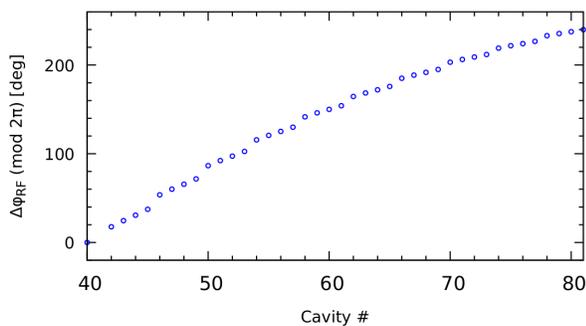

FIG. 13. Change in the control system input rf phase settings relative to those in regular operation for the SRF cavities downstream of the disabled cavity.
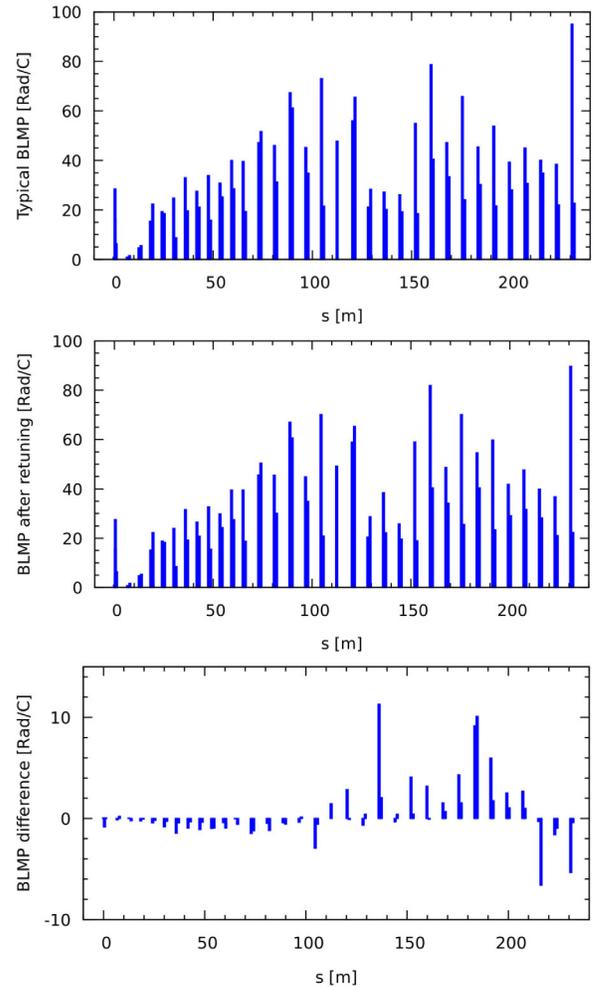

FIG. 14. Beam loss monitor patterns (BLMPs) during regular operation (top) and after retuning with the SRF cavity off (middle). The bottom graph shows the difference of the two patterns.

procedure, we chose to keep the same synchronous phases of the downstream cavities to maintain their energy gains approximately constant barring the effect associated with the TTF energy dependence. That was achieved by adjusting the rf input phases of the downstream cavities as shown in Fig. 13 to account for the change in the bunch arrival time. The beam loss monitor patterns (BLMPs) corresponding to the original and retuned SCL configurations are shown in Figs. 14 (top) and 14 (middle), respectively. The relatively small amount of difference between them shown in Fig. 14 (bottom) confirms that safe SCL operation consistent with the original setup has been fully restored.

V. CONCLUSIONS

The SNS has been routinely and reliably operating for the past several years at the design beam energy of 1 GeV and power of 1.4 MW. To date, this is the highest pulsed power superconducting proton linac in the world. It paves

the way for upcoming and future even more ambitious high-power linac projects. In FY20, the weekly SNS availability was consistently above 90% with a year average of 95%. This is a remarkable result given that the linac energy margin corresponds to a single SRF cavity. The majority of the linac downtime can be attributed to SRF cavity trips. 85% of them have been recovered from in less than 6 minutes by operator staff. The timing between significant events of more than 3 hours exceeds 100 days. This experience suggests that high-power hadron SRF linacs can be operated reliably and encourages their industrial application.

We successfully demonstrated that there are no conceptual problems with retuning an SRF linac to the original energy in case of a cavity failure. The energy overhead in the SNS linac corresponds to a single SRF cavity. Therefore, application of this procedure is the most beneficial when a serious failure occurs and the potential downtime is substantial. In other circumstances, it is more practical to recover the tripped cavity. Proper diagnostics techniques must be developed to identify potential long-downtime events. In case, when exceptional reliability and availability are needed, such as supply of an accelerator beam to an ADS reactor, our experiment supports the concept that the required performance level can be reached with appropriate hardware redundancy and energy margin.

Currently, the shortest time of restoring the SNS SCL from an SRF cavity failure is limited to 10 sec and is dominated by the time it takes to shift the disabled cavity's frequency sufficiently far from a harmonic of the bunch frequency. Thus, to overcome this limit in the future, techniques for faster detuning of an SRF cavity must be developed.

It is not practical to prepare in advance SC linac configurations for all possible combination of the missing cavities. Even at the SNS SCL consisting of 81 cavities and having one-cavity energy margin, this would correspond to 81 pre-saved linac states that would take an impractically long time to prepare and would lose their validity with any change of the cavity calibration. The calibration changes occur on the scale of a week due to modification and maintenance of the linac hardware. It becomes virtually impossible with a large hardware redundancy and a greatly increased number of setup scenarios. For example, with a two-cavity energy overhead, the number of potential configurations would be $81 \cdot 80 = 6480$. Therefore, the SNS linac model needs to be capable of real-time machine optimization and reconfiguration. It currently accounts for the most significant setup parameters such as the transit time factors, accelerating voltages, synchronous phases, optics, and space charge. Its further development needs to include additional parameters in the optimization process such as cavity loading, beam loss, and the trip rates and voltage performance of the individual cavities. Modeling

these parameters presents a challenge and may be best approached using machine learning techniques. It is one of our further research directions.

ACKNOWLEDGMENTS

[1] P. Ostroumov and F. Gerigk, Superconducting hadron linacs, Rev. Accel. Sci. Techol. **06,** 171 (2013).

[2] S. Henderson *et al.*, The Spallation Neutron Source accelerator system design, Nucl. Instrum. Methods Phys. Res., Sect. A, **763,** 610 (2014).

[3] J. Wei *et al.*, The FRIB superconducting linac: status and plans, *Proceedings of LINAC'16, 2016, East Lansing, MI* (JACoW, Geneva, 2016), MO1A01.

[4] K. Andersen *et al.*, ESS conceptual design report, Report No ESS-2012-001, 2012.

[5] R. W. Garnett *et al.*, High-power options for LANSCE, *Proceedings of NAPAC'11, 2011, New York, NY* (JACoW, Geneva, 2011), THOCN4.

[6] N. Solyak *et al.*, Physics design of the Project X CW linac, *Proceedings of NAPAC'11, 2011, New York, NY* (JACoW, Geneva, 2011), MOP145.

[7] L. Arnaudon *et al.*, Linac4 technical design report, Report No CERN-AB-2006-084 ABP/RF, 2006.

[8] F. Gerigk *et al.*, Conceptual design of the low-power and high-power SPL, a superconducting H^- linac at CERN, Report No. CERN-2014-007, 2014.

[9] M. Seidel *et al.*, Production of a 1.3 MW proton beam at PSI, *Proceedings of IPAC'10, 2010, Kyoto, Japan* (JACoW, Geneva, 2010), TUYRA03.

[10] K. Hasegawa *et al.*, Status of J-PARC accelerators, *Proceedings of IPAC'13, 2013, Shanghai, China* (JACoW, Geneva, 2013), THPWO031.

[11] E. Pitcher *et al.*, Accelerator and target technology for accelerator driven transmutation and energy production, DOE white paper on ADS (2010).

[12] W. Pan, Overview of worldwide accelerators for ADS, *Proceedings of IPAC'14, 2014, Dresden, Germany* (JACoW, Geneva, 2014), FRXCB01.

[13] H. Podlech *et al.*, The MYRRHA project, *Proceedings of NAPAC'19, 2019, Lansing, MI* (JACoW, Geneva, 2019), THZBA2.

[14] S. Liu, Z. Wang, W. Chen, W. Dou, H. Jia, Y. Qin, Y. He, and H. Zhao, Commissioning of China ADS demo Linac

and baseline design of CiADS project, J. Phys. Conf. Ser. **1401**, 012009 (2020).

[15] B. Yee-Rendon *et al.*, Present status of the R&D of the superconducting linac for the JAEA-ADS, J. Phys. Soc. Jpn. Conf. Proc. **33**, 011043 (2021).

[16] J. R. Delayen *et al.*, Operational optimization of large-scale SRF accelerators, *Proceedings of the 1999 Particle Accelerator Conference (Cat. No. 99CH36366)* (IEEE, Piscataway, 1999), p. 940.

[17] C. E. Reece *et al.*, Dealing with a large installation of SRF cavities: characterizing limitations and exploiting operational flexibility, *Proceedings of the 1999 Workshop on RF Superconductivity, Santa Fe, NM* (JACoW, Geneva, 1999), WEP038, p. 524.

[18] B. Terzic, A. S. Hofler, C. J. Reeves, S. A. Khan, G. A. Krafft, J. Benesch, A. Freyberger, and D. Ranjan, Simultaneous optimization of the cavity heat load and trip rates in linacs using a genetic algorithm, Phys. Rev. ST Accel. Beams **17**, 101003 (2014).

[19] A. Shishlo, Model and beam based setup procedures for a high power hadron superconducting linac, *Proceedings of LINAC'14, 2014, Geneva, Switzerland* (JACoW, Geneva, 2014), MOIOC03.

[20] A. Shishlo *et al.*, Model benchmark with experiment at the SNS linac, *Proceedings of HB'16, 2016, Malmoo, Sweden* (CERN, Geneva, 2016), WEPM2Y01.

[21] J. C. Wong, A. Aleksandrov, S. Cousineau, T. Gorlov, Y. Liu, A. Rakhman, and A. Shishlo, Laser-assisted charge exchange as an atomic yardstick for proton beam energy measurement and phase probe calibration, Phys. Rev. Accel. Beams **24**, 032801 (2021).

[22] A. Shishlo *et al.*, The XAL infrastructure for high level control room applications, *Proceedings of ICAP'09, 2009, San Francisco, CA* (JACoW, Geneva, 2009), WE3IODN01.

[23] http://energy.gov/downloads/doe-public-access-plan.