ICANS-XIV
14th Meeting of the International Collaboration on
Advanced Neutron Sources
June 14-19, 1998
Starved Rock Lodge, Utica, Illinois, USA

Front End and Linac Developments for the Spallation Neutron Source*

Jose R. Alonso, for the SNS Collaboration
Oak Ridge National Laboratory
SNS – 104 Union Valley Road
Oak Ridge, TN 37831-8218

1. ABSTRACT

The Spallation Neutron Source (SNS) at the Oak Ridge National Laboratory will be the world-leading accelerator-based neutron-scattering research facility when it comes on line at the end of 2005 [1]. By providing 1-MW of beam power on a heavy-metal target in short (<1 μs) bursts of 1-GeV protons, intense neutron beams will be available at flux levels at least a factor of five over presently-operating spallation sources. This paper will address the requirements driving accelerator parameters for the SNS, and will discuss how the SNS design meets these requirements. As the ultimate power level for the SNS is expected to be significantly higher than 1 MW, the roadmap for achieving these higher powers and increased flexibility for users will also be addressed.

2. REQUIREMENTS

The chief advantage of accelerator-generated neutrons over those from a reactor is the ability to generate sharp "start" pulses of neutrons, to characterize velocity (hence wavelength) of the neutron by pure time-of-flight from the arrival of the proton pulse on the target. Parameters that need to be explored to set the baseline for accelerator performance are: efficiency of neutron production (conversion of proton beam power into neutron flux), the time-length of the proton pulse (how much contribution to the uncertainty in the "start" pulse can come from the accelerator pulse), and repetition rate.

2.1 Neutron Yield

The flux of neutrons emanating from the target depends on proton beam energy, total power deposited on the target, and the target material itself. As to the latter, common wisdom calls for a heavy metal, with the highest possible nuclear density commensurate with the ability to absorb the heat and radiation damage from intense proton bombardment. As is well know, these considerations have driven the SNS team to select liquid mercury as their target material [1]; this matter is discussed at substantial length in other parts of these proceedings.

A rather straightforward [2] relationship for neutron yield exists between beam energy and total power, as shown in Figure 1. The curve in this figure represents neutron yield versus beam energy, for a constant beam power. The curve shows neutron yields as being approximately equal for a 1 MW beam that could consist of, for instance, a 1 mA average beam current at 1 GeV, or a 0.5 mA, 2 GeV beam. As is seen, the curve is fairly flat, with an optimum range between 1 and 3 GeV. (For purposes of normalizing the curve, approximately 26 neutrons are produced per proton at 1 GeV on a

---

* This work is supported by the Director, Office of Science, Office of Basic Energy Sciences, of the US Department of Energy under Contract No. DE-AC05-96OR22464

mercury target.) Optimizing the choice of proton beam energy for the SNS is driven by cost and technology factors: the SNS team has selected a full-energy linac and an accumulator ring (AR) as the proton generator; cost optimization for this scenario drives design to a lower-energy, higher-current configuration. Were one to select a shorter linac and a rapid-cycling synchrotron (RCS) to accelerate the beam during each storage cycle, the better cost-performance tradeoff would occur for higher energies and lower currents. From the standpoint of neutron production alone, either configuration would be equally satisfactory. However, the SNS team strongly endorses the full-energy linac option in that it provides a better platform on which to build for higher powers, and although initially about 10-15% more costly than the RCS option, provides much lower technical risk towards achieving a reliable, user-friendly facility.

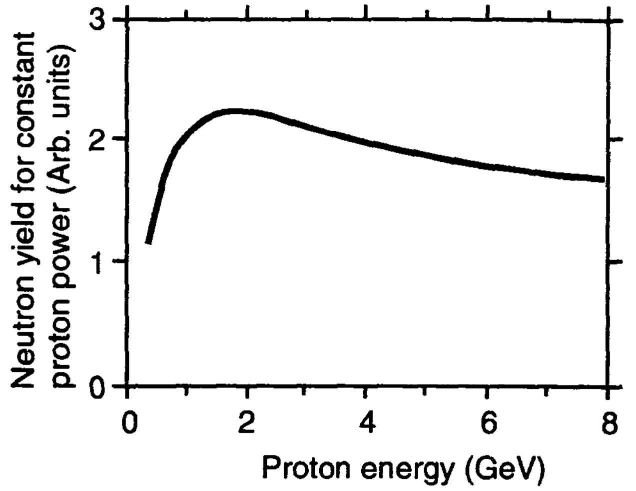

Figure 1. Neutron yield as a function of proton energy, for constant proton power deposition on target.

2.2 Timing Considerations

The sharpness of the "start" signal for the time-of-flight (TOF) measurement is dominated by the characteristics of the slow neutron pulse emerging from the moderators towards the measurement instruments. Although somewhere in the order of 10^9 reduction in neutron energy must be accomplished in the moderator (from MeV to meV), this takes place in a surprisingly small number of collisions, and a very small time period. Figure 2 shows calculations [3] for the full-width at half-maximum (listed as "W") of the neutron pulse emerging from an ambient (room-temperature) water moderator, assuming all neutrons entered the moderator at exactly the same time. The width of this pulse is thus a measure of the sharpness to which the neutron time-of-flight can be determined. While the sharpness of the pulse can be affected by coupling and poisoning the moderator (almost invariably a sharper pulse can be obtained at the expense of flux), natural widths for neutrons in the milli-electron volt range are in the tens-of-microseconds range. If an accelerator pulse were to be of the order of one microsecond, there would be very little contribution to the time-jitter from the width of the proton pulse on the target. It is not until one looks for neutrons at energies greater than a few 100 meV (e.g. a non-thermalized cut) that a 1 μs proton pulse might contribute to time uncertainty. Note that widths for cryogenic moderators are even larger, so again a 1 μs proton pulse contributes almost insignificantly to the time uncertainty for thermally-equilibrated neutrons from these moderators. General consensus is that a specification for proton beam width of "less than 1 μs" is appropriate. By a happy chance of nature, 1 μs is close to the revolution period of beam in an appropriately-sized 1 GeV accumulator ring, making this particular timing requirement on the beam a straightforward one to meet.

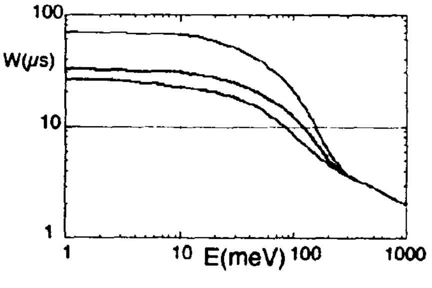

Figure 2. Time widths W (FWHM) of neutrons emerging from an example of a room-temperature water moderator, for (top-to-bottom) coupled, decoupled and decoupled-poisoned conditions.

An interesting constraint arises when considering the design of the overall timing system for the accelerator. The expected scenario for the SNS' pulse-repetition rate (discussed in the following section) will be a multiple or subharmonic of 60 Hz. Efficient operation of the accelerator power systems indicate it best to synchronize the pulses with the AC line frequency, triggering off of a zero-crossing of one of the AC line phases. For most experimental conditions this is quite satisfactory; a beam-current transformer located in the beamline just upstream of the target produces a signal when the proton pulse passes through it on its way to the target, providing a clean "start" trigger for the TOF measurement. However, for instruments measuring inelastic scattering, where an initial neutron velocity determination is made by measuring the time between the beam arriving on target and the opening of a Fermi chopper window close to the sample, phase errors between the Fermi chopper and the AC line can produce potentially-unacceptable variations in the beam velocity measurement. This is likely to occur since variations in both long-term, and pulse-to-pulse reproducibility of both chopper revolution frequency and a typical AC line trigger can be many times greater than the acceptable accuracy limits. Achieving needed timing accuracy then suggests that the master accelerator timing signal for the accelerator beam be derived from the most critical (if there are more than one in operation) Fermi chopper. A sophisticated chopper drive system is required to ensure that all operating choppers run as closely as possible to the same revolution phase, and that this phase closely tracks variations in the AC line. Rotation frequency of the choppers must be adjusted at a slew rate commensurate with the inertia and drive-system capabilities of these choppers.

2.3 Pulse Repetition Rate

The best repetition rate for proton pulses is closely connected to the type of measurements to be made, to the velocity and flight-path-length of the neutron beams, and to the requirement for total neutron flux from the target. For the accumulator ring concept, total proton beam-power in a single pulse is a very steep function of dollars. For instance, the SNS design calls for 17 kJ of energy to be delivered in a single pulse, and while delivering less than this is relatively easy to do, any increase over the 17 kJ limit would require very significant and quite costly design modifications to the accelerator system. The result of this is that a (1 GeV) 1 MW facility based on 17 kJ pulses at a 60 Hz rate is a very different animal from a (1 GeV) 1 MW system running at 30 Hz that would require 34 kJ pulses. While it would be very easy to run the SNS accelerator complex at 30 Hz, the total power delivered would only be 500 kW. The initial capital investment is driven primarily by the power in each pulse, increasing the number of pulses comes at only a modest cost. Not to say that one would wish to run at 120 Hz, but it does imply that achieving the highest possible overall facility power for the lowest cost is achieved by designing for the highest practical repetition rate.

The design decision then to run the SNS at 60 Hz is based on the fact that a significant portion of the user community can utilize neutrons arriving at this rate. Note that it is very cost-effective to add a second target station that can operate at a lower repetition rate in concert with the high-rep-rate target, providing more optimal service to those users requiring longer flight-paths or slower neutrons.

Such a second target station also doubles the number of available beam ports, in addition to allowing for customization of moderator configurations for specifically tailored neutron spectra on both targets. The second target station is, as a consequence, the very first item on the priority list for facility upgrades.

2.4 Reliability and Availability

An extremely important requirement is the highest-possible reliability and availability. Neutron-scattering experiments are by nature short, and are typically planned far in advance. As a result, it is most important that when an experimental group comes to the facility, neutrons are available as promised. In short, the SNS must be a facility built for users, to ensure maximum experimental throughput and productivity, and is not to be a proving ground for development of new accelerator frontiers.

Reliability results from several factors: a) use of well-developed, mature technology, b) conservative design parameters with adequate performance margins to ensure that systems are not being stressed unduly, c) high-quality components that will not be likely to fail, and d) a highly-trained and motivated operations staff.

All of these factors, unfortunately, add to the facility cost. One must be extremely careful, in the inevitable cost-optimization processes that go on throughout the life of the construction project, that economies are not taken which introduce unacceptable risk to the operational reliability of the facility.

One specific area sets the SNS apart from other user-based accelerator facilities; the beam-currents are so high that extreme care must be taken to minimize beam losses during the acceleration and accumulation processes. The consequences of beam loss are either catastrophic failure of components that might be hit, in the event of accidental gross mis-steering of the beam, or activation of components in the accelerator tunnels to the point where access must be limited, and maintenance of accelerator systems becomes very difficult. Both will adversely impact machine availability. Typical fractional beam losses in existing high-power proton accelerators are significantly higher than can be tolerated in the SNS, leading to significant challenges to the design teams to understand mechanisms for beam loss in these existing accelerators, and to provide designs that will perform better. The goal of our design is between one and two orders of magnitude lower fractional beam losses than are presently experienced at commensurate accelerator systems of LANSCE, AGS or Fermilab.

Table 1. Design Parameters for the SNS

Beam Species on Target	Protons
Beam Species in Linac	H^-
Proton Beam Energy	1 GeV
Beam Current (average)	1 mA
Linac Duty Factor	6%
Pulse Repetition Rate	60 Hz
Pulse Width on Target	≈600 ns
Particles per Pulse on Target	1×10^{14}
Instantaneous Beam Energy on Target	17 kJ
Instantaneous Power on Target	≈30 GW

3. SPECIFICATIONS

Translating the above requirements into performance specifications for the SNS accelerator, we derive the operating parameters listed in Table 1. The biggest challenge lies in the large number of particles that must be delivered to the target in the very short time period. The best way of achieving this is to use well-established RF linear-accelerator technology, which can easily handle tens-of-milliamps of current, and to use a pulse-compression system such as an accumulator ring to build the peak current up to required levels.

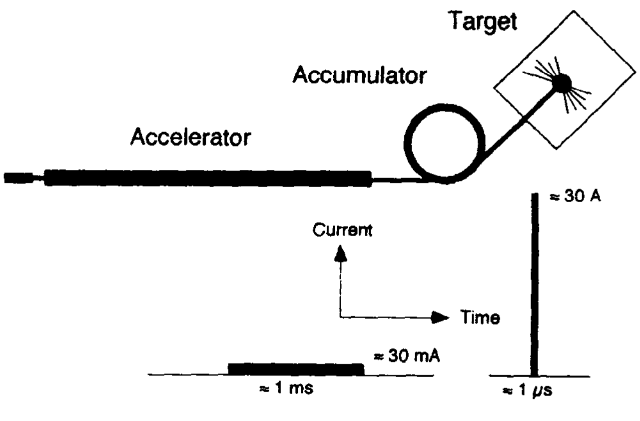

Figure 3. Configuration of a short-pulse spallation source, indicating the use of an accumulator ring to achieve the needed factor of 1,000 amplification in peak current.

The planned instantaneous current through the SNS RF linac structures is around 30 mA, driven both by ion-source performance at the emittance and brightness required, as well as by linac-structure sizes and costs. The instantaneous current delivered to the target, however, must be of the order of 30 amperes (to meet the criterion of 1 MW in 60 pulses of 1 μs duration). To produce this 30-A pulse, one wraps the beam from the linac many times around the circumference of the accumulator ring, then kicks this accumulated beam out in a single turn. By stacking 1000 turns into the ring, the required instantaneous power amplification is achieved. Figure 3 shows a schematic for this configuration; the SNS design calls for the number of particles stored in the ring, around 1 x 10^{14}, to enter the ring as a low-current, long string and to leave as a short, high-current pulse.

In the context of minimizing beam losses, one of the most critical areas is the injection and stacking into the accumulator ring. The most important element in this regard is the absolute necessity of accelerating H^- ions in the linac. Charge-changing, by passing these negative ions through a thin stripper foil, is the most effective way of injecting protons into a ring, and is now used almost exclusively in the world's high-energy proton accelerator facilities. Increases in phase-space density are possible with this technique that are not possible using positive-ion injection, as the circulating beam can pass through the stripper foil used to convert the negative hydrogen into protons (by removal of the two electrons), with little consequences to the overall beam quality. A magnetic field, bending opposite charges in opposite directions, merges the two beams to a common point at the stripper foil. As both beams have the same charge following the foil, their orbits are now coincident. This is shown schematically in Figure 4. While this technique has allowed an efficiency increase in stacked and captured particles from the 70% typical for optimized positive-ion multi-turn injection to well over 90%, the required efficiency for the SNS of >99.9% has required added refinements to the concepts. These will be discussed briefly below.

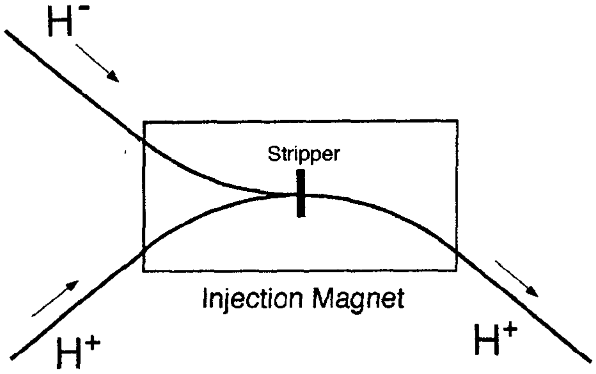

Figure 4. Schematic representation of charge-changing injection into a ring. The injection line (H^- beam) is bent upwards in the magnetic field, and merges with the circulating (H^+) beam that is bent downwards. The two orbits meet in the stripper foil. After the two electrons are removed from the incident ions, the two beams follow the same trajectory through the remainder of the magnet.

4. SNS BASELINE DESIGN

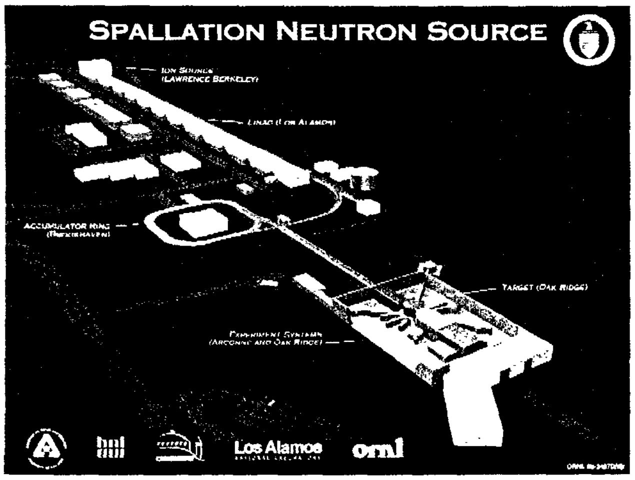

Figure 5. Schematic layout of the SNS Facility

The baseline design for the SNS is described in the Conceptual Design Report [1], and is shown in Figure 5. Although over a year old, the baseline design has changed little since this original study. The scope of this paper covers the Front End and Linac portions, only brief mention will be made of the components beyond the linac. The reader is referred to the Conceptual Design Report for in-depth descriptions of these components. In addition, ring issues (in the context of the LANSCE - PSR) are covered in R. Macek's paper in these proceedings.

4.1 Front End

Lawrence Berkeley National Laboratory is providing the “front end” of the SNS, consisting of components that produce the H⁻ beam and accelerate it to 2.5 MeV. Figure 6 shows a schematic of these components, including the ion source, LEBT (Low Energy Beam Transport), RFQ (Radio-Frequency Quadrupole) accelerator, and MEBT (Medium-Energy Beam Transport) line. The total length of the front end is around 10 meters. Although a very small fraction of the total system dimensions, the front end plays a crucial role in defining beam quality necessary to meet performance goals.

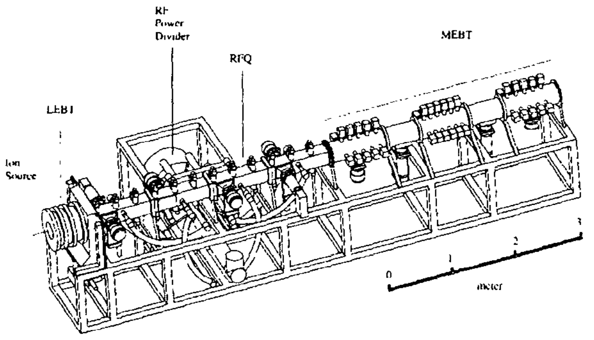

Figure 6. Schematic layout of Front End, showing ion source, LEBT, RFQ and MEBT

The ion source, a cesium-enhanced volume-production source is now a very mature technology that has been applied in many accelerators and other beam-forming applications. Plasma formed by a high-power (≈30 kw) RF discharge is contained by a multi-cusp magnetic field configuration of permanent magnets. Ions drift through a barrier field into a low ion-temperature region, where cesiated surfaces enhance the population of negative ions. Beam is extracted through a 6 mm diameter hole. A sophisticated magnetic-field configuration separates electrons extracted with the ions. Required peak beam current is 35 mA, this level has already been demonstrated in R&D sources operating without Cs. The introduction of Cs will increase current by about a factor of 3, allowing much lower RF discharge powers, significantly increasing source lifetimes and reliabilities.

The LEBT consists of a short (11 cm) column of electrostatic lenses which accelerate the extracted beam to 65 keV and match it to the RFQ. While traditional LEBTs are longer and employ solenoid-magnet focusing, the use of an all-electrostatic system eliminates problems associated with time-dependent space-charge neutralization that can arise with pulsed beams.

The RFQ, operating at 402.5 MHz, bunches the beam from the ion source into the packets that fit within the accelerating RF “buckets” (every 2.5 ns), and accelerate the beam to 2.5 MeV. Careful attention is paid to maximizing the efficiency of this bunching process, and ensuring that the emittance of the beam (its spatial and time characteristics) is preserved. This is particularly difficult because of the strong space-charge forces from the very high bunch densities that tend to drive the beam apart. RF quadrupoles have been long-used in mass spectrometry as strong-focusing transport channels, but by careful modulation of the shapes of the vane-tips (introducing “wiggles”) selective longitudinal components to the strong transverse electric fields can be introduced which first bunch, then accelerate the beam. Again, this is now a very mature technology, with almost all of today’s high-energy accelerators employing these structures for initial

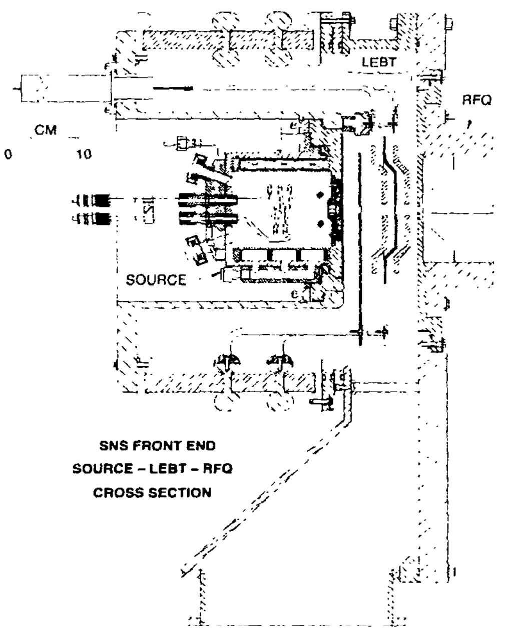

Figure 7. Cross section of ion source and LEBT. Large pumping port in base allows for handling of gas loads from ion source

stages of acceleration.

The MEBT serves as a 3.7 meter channel to transport the beam from the RFQ to the next stage of acceleration. While it is normally desired to close-couple each acceleration stage (butting the accelerating sections right up against each other), the SNS requires a sophisticated "beam-chopping" system that mandates the use of the long MEBT.

4.2 Beam Chopper

Three basic time scales can be identified for the beam in the linac: the 1 millisecond "macro-pulse" length providing all the particles that are compressed into the single very short pulse on the target, the RF bunches ("micro-pulses") that come every 2.5 nanoseconds, and the natural revolution frequency of the beam in the accumulator ring (AR), which introduces a time scale in the microsecond range.

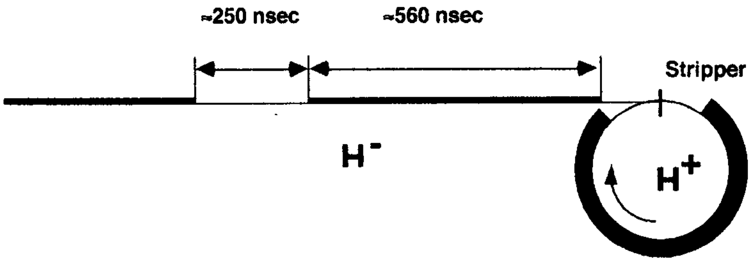

Figure 8. Schematic of beam chopper showing train of 560 ns pulses from linac stacked into accumulator ring.

As noted earlier, the 1 ms macro-pulse is wrapped into ≈1000 turns in the AR, and then is kicked out in a single turn. Thinking of the mechanics of the kicking process, which is performed by a fast "kicker magnet", one must ask how "fast" the fast kicker really can be. The best engineering designs, for a ring of the required aperture, indicate the quickest rise time could be around 200 nanoseconds. This is almost 25% of the ring revolution period. If any beam is in the kicker during the time the magnet is transitioning from "off" (beam circulating) to "on" (beam extracted), this beam will be sprayed into the region between the ring and the extraction channel. As the total beam power is 1 MW, this implies that approximately 250 kw of beam will be lost in this area, causing completely unacceptable activation and component damage.

To mitigate this problem, the linac beam is "chopped," placing 250 ns holes in the otherwise continuous macro-pulse that correspond to what is called the "extraction gap." By synchronizing the beam chopper with the ring circulation frequency, the overlap of all 1000 such holes is ensured. Chopping is performed by passing the beam through a set of parallel plates on which a high-voltage square-wave is placed. This pulse deflects the unwanted 250 ns of beam into a stopper, and allows the remaining 560 ns of beam to pass through undeflected. Most efficient chopping is performed at a low energy, where reasonable fields can cause adequate deflection of the beam, where beam power of the waste beam can be effectively absorbed, and where this waste beam will not produce neutrons. For these reasons, the 2.5 MeV point has been selected for beam chopping.

Several issues must be addressed: it is important that the rise-time of the square wave be very sharp, to separate adjacent RF micro-pulses. This would imply a better-than-2.5 ns rise-time. Any micro-pulse that is caught in the transition time of the square wave will be partially deflected, and at least a portion will miss the stopper, producing beam that might be transmitted off-center through the remaining accelerating components. These partially-chopped bunches might possibly lead to significant beam-loss problems at higher energies. For this reason, the MEBT is designed with an "anti-chopper" that is designed to bring any partially-chopped bunches back to the beam axis.

Another issue is that the chopper plates are about 50 cm in length, and the transit time of a 2.5 MeV beam through these plates is around 40 ns. It is important then that the plates be not continuous, but be segmented into a travelling-wave structure, so that the transitioning edges of the square wave can propagate down the length of the structure at the same velocity as the beam bunches. This complicates the design of the chopper, and increases the problem of maintaining the required very fast rise-time.

Nonetheless, a chopper with close to the required performance specifications has been in operation at LANSCE for many years, and Los Alamos will provide the SNS chopper system based on this experience.

4.3 Linacs

Three different linear-accelerator structures will be provided by Los Alamos National Laboratory for the main acceleration stages: a DTL, or Drift-Tube Linac brings the beam to 20 MeV, a CCDTL (Coupled-Cavity Drift-Tube Linac) takes the beamto 93 MeV, and a CCL (Coupled-Cavity Linac) completes the chain to 1 GeV. Many examples of DTL and CCL structures exist, designs are very well characterized and their performance is quite well understood. The CCDTL is a new concept, with excellent properties to provide gentle matching between the DTL and CCL structures, which produces, at least in simulations, very high quality beams. While no operating CCDTL exists, models have been constructed to verify field calculations, and prototypes are being built for both APT and SNS to ensure proper operating characteristics.

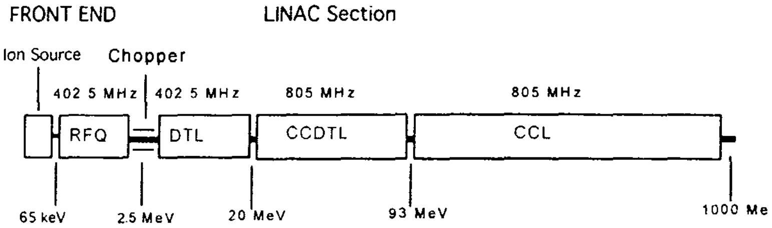

Figure 9. Schematic (not to scale!) of Front End and Linac elements.

The DTL, or Alvarez Linac, consists of a long tank (actually, the DTL is in two sections, so there will be two tanks) resonating at 402.5 MHz in a mode where the electric field parallels the beam direction. As the electric field switches directions every 2.5 ns, copper "drift tubes" are provided at appropriate intervals to shield the beam bunches from the adverse fields. Bunches emerge from the end of each tube to receive a kick before entering into the next tube. There are 84 drift tubes in the tanks, these tubes increase in length down the structure to match the increased beam velocity. Each tube also contains a quadrupole magnet (SmCo) to provide transverse beam focusing.

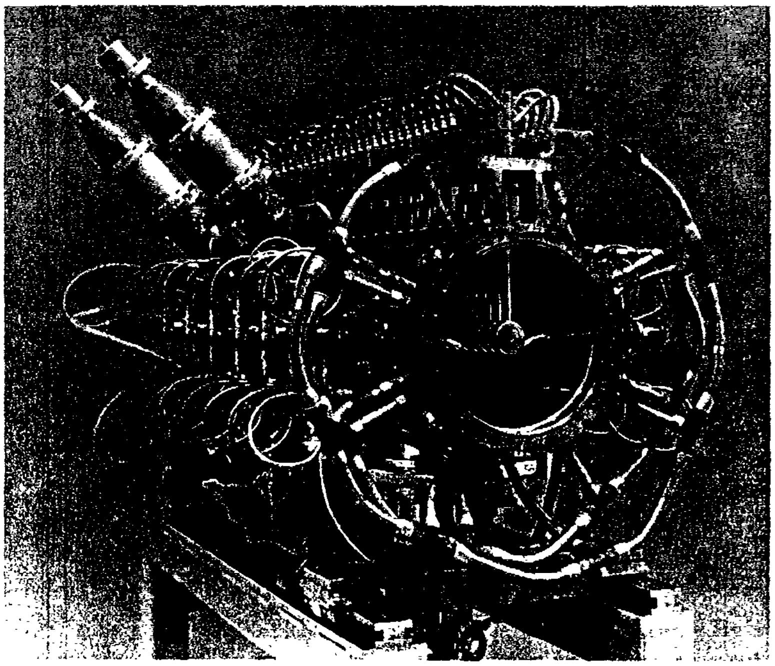

Figure 10. DTL structure built by LANL for the GTA program, very similar to design for SNS. Beam passes through small aperture in drift tubes.

The CCDTL structure is directly coupled to the end of the DTL, and operates at 805 MHz. The doubling of the operating frequency allows for more efficient size and shape of the accelerating structures, and as long as the designs are carefully done, the transition in frequency will have very little effect on the beam quality. It should be noted that LANSCE does not have this intermediate CCDTL structure, but instead transitions directly from a 200 MHz DTL to an 800 MHz CCL. Addition of the new structure, and reduction in the frequency jump factor from 4 to 2 are expected to produce much higher quality beams than are currently available from LANSCE.

After approximately 68 meters of CCDTL structure, at the 93 MeV point, the CCL structure takes over bringing the beam, after an additional 473 meters, to its full energy of 1 GeV. The transition between these two structures occurs seamlessly, ensuring very smooth transitions in both longitudinal and transverse forces on the beam.

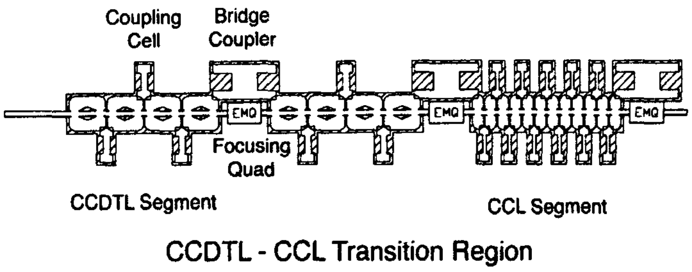

Figure 11. Schematic of CCDTL and CCL structures showing individual components

As seen schematically in Figure 11, the CCDTL and CCL structures externally look very similar: brazed segments approximately 1.5 meters in length each containing several accelerating cells, are mounted on strongbacks, with approximately 30 cm spacing between segments in which focusing magnets and beam diagnostics are placed. Segments along the 540 meter length of the two structures are grouped into 26 distinct RF "modules," each fed by two klystrons, the RF feed occurring in a bridge-coupler between two segments (at the 1/6 and 5/6 point of each module). RF energy is transferred between the individual cells by means of the coupling cells on the top or bottom of adjacent accelerating cells. Fields on the beam axis in the cells is parallel (or antiparallel) to the beam, phasing of the beam bunches is such that each sees an accelerating field as it traverses the cell and is in an inter-cell "bore tube" as the fields are reversing. Energy of the beam in the CCDTL is low enough that an additional drift tube must be inserted in each of the cells to ensure proper matching of structure geometry, resonant fields and beam velocity. Thus, each of the CCDTL cells will have two accelerating gaps.

4.4 Accumulator Ring, Target, Instruments

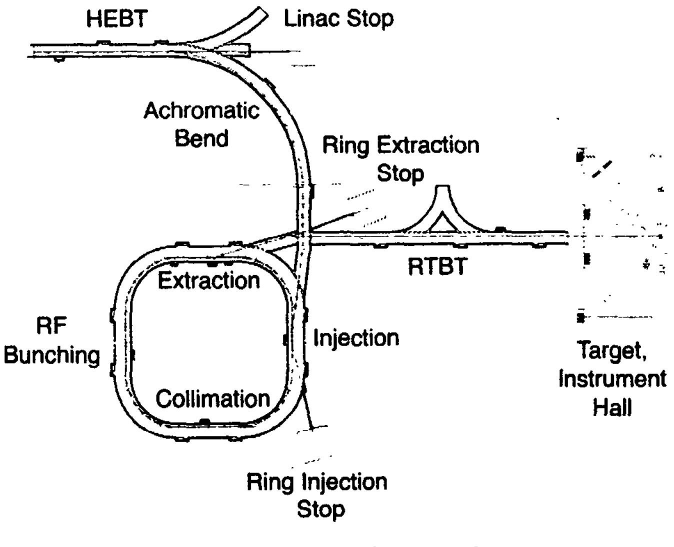

Figure 12. Simplified schematic of accumulator ring systems, including the HEBT transport from the *linac to* the *ring, and the RTBT between the ring and the target building.*

At the 1 GeV point, beam enters the HEBT (High-Energy Beam Transport), is bent through 90° to provide beam analysis and clean-up opportunities before reaching the ring injection point. An elaborate series of pulsed and fixed bump magnets provides flexibility for "painting" the 1000 turns as they are injected through the carbon stripper foil, to provide the best distribution of beam circulating in the ring. Optimization of this beam distribution can have significant impact on beam losses during the accumulation process, and also affects the distribution of beam density on the target, important for thermal and radiation-damage management on target components.

A fast kicker bumps the beam up vertically into a special septum magnet that provides horizontal deflection into the RTBT (Ring-to-Target Beam Transport), the beam bunch is now 590 nanoseconds long. This beam is transported to the target, where it is spread out into a relatively-uniform rectangular distribution roughly 20 cm wide by 7 cm high. Again, because the beam is kicked out in a single turn, distribution on the target is closely related to the distribution of beam inside the ring, and hence to the injection painting scheme.

These systems will all be provided by the Brookhaven National Laboratory.

The mercury target, provided by the Oak Ridge National Laboratory converts proton power to neutrons via spallation reactions; intra- and inter-nuclear cascades yield approximately 26 neutrons for every incident 1 GeV proton. Liquid mercury has been selected to mitigate thermal shock effects from the ≈30 GW proton pulses that might cause unacceptable damage to solid target materials. These pulse loads will be substantially higher as well, once the SNS is upgraded to its higher-power operating levels.

Two room-temperature water moderators, and two supercritical hydrogen cryogenic moderators will deliver neutron beams to 18 beam lines, and an initial complement of 10 instruments will be built under the supervision of the Argonne National Laboratory as part of the SNS project line-item scope.

5. CONTROLLING BEAM LOSS

As has been mentioned before, a most important element of the SNS design is controlling beam losses. Particularly at the higher energies, beam that is lost will cause activation, or even component damage that can have severely deleterious effects on operational reliability and efficiency of maintenance activities. Mentioned above as well was that design requirements for beam loss for many of the systems employed in the SNS surpass by a large factor the base of current experience. As a consequence, very careful attention has been paid to parameters and designs to maximize the chances for controlling beam loss.

A very large experience base has been collected from the operation of the LANSCE linac, which has beam power capabilities very similar to SNS'. Although the power actually delivered to the neutron-production target in the Lujan Center is only about 80 kw, the total beam power available from the LANSCE linac is of the order of 1 MW. Typical losses experienced are of the order of 1 na/meter, producing dose rates in the vicinity of linac components of the order of 30 mr/hr. This level does not preclude hands-on maintenance, provided exposure times are kept to the barest minimum. As a matter of course, good engineering practice calls for accelerator component designs that allow for accomplishing required maintenance activities as rapidly and efficiently as possible.

It is anticipated that the SNS will operate at significantly lower beam loss levels than LANCE's. Since the original design of LANSCE, substantial experience has been gained in understanding of beam-loss mechanisms, particularly in mechanisms for driving beam particles into the wings of the transverse distributions that can yield orbits for particles that go beyond the containing boundaries of the beam pipe. In particular, ensuring smooth transitions between different accelerating structures is predicted to very significantly reduce population of beam halos. The size and extent of beam halos are also much better understood now, and required structure apertures to contain most readily-populated halo trajectories can be calculated. As a result, the SNS apertures are significantly larger than LANSCE's, and in addition the beam diameter is quite a bit smaller.

One area of linac engineering design that must receive careful attention is the vacuum system. For the H^- ions, interactions with residual gas atoms can strip electrons from the ions, leading to beam loss. Base pressures in the 10^{-8} torr range are required to keep these losses within acceptable limits. Such levels are achievable with proper attention to cleanliness, and good pumping systems.

There are several areas in the ring where beam losses could present significant problems. As mentioned earlier, the injection system must perform at new levels of efficiency. As a first step to achieving this, the HEBT is designed to ensure ample opportunities for cleaning halos and outlying particles from the beam prior to reaching the injection foil. In addition, specialized RF cavities in the

HEBT line can provide fine adjustments in beam characteristics to improve injection matching. Good simulation codes are also now available for modeling the injection process.

Of particular note is the atomic physics associated with the stripping process. Even at 1 GeV, the removal of both electrons from the H^- ion is not a completely trivial matter. Foils must be very thin, resulting in an appreciable percentage of the beam remaining in excited H° states. The interaction of these excited states with downstream magnetic fields can lead to high beam losses through Lorentz stripping. This is the dominant mechanism for beam loss in the LANSCE PSR, and has been one of the factors limiting performance of this facility. As a specific example, at 1 GeV, a 200 $\mu g/cm^2$ carbon stripper foil can leave almost 8% of the beam in various un-stripped states. While a 400 $\mu g/cm^2$ foil reduces this fraction to well-below 1%, the operating temperature of this foil may reach over 3000°, causing concern for foil lifetime and reliability. A novel layout for the injection region has been developed which is particularly effective in mitigating Lorentz stripping effects, allowing the use of foils thin enough to where thermal effects can be controlled. This, coupled with the better understanding of the physics of injection processes and improved simulation codes, gives us confidence that the SNS injection efficiency will be within specifications.

There are many other mechanisms that can cause orbits for beam particles in the accumulator ring to grow beyond the available machine aperture, many of them as yet not fully analyzed. However, a properly-designed collimator system capable of absorbing such tails, can keep the errant particles from activating portions of the machine where routine access for maintenance activities is required. As seen in Figure 12, the SNS ring is designed with four straight sections, one dedicated to collimation systems. The aperture in this region is by design smaller than other portions of the machine, so that any particles with aberrant orbits will be intercepted in this region. As a result, it is expected that the required specification of 1 part in 10^4 of *uncontrolled* beam loss can be satisfactorily met in the ring.

6. UPGRADE PATHS

The initial 1 MW operating level is considered the base starting point for the SNS performance. A clearly-stated requirement for the next-generation spallation source has been [4] that if the initial performance level was 1 MW, it must be upgradable to significantly higher powers. The Kohn report [5] drew a parallel between an accelerator-based spallation source and a top-class research reactor, stating that a rough measure of scientific productivity would equate the two were the power level for the spallation source to be 5 MW. As a consequence, plans have been formulated to ensure that an upgrade path to powers at least in the 4 to 5 MW range is clearly defined for the SNS.

6.1 Second Target Station

Recalling the discussion above on specifications for repetition rates for the SNS, it was concluded that the broader neutron-scattering community could be best served were the facility to include two target stations. The first would receive pulses at the basic 60 Hz SNS operating rate, while the second could be optimized for a lower repetition rate, of say 10 Hz. Such a configuration would not only allow for a doubling of available neutron ports, but would also double the number of moderators, and would enable the optimization of each moderator design to best suit the particular neutron spectral characteristics for instruments viewing that moderator. As a result, this item is clearly the highest priority for enhancing the overall capabilities of the SNS.

Space has been reserved on the proposed site (see Figure 13) to accommodate a second target hall, and a beamline from the RTBT to the second hall. The operating mode for the facility would not change the basic 60 Hz rate for pulses being produced by the SNS accelerator system, but every sixth pulse would be directed to the second target station. Note, any combination of the 60 pulses per second could be distributed between the two target stations, but it is expected that once optimization has occurred in target, moderator and instrument design, the utilization of this flexibility is unlikely.

6.2 Two MW Operating Level

Upgrade of the accelerator complex to deliver 2 MW is quite straightforward. All of the accelerator components are designed to handle this level of beam, only a few elements need to be upgraded to reach this level of performance.

Specifically, the ion source must be improved to deliver 70 mA of beam from its original 35 mA. This is not viewed as difficult: this type of ion source has performed at this current level already [6], but for lower duty factors, so only engineering enhancements are required for management of extra heat loads from the longer SNS duty factor. Beam dynamics simulations for the front end, linac and ring sections have verified that designs of these elements are suitable for accelerating and transporting 2 MW beam levels. The extra beam-loading will require some enhancements in the linac RF system, but a very efficient scheme for addition of extra klystrons has been developed which will accomplish this upgrade in a very straightforward manner. Some very minor improvements to the target thermal-management systems will be required, but basic shielding is designed to handle a full 4 MW of beam power.

It is anticipated that this upgrade will occur as a relatively routine extension of operations. It can probably be accomplished with very little if any disruption of service, and might even happen incrementally over the first few years of operation.

6.3 Four MW Operation Level

Upgrading to 4 MW is somewhat more involved, but is still of much smaller scope than the initial construction project. As the accumulator ring is designed for a maximum of 2 MW, a second ring will be required. Again, space has been reserved on the planned site for this ring, to be built in mirror-image position from the first ring. Building a separate tunnel, while adding to the costs of the upgrade (as opposed to placing the second ring in the same tunnel as the first ring), provides for continuing operation of the first ring during construction and installation. In addition, it also adds significant flexibility to the facility for operating modes (e.g. continuing operation of one ring while handling maintenance problem in the second ring), as well as for future upgrades (placing double rings in one or both tunnels).

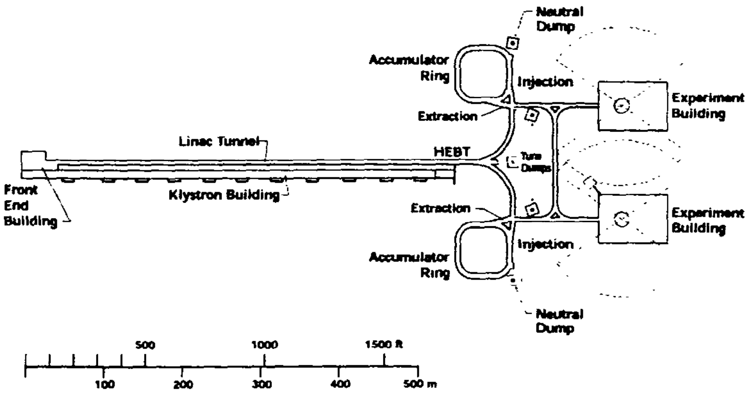

Figure 13. Footprint of SNS upgraded to 4 MW, with two target stations. Note second ring, and cross-over transport line to bring beam from both rings to either target station.

It is unlikely that a 140 mA ion source can be developed with suitable operating characteristics. As a result, the 4 MW upgrade plan calls for building a second front end system, up to and including a second DTL set. This second system would be placed parallel to the first, and a new merging section installed at the 20 MeV point to serve as a "funnel." As this is also the point where a frequency jump between 402.5 MHz and 805 MHz occurs, phasing of the two front ends will allow a doubling of the pulses accelerated in the 805 MHz section, hence a net doubling of the beam current. Note, that with only one front end operating at 402.5 MHz, only half of the available buckets in the 805 MHz section are utilized. The funnel will alternate pulses from the two front ends, filling all available buckets.

As was the case with the 2 MW upgrade, some further enhancements in the linac RF system will be needed to provide for the increased beam loading.

The switchyard will become more complex as well. The cross-over to the second target station will now become bi-directional, so that beam from either ring can be delivered to either of the two target stations. The operating mode will be as follows: the first 500 μs of the linac pulse will be used to load 2 MW of beam into one ring, the second 500 μs loads the second ring. Then the switchyard is set up for whichever target will receive the pulse, and both rings will be emptied into that target. This can be done in one of two ways: the first would have the tail of the pulse from the first ring immediately followed by the head of the second ring's pulse (with a 250 ns gap to allow for switching of the proper kicker). This would produce a pulse of almost 1.3 μs duration. A second method would be to have a parallel transport line into each target, so the heads and tails of pulses from both rings would be aligned. This would produce the shortest pulse, ≈590 ns, but would require a very sophisticated transport line design.

6.4 Further Power Increases

As mentioned earlier, having a full-energy, highly flexible linac as a base platform provides excellent opportunities for expansion of facility capabilities. The well-defined paths to 2 and 4 MW have been described, but in addition it is possible to consider inclusion of second rings in each tunnel, again doubling beam power. One could even contemplate using one of the accumulator rings as an injector for a higher-energy synchrotron that could deliver, say, 10 GeV protons to the target. While fraught with potential difficulties, the time scale for such a further enhancement is sufficiently far in the future that other projects may have paved new technology paths that would be applicable for either this scenario, or for as-yet unanticipated designs.

7. SUMMARY

The design requirements for the SNS are very well understood, and a well-conceived design has been developed which is expected to meet the required specifications with very conservative safety margins. The technologies selected are well established, leading to an overall level of technical risk which is substantially lower than normally expected for a project of this magnitude. Areas where extrapolations are called for, such as the very tight specifications on beam loss, are believed to be well understood. Extra conservatism has been incorporated into designs affected by these areas.

An overriding criterion is designing for reliability. Engineering specifications and analyses to ensure meeting this goal are an important part of the design effort. It is our goal to produce the best facility in the world for neutron scattering!

It should be noted that at the time this manuscript was prepared, construction approval for the SNS has been granted, and an allocation of $130M has been made by the United States Congress for the first year of construction. (The author is particularly grateful for the patience of the ICANS organizers, which has allowed the sharing of this wonderful news!) The project is now underway, and we fully expect that the year 2006 will see full utilization of neutron beams from the SNS.

REFERENCES

[1] The NSNS Conceptual Design Report, ORNL / NSNS / CDR, May 1997 <http://www.ornl.gov/~nsns/nsns.html>

[2] "Los Alamos Next-Generation Spallation Source," LA-UR-95-4300, December 1995, Vol 1, Figure 3-12, page 3-20.

[3] Lowell Charlton, ORNL, Private Communication, and NSNS CDR, Section 5 p. 5-97

[4] 'Next Generation Spallation Source Committee Report" BESAC, Feb 1996, Thomas Russell, Chair.

[5] "Neutron Sources for America's Future – Report of the Basic Energy Sciences Advisory Committee Panel on Neutron Sources" January 1993, Walter Kohn, Chair. DOE/ER-0576P

[6] K. Saadatmand, K-N Leung, private communication. The source provided by LBNL for the SSC actually operated at over 100 mA, but for duty factors of the order of 0.1%.