

FIRST EXPERIENCE WITH HANDOVER AND COMMISSIONING OF THE SNS CONTROL SYSTEM*

D.P. Gurd and W.R. DeVan, Oak Ridge National Laboratory,
L. Dalesio, Los Alamos National Laboratory
L.T. Hoff, Brookhaven National Laboratory, and
S.A. Lewis, Lawrence Berkeley National Laboratory

Abstract

As reported in several earlier instances of the ICALEPCS series [1], the Spallation Neutron Source (SNS) control system is being designed by a team distributed among the SNS partner laboratories, including Lawrence Berkeley National Laboratory (LBNL), Los Alamos National Laboratory (LANL), Brookhaven National Laboratory (BNL) and Oak Ridge National Laboratory (ORNL). In the past year the Front End systems, including an ion source, radio frequency quadrupole (RFQ) and medium energy beam transport (MEBT), have been delivered, re-installed and commissioned at the SNS site in Oak Ridge; as have the first components of the LANL-built drift tube linac (DTL). This paper describes the management issues involved as well as our experience and lessons-learned in handing over equipment designed and built at one laboratory for installation and commissioning at another. There is good news and bad news.

INTRODUCTION

The Spallation Neutron Source (SNS) facility currently under construction in Oak Ridge, Tennessee, is being built as a collaboration of six US Department of Energy National Laboratories. When it comes on line in 2006, the SNS will be the largest source of spallated neutrons in the world, exceeding the present world leader, ISIS in the UK, by a factor of approximately 80. The Ion Source, RFQ and Medium Energy Beam Transport (MEBT) systems were designed and constructed by Lawrence Berkeley National Laboratory (LBNL) and have already been delivered to the Oak Ridge site and recommissioned. A Drift Tube Linac (DTL) and Cavity Coupled Linac (CCL) will be provided by Los Alamos National Laboratory (LANL), which will also provide components of the warm sections of a Superconducting Linac (SCL). The cryomodels for this SCL and the associated Cryoplant for the manufacture of liquid helium are coming from the Thomas Jefferson National Accelerator Facility (TJNAF). A compressor ring is to be provided by Brookhaven National Laboratory. Conventional Facilities and a 1.4 MW liquid mercury target for neutron production are the responsibilities of the home institution, Oak Ridge National Laboratory (ORNL). Finally, the design and assembly of a suite of neutron scattering instruments has been the responsibility of Argonne

National Laboratory (ANL).

As of July, 2003, the project is 68% complete. Civil construction is well advanced and the accelerator tunnels and buildings are largely complete and rapidly filling with equipment. Only the Central Laboratory and Office (CLO) building remains to complete. The Front End Systems from LBNL have been accepted and recommissioned in Oak Ridge, and the team at LBNL has been disbanded. The first of the DTL modules has also received beam. Most of the design work on the warm linac is complete, and the SNS Division at LANL will be disbanded in April, 2004, although a small amount of remaining work will be continued under a Memorandum of Agreement (MOA) between LANL and ORNL. This work explicitly includes continued participation by the LANL controls team. Similar MOAs have been or are being negotiated with the other partner laboratories, and the entire instrument team has moved from Argonne to Oak Ridge. By next May, the active partnership will consist only of BNL, TJNAF and ORNL.

MANAGEMENT APPROACH

Development of the SNS control system has been organized and managed in a way that reflects on a smaller scale the organization and management of the project as a whole. Thus the distributed controls for each partner laboratory-provided subsystem is developed at that laboratory, and delivered with the subsystem it controls. It is the role of the controls team at ORNL to work with its distributed partners to accept, integrate and then commission these subsystems as they are "handed-over" to the Oak Ridge team.

At previous conferences in this series, we have outlined and discussed the organization and management issues inherent in such a collaborative and distributed control system development. Before reporting on the handover successes and difficulties to date, we summarize some of the issues and the approaches we have taken.

Management and Organization

The most important idea in the organization of this collaborative control system development was to treat the control system as a "Level 2" activity in the project management system. That is, the control system was treated at the same management and reporting level as each of the partner laboratories or major subsystems. (Linac, Ring, Target, etc.). (See figure 1 below.) This was by no means obvious, and was not accepted without considerable discussion. There have been two,

* Work supported by the US Department of Energy under contract DE-AC05-00OR22725

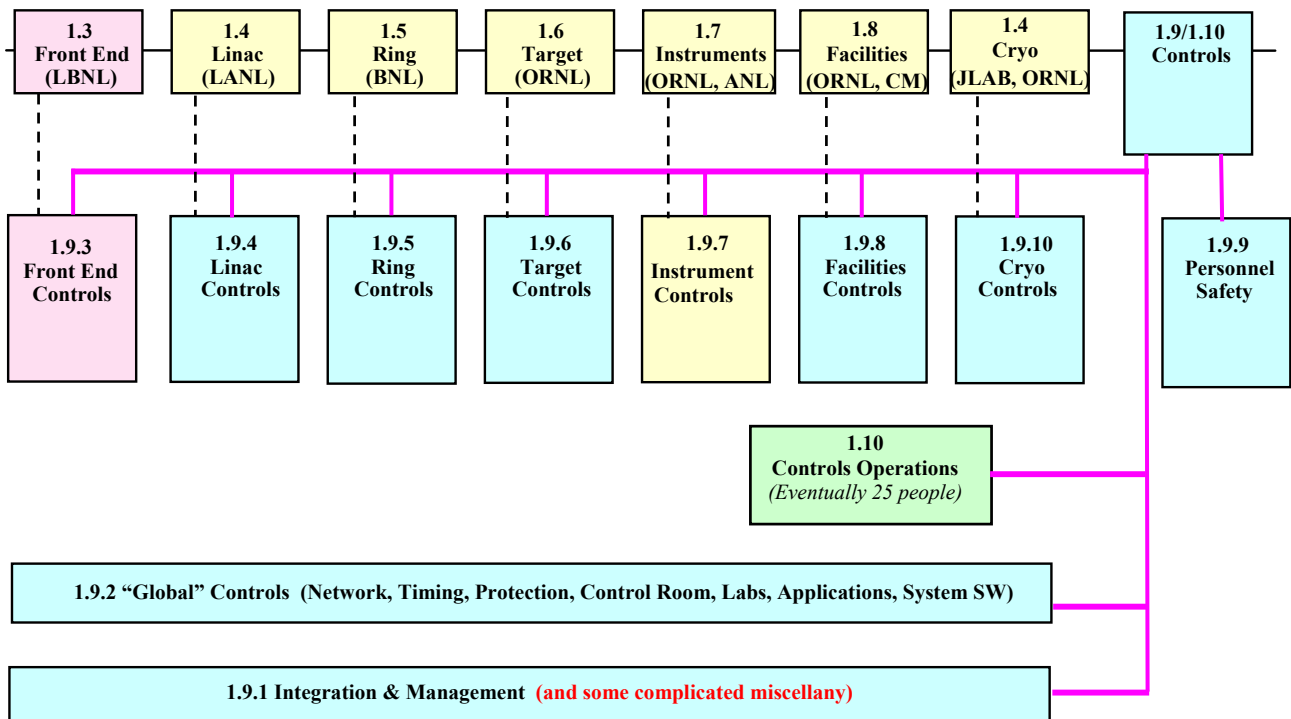


Figure 1: Organization of the SNS Controls Group

“downsides” to this approach for controls management. First has been the complexity of financial reporting which requires the integration of diverse reporting tools and methods from several different partner laboratories; and second has been the high visibility of the control system. (We report to the Department of Energy at the same level as LANL or BNL for example.) Nonetheless this approach has given the controls team a degree of autonomy and authority it would not otherwise have had, and can be recommended to collaborative projects in the future. Successful standardization and integration is hard enough as it is, and would be far harder with a different management model.

Standardization

Another important aspect of the management approach was to standardize as extensively as possible. This is self evident, but more difficult than might be supposed. The SNS has had moderate success in this area. At previous conferences we have discussed standards in the area of software, hardware and process. The most important single standard adopted was the use of EPICS as the underlying system. This was not completely evident, as some of the partner laboratories were not previously familiar with this toolkit, and on their own might have chosen differently. We also successfully standardized on the display manager (EDM), Linux and, for higher-level applications, XAL. Oracle was accepted as the project database, but the tools were not provided to facilitate its use in the generation of EPICS databases, and the use of Oracle in the control system cannot yet be counted as a success. More successful, but not without caveats, was the

agreed imposition of the use of CVS as a mechanism for controlling software developed by the partners. By now this is the process universally used to hand over software to ORNL, and software to be loaded on the SNS control system is built from modules in CVS. Still, there remain differences in detail over which we occasionally trip.

Hardware standardization has been quite effective. The SNS architecture includes many PLCs, and with only a few exceptions these are ControlLogix PLCs from Allen Bradley. Only one vendor-supplied PLC in the Cryoplant will not be ControlLogix (although two partner-provided still remain to be converted.) Input-Output Controllers (IOCs), both crates and processors, are all purchased on standardized ordering agreements, as are instrument racks. Nonetheless, there has been an unwanted proliferation of fieldbuses, which now includes Allen Bradley “blue hose,” Group 3, FlexIO and Beckhoff. At what point does one decide to compromise the best solution for the problem at hand in favor of more stringent standardization?

As discussed at the last ICALEPCS, the most problematic attempted standard has been what is arguably the most important – device and signal naming. This failure occurred in spite of honest attempts by all parties to follow a standard that was published early in the project. Addressing the problems related to naming has been frustrating, time-consuming and expensive.

HANDOVER

The third general management topic which we have presented at earlier conferences is that of the “handover”

philosophy – that is, the process by which systems developed at the partner laboratories are handed over to the Oak Ridge Team for installation, testing and eventually commissioning with beam. The model has been christened “Lead, Mentor, Consult.” The idea is that engineers and developers from the partner laboratories take a lead role in the installation of the first systems of a kind, a secondary, mentoring role for the next few, and then retire to their home laboratories and act as consultants only as required as systems continue to be installed. As we shall see, this model, while reasonable, has been interpreted and applied differently for different subsystems. As a model it may be ideal, but in practice there have been and will be many variations. A description of handover experience with some of the subsystems follows.

Conventional Facilities Handover

Conventional Facilities (CF) – the power, Heating, Ventilation, Air Conditioning and other building controls – represent a very special case. It was first agreed to integrate the Conventional Facilities controls with the Accelerator controls, and to use EPICS to do it. This was a difficult case to win, and was eventually agreed only with serious reservations expressed by the contractor responsible for these systems. As a compromise, the IO level was implemented using PLCs, with which the contractor and our own CF personnel were comfortable. Next, it was agreed that the control system, including the EPICS layer – both databases and screens – would be implemented by a commercial contractor, Sverdrup Control Systems in Tullahoma, Tennessee. An agreed scope of work was defined, a contract was signed and an EPICS training program was carried out for Sverdrup employees. A very limited and conservative set of EPICS tools was to be used. Like our National Laboratory partners, CVS would be the mechanism for handing over software.

Notwithstanding rampant skepticism, this approach for Conventional Controls has been a resounding success. The first operational IOC at the SNS site was for building controls. The Sverdrup team came to the SNS site to install and test the system they had developed and there were very few problems. They have been readily available and responsive on the phone when there have been issues (very few) and on one occasion did come to the SNS to help resolve a problem.

The success of the CF handover was entirely the result of good contract management. One control system engineer took this responsibility. The level of documentation required when dealing with an industrial partner is greater than for in-house development, and frequent visits to the supplier’s site were undertaken. (It was a drive of a few hours.) There are no more skeptics. The same model has been adopted for parts of the target control system. A spin-off of this approach has been the availability nearby of a pool of trained EPICS engineers who could be contracted to do EPICS work in the event of a shortage of personnel.

Front End Handover

The Front End Systems (Ion Source, RFQ and MEBT) were developed at LBNL in Berkeley. The controls team was very small – essentially only two full-time people. The front end was also unique in that it was completely assembled, operated and characterized at Berkeley before being shipped to Oak Ridge for recommissioning. This will be the only system to have been operated as an accelerator before shipment to ORNL.

The original beam operation of the Front End Systems at Berkeley afforded SNS and the controls team a unique opportunity to participate in its commissioning and learn its idiosyncrasies before it was shipped to ORNL. It didn’t hurt that with the shipment was included one of the two LBNL control system engineers who developed the system originally and who will stay in Oak Ridge for two years. Not surprisingly, his understanding of both the hardware and software architecture of the system greatly facilitated the handover.

A particular issue with the Front End Control System was the fact that its early development in the SNS schedule resulted in many design decisions being made before SNS standards had been established. A different PLC was used; a different display manager was used; a different approach to machine protection was used and the device and signal names were eccentric. All of these have been or will be converted to SNS standards, however this task seldom reaches top priority, and is lagging behind schedule.

Overall, the handover and recommissioning of the Front End Systems from Berkeley has been a success, but because of the special circumstances it cannot be considered a realistic model for later subsystems.

DTL Handover

The handoff of the control system for the warm linac from LANL, and the cold linac from JLab and LANL together, will follow more closely the “Lead, Mentor, Consult” model, which was in fact originally proposed by LANL. Although we have operated only the very first warm linac component (DTL1) thus far, we have installed most of the DTL subsystems and some of the CCL subsystems, and therefore have experience with both the “lead” and the “mentor” phases of this model. It is the handover of the DTL subsystems which has been the most instructive. The bottom line has been success: accelerated beam – relying on the control system for both control and data acquisition – on the first attempt. There is value nonetheless (we hope) in noting some of the difficulties incurred.

The idea of the “Lead” phase is that installation and testing of the first instances of each subsystem should be done by the partner lab engineers. The DTL, however, had many unexpected technical problems and schedule slippages and it was difficult or impossible to correctly schedule visits by the LANL controls team. On occasion they would come according to the best laid plans, and find the subsystem was not ready for one reason or another. A frustrating week wasted. On other occasions a window of

opportunity would be presented on short notice, and the particular installation or test would have to be done by the ORNL team, without the intended involvement of LANL. The dynamic schedule also made it difficult for the design engineers to know when a particular subsystem would be required – there was a temptation to wait until the last minute, possibly gilding the lily unnecessarily. All of this led to some tension between the ORNL team “on the scene” and the LANL team who, from their distance, could not appreciate the reasons for change.

When engineers did arrive bearing software, we quickly learned the importance of following the CVS protocol. Software remotely checked in to CVS from LANL and then built into IOCs at ORNL generally worked well. Software that had worked successfully at home and was carried in laptops to ORNL often rebelled against a foreign environment. There was an analogous circumstance with some hardware delivery. In some cases assumptions about wiring were made that proved different in the field. This resulted in the need for some field rewiring. LANL technicians visited ORNL to help with this work, but it was an additional cost and inefficiency. Making use of the “Lead” concept, we now make sure LANL technicians help with the installation of the first system, so that they can effect any necessary rewiring before the delivery of subsequent systems. In one case there was a difference of philosophy between LANL and ORNL for field wiring. A command decision was made to follow the ORNL preference. LANL drawings were not changed however, and this led to confusion and more rework. Once a decision is made, all partners must act accordingly.

The issue of software handoff was particularly difficult. After initial installation and testing by LANL engineers, it was frequently the case that changes were required in response to operational needs. The LANL engineer – not present in the control room – might not appreciate the operational imperative. There was understandable frustration that a carefully designed program was being modified in an ad hoc way “on the fly.” Or the delivered IOC program might be combined with something locally-developed (each IOC includes Machine Protection and Timing databases, for example) and the LANL engineer would have difficulty feeling responsibility for “his” working IOC that had been modified. Or the locally made change might conflict with upgrades being made at home. Who’s in charge here? These situations were worked out on a case-by-case basis. It is hard to imagine a general rule that would be effective.

Of necessity, the first screens developed are “engineering” screens, intended to help develop and test new hardware. These screens are essential, but they are not well suited for operations. Moreover, the interpretation of a “high-level overview screen” is different for a remote subsystem designer than it is for an accelerator operator. Indeed, it is scarcely possible for the partner laboratory to deliver overview and “comfort” screens, because these by definition cross subsystem

boundaries and they require the input of the operations staff. This work has to be done by the on site programming team.

This is a long litany of issues and lessons-learned. Lest the resulting impression be negative, be it known that the bottom line thus far has been decidedly positive. DTL systems have been available on time and used for operations from the word go. The lessons learned have been applied effectively, and the CCL installation has gone far more smoothly. The LANL team has become increasingly engaged in operations, and sympathetic to the operational imperatives – all be they distant. The principle reason for success has been the determined desire of the combined controls team to work together to address issues as they arose, and their willingness to find workable compromises that would assure success. An extremely important tool in making this work has been the ability to monitor operations remotely, and effect necessary changes from LANL – always of course with permission from and/or in contact with local personnel.

We have as yet no experience with the handoff of the Cold Linac or Ring control systems. Lessons learned from our experience with the warm linac are being applied however – the most important of which has been the application of far more formal system engineering and documentation. Several members of the BNL controls team have also spent time in the control room during linac commissioning, and are being conditioned to be sympathetic to unexpected new or changed requirements needed “right now.”

A final and particularly difficult issue related to the partnership is personnel management as the project draws to a close. In the case of LANL, the specially-created “SNS Division” is scheduled to disband in April 2004. The role of the Controls Team will continue, but the special overhead rate negotiated to support SNS goes away, and the cost goes up. As there is no additional budget to support this cost, we have had to reduce the planned level of support from LANL for the remainder of the project. This in turn creates a problem for LANL management – how to support staff that has been prematurely removed from the project.

One rationale for a partnership in the first place (other than political expediency) is that the partner labs were to be both a source and an eventual sink for trained accelerator talent. The source part worked well; the sink less so. As the responsibilities of the partner laboratories are completed, there is not always an assurance of work for people coming off the project. Some will leave early when a more permanent position presents itself. Others may be distracted by a concern for their future.

The SNS has put in place “Memoranda of Agreement” (MOAs) with each of the partner laboratories which come into effect when the official scope of work of the partner is completed, and which define the level of support which ORNL undertakes to support, and which the partner undertakes to provide. It is not much.

OPERATIONAL EXPERIENCE

Notwithstanding the issues cited above, the bottom line is that the contractor and partner lab-delivered control systems have in each case been installed, tested and handed over in time to support initial operations. The Front End was recommissioned, two DTL modules were rapidly conditioned and beam was accelerated in the first DTL module – all fully under the control of the control system. Accelerator old-timers who don't realize that this should be expected in a modern control system have expressed surprise and offered praise. We on the other hand see the deficiencies, and they are briefly mentioned below.

Not surprisingly, the partner labs provided primarily engineering screens that had been designed originally to support the development of new equipment or subsystems. There are too many screens (over 1100 screens already!), navigation between the screens is complicated and unintuitive and there are few if any operator-friendly high-level “comfort screens.” Such screens would be difficult to design without an interaction with operators, and have become the responsibility of the ORNL team. This should have been anticipated from the outset.

The archiver has been an absolutely invaluable tool during early commissioning. It was not available at the start of the first run – a mistake – but has been depended upon during all subsequent running periods. It has been difficult with the design engineers at a distance to get the right Process Variables into the archive configuration files at the right rates, and frequently we discover that something has not been archived that would have helped a later analysis. Initial archive files should be the responsibility of the partner lab design teams.

We operate still without benefit of an Alarm Handler, although an initial configuration of the EPICS Alarm Handler tool (ALH) has been made available to

operations. During this early period of rapid change many parameters are operating in off-normal conditions, the alarm screens are covered with red and yellow and are simply not useful. It has not yet been accepted as useful, even though it clearly provides an easier navigation mechanism to problem sources than speculating one's way through the maze of screens.

We have had difficulty getting the EPICS Save and Restore and Bumpless tools to work in a foolproof manner, and this was often a problem when IOCs required rebooting. It is somewhat of an embarrassment that the Save and Restore tool of preference for operations is a new tool created in Java by the SNS Accelerator Physics Team

The most serious operational problem experienced at SNS has been what has come to be called “IOC Disease.” Under certain not understood conditions, IOCs stop intercommunicating, buffers fill up and data is lost. Reboots help only temporarily. Although we can find stable configurations, these alarming symptoms are never far from the surface. Except for two virulent outbreaks, the EPICS communication protocol has been robust enough to handle the level of errors experienced; however it does not augur well for a larger, more stressed system.

SUMMARY

Notwithstanding a number of management hurdles, the handoff of partner or contractor-built control subsystems has been successful. We are demonstrating the feasibility of building and integrating a large control system using a widely distributed partnership, and we are learning as we go. We should really know how to do this the next time...

REFERENCES

- [1] D.P.Gurd. “Management of a Large Distributed Control System Project,” ICALEPCS 2001, San Jose CA 27-30 November 2001.