Experience with collaborative development for the Spallation Neutron Source from a partner lab perspective

L.T. Hoff, BNL, Upton, NY, USA*

Abstract

Collaborative development and operation of large physics experiments is fairly common. Less common is the collaborative development or operation of accelerators. A current example of the latter is the Spallation Neutron Source (SNS). The SNS project was conceived as a collaborative effort between six DOE facilities. In the SNS case, the control system was also developed collaboratively. The SNS project has now moved beyond the collaborative development phase and into the phase where Oak Ridge National Lab (ORNL) is integrating contributions from collaborating "partner labs" and is beginning accelerator operations. In this paper, the author reflects on the benefits and drawbacks of the collaborative development of an accelerator control system as implemented for the SNS project from the perspective of a partner lab.

Introduction

Project collaborations are formed for a number of reasons. Often collaborators hope to gain preferential access to the completed device or facility. Sometimes collaborators have specialized competency to offer or available capacity to perform work, or may wish to develop specialized competency. Collaborators may even have complete designs available, requiring little or no modification. Collaborations may be formed for non-technical reasons, such as to gain access to additional funding sources, or sources of political clout.

An important aspect of any non-trivial project is the effective partitioning of work among the various participants, including collaborators, subcontractors, and vendors. In his famous essay "the mythical man month" [1], Frederick Brook observes that "In tasks that can be partitioned but which require communications among the subtasks, the effort of the communication must be added to the work done.", and "The added effort of communicating may fully counteract the division of the original task." These issues are magnified in large, geographically distributed projects. A collaborative arrangement suggests more of a joint intellectual effort than a vendor or subcontractor relationship. This implies a continual refinement of project design and work scope, further magnifying such issues. To some degree technology, including telephone and video conferencing, e-mail and other network file transfer protocols, has evolved to mitigate long-distance communication barriers [2], but imperfections in these solutions remain.

Collaborative efforts face additional unique challenges. One such challenge is the difficulty capturing and retaining institutional knowledge gained during the development phase of the project. This can be especially true for collaborations which disband shortly after such knowledge is gained, and is less true for collaborations which persist through commissioning and into facility operation. Another challenge is addressing the potential for widely varying development practices and styles used by different collaborators. This situation can be addressed by agreeing to one common set of development practices, or by finding techniques for allowing different styles to harmoniously coexist.

At the time the author was asked to assume the role of BNL/SNS team leader, the work breakdown had already been organized, budgets had been established, and the majority of the staff were in place with tasks assigned. This sequence of events did not provide the author with particular insight into the process leading to these decisions. However, it did provide the author with a uniquely

*Work performed under the auspices of the US Department of Energy

unbiased perspective from which to observe and analyze the benefits and drawbacks of the SNS collaborative organization.

SNS project

Detailed information regarding the SNS project can be read from the project's web site: http://www.sns.gov. The following paragraph is reprinted from the page: http://www.sns.gov/aboutsns/source.htm:

The SNS project is a partnership involving six DOE national laboratories (Argonne, Brookhaven, Jefferson, Lawrence Berkeley, Los Alamos, and Oak Ridge) in the design and construction of what will be the most powerful spallation source in the world for neutron-scattering R&D. The baseline design calls for an accelerator system consisting of an ion source, full-energy linear accelerator (linac), and an accumulator ring that combine to produce short, powerful pulses of protons. These proton pulses impinge onto a mercury target to produce neutrons through the spallation nuclear reaction process. The SNS will deliver 1.4 million watts (1.4 MW) of beam power onto the target, and it has been designed with the flexibility to provide additional scientific output in the future.

SNS work breakdown structure

The SNS work breakdown structure (WBS) was organized to reflect the loosely coupled nature of the work performed by each partner laboratory. Each partner laboratory worked within separate "level 2" work breakdown structures. Oak Ridge National Laboratory (ORNL) retained overall project responsibility, performed site civil construction, approved partner laboratory work packages, and authorize budget authority.

The control system, however, was organized with a more global philosophy. [3] The entire control system was treated as a WBS "level 2" activity. "Satellite" controls groups were established at partner laboratories operating at "level 3" of the work breakdown structure. That is, the partner laboratory controls effort was *not* organized as a sub-activity of the partner laboratory effort. It was organized as a sub-activity of, and therefore somewhat more beholden to, the global controls effort. By establishing these "satellite" controls groups, control system support could be co-developed with accelerator equipment at the partner laboratories, prior to delivery to ORNL.

Global controls activities, performed at ORNL, included selecting hardware and software infrastructure. This included selection of EPICS as the control system framework, VME64X crates to house distributed embedded systems, MVME2100 processors and Allen Bradley ControlLogix PLCs, LINUX operator interfaces (OPIs), and concurrent versioning system (CVS) for configuration management. ORNL was also responsible for the global timing, machine protection, and personnel protection systems, as well as developing and documenting global standards for device naming, human-machine interface (HMI), etc., and for coordinating installation efforts, and integration of partner laboratory contributions.

Control system characteristics

The primary purpose of an accelerator control system is to provide accelerator operators, typically located in a centralized control room, with a view of the accelerator equipment, so that they can understand and adjust accelerator performance. This implies that the control system acts as "glue", bridging the gap between accelerator devices and operator screens.

The "principle of least astonishment" requires not only that the control system provides a global view of the accelerator, but that such a view is uniform across the entire facility. Uniformity should include nomenclature, system performance, and control screen organization and other HMI issues. E.g. adjusting the gain of a beam loss monitor should appear substantially similar to the operator whether the beam loss monitor is located in the SNS accumulator ring, developed by BNL, or near the ion source, developed by LBNL, even if the loss monitor systems use different underlying technology. In this case, the control system must “wallpaper over” the differences in technology to achieve a uniform view.

During development, installation, and even early commissioning of accelerator equipment, control screens need to provide an engineering view of a particular piece of accelerator equipment. Such screens are vital for allowing technical specialists to verify the fundamental performance of the accelerator equipment, and typically provide a comprehensive array of performance diagnostics. However, accelerator operators more commonly require task-oriented control screens, which span multiple pieces of accelerator equipment, but assume proper operation and performance of each piece of equipment. E.g. a control screen for adjusting beam orbit may include measurements from beam position monitors (BPMs) and control magnet power supply levels. Such a screen might well omit details of power supply interlocks or BPM timing configurations. It is rare for one control screen to suit both purposes without unreasonable compromise. Typically, “engineering screens” are developed in conjunction with the technical specialist designing the accelerator equipment, and “operator screens” are developed in conjunction with accelerator operators as operational procedures and tasks are defined.

The BNL/SNS controls perspective

In general, the decision to create a “satellite” controls group at BNL, and to co-develop control system support with accelerator equipment, proved to be sound. It is difficult to imagine managing the inevitable tactical decisions necessary during such co-development without close and continual contact between controls staff and technical groups developing SNS accelerator equipment. Close and continual contact requires a degree of trust and understanding that can only be developed over time. When possible, the BNL/SNS controls group was staffed with BNL employees with prior accelerator controls experience. Additional controls staff were hired as ORNL employees sited at BNL. The use of veteran BNL employees leveraged established working relationships with other BNL technical groups developing SNS accelerator equipment.

Early integration of control system and accelerator equipment can yield enormous benefits. It can uncover misunderstandings regarding signal levels, connector and cable types, unexpected sources of noise, or any other issues of compatibility or quality. A process of gradual refinement [4] is possible during which such issues are uncovered, then corrected either by adjusting the accelerator equipment or the control system support. During this process, idiosyncrasies of the equipment are invariably uncovered. This “institutional knowledge”, gained during the development phase, can be valuable during installation, commissioning, and early operation of the equipment. This type of knowledge is also very difficult to capture and convey to those who would most benefit from it. One approach, used to some degree by the SNS global controls group, was to visit partner laboratories and participate in development or witness demonstrations of accelerator equipment. This seemed to be most successful with LBNL, who operated and characterized a complete ion source at LBNL prior to delivery to ORNL. The most successful approach to this issue from the BNL perspective was to transfer some of the BNL/SNS staff to ORNL as BNL activities wound down. This was facilitated somewhat because several of the BNL/SNS staff were already ORNL employees. This approach was fairly widely adopted. Controls staff also transferred to ORNL, either on a temporary or permanent basis, from both LBNL and LANL.

The decision to organize the “satellite” controls group as a sub-activity of global controls instead of a sub-activity of the partner laboratory effort proved equally sound. The ORNL global controls group could maintain a broad perspective when making strategic decisions establishing policies and priorities. For instance, the two largest accelerators at BNL, The Relativistic Heavy Ion Collider (RHIC), and the National Synchrotron Light Source (NSLS), both use proprietary control systems for accelerator operations. The experienced BNL employees in the BNL/SNS controls group were comfortable working within the context of these control systems. If BNL delivered SNS control system support using one of these control systems, ORNL would face the prospect of porting the support to a common control system context, or implementing appropriate “bridging” techniques to create a uniform view across the entire facility. Instead, the global controls group selected a common control system infrastructure (in this case, EPICS), and provide the partner laboratories with necessary training and documentation to be able to work effectively in an unfamiliar environment.

The mythical man-month

This division of roles, with tactical decisions made by partner laboratories, and overall strategic guidance from ORNL, is consistent with the concept of maximizing the use of high-bandwidth and low-latency channels for most required communication. This philosophy was not, and perhaps could not be, universally applied. In cases where it was not applied, there was generally noticeable strain. Normally, ORNL conducted a weekly telephone conference with the “satellite” controls groups. This conference was used to report status, discuss issues of global impact, and generally keep all groups apprised of the activities of other groups. Some amount of e-mail, private phone conversations, and even video-conferences were used as needed for specific purposes. When the frequency of inter-laboratory communication increased, or when people as if they could not proceed with their work until they received information from a remote source, there was a strong likelihood that tactical and strategic roles were becoming blurred, perhaps unnecessarily and unwittingly.

One such case involved the inherent conflict between the “engineering screens”, which are vital for technical specialists developing and commissioning accelerator equipment, and “operator screens”, alarm displays, and data archiving configurations, which are the mainstay of the accelerator operator. Partner laboratories generally understood their role to include developing “engineering screens”, working closely with technical specialists. Conversely, partner laboratories generally expected that ORNL would develop appropriate “operator screens”, working closely with the accelerator operators at ORNL. For a variety of reasons, including both controls group staffing levels and initially only an embryonic accelerator operations group, ORNL did not develop needed “operator screens” in time for early commissioning efforts. Several less-than-ideal situations emerged. Operators were forced to make use of “engineering screens”, which were dense, difficult to navigate, and not task-oriented, leading to considerable clutter. During early SNS commissioning efforts, BNL/SNS controls staff attempted to work directly with ORNL operations staff to develop appropriate alarm displays. This was somewhat successful while on-site, when effective communication was possible, but failed completely when remote communication was necessary. Finally, due to the extensive operational experience of some of the BNL employees, the BNL/SNS group made some attempts to develop “operator screens” according to ORNL HMI standards. From a BNL perspective, this was most successful when the standards were stable over time. This was least successful when such standards were continually refined as operational experience was gained, especially if partner laboratories were expected to keep current with all such refinements and to upgrade previously developed “operator screens”.

Another example where the distinction between strategic and tactical roles became blurred, resulting in “coupling” that was somewhat too tight between the partner laboratories and the global controls group, involved the specific release level of EPICS used when developing, delivering, and continuing to maintain control system support.

The global controls group had two, somewhat conflicting, simultaneous roles. ORNL was responsible for receiving and integrating contributions from partner laboratories, and also for supporting accelerator operations during early commissioning efforts. ORNL necessarily needed to make tactical decisions regarding the most appropriate EPICS release level to support accelerator operations. The impression at BNL is that partner laboratories were expected to maintain fairly strict lock-step agreement with EPICS release level, even at the “minor” level (e.g. 3.13.6 vs. 3.13.7). The use of a centralized source-code repository (CVS) combined with a “best practices” philosophy of regular source code check-ins could make it quite visible when partner laboratories deviated. A more ideal approach might have been to require reasonably tight coupling of EPICS release levels only at the “major” level, but permit each partner laboratory to make tactical decisions at the “minor” level, driven by local needs (including accelerator operations). This might require the global controls group to develop procedures for reconciling differences at the “minor” level.

A final example of blurring between tactical and strategic roles involved system installation. The LANL/SNS controls group was deeply involved with the installation of controls equipment in the SNS LINAC. Installation schedules and manpower allocations required tactical decisions on the part of ORNL, based on local issues including budget, delivery delays, unexpected equipment failure, and even weather. The LANL/SNS controls group attempted to adapt their travel and installation schedules to match the necessary ORNL tactical decisions with limited success. In light of that experience, the BNL/SNS controls group intentionally became much less involved in installation of controls equipment, shifting that responsibility to ORNL. This required much more up-front documentation and knowledge transfer between BNL and ORNL, but loosened the coupling between ORNL installation schedule and BNL development and delivery schedules.

The wisdom of crowds

A side effect of staffing “satellite” controls groups with veteran accelerator controls professionals is that it afforded what is commonly known as “the wisdom of crowds”. [5] The background and experience of the combined staff in all the “satellite” controls groups was more diverse than might be expected in one centralized group. Such broad experience can be used to avoid unfruitful pursuits, more quickly recognize and resolve unwanted or unexpected behavior, and recognize where existing solutions might be applicable.

For example, although RHIC and the NSLS do not use EPICS for accelerator operations, they both use VxWorks and VME for their distributed, embedded systems. This experience became valuable to SNS when a bug in the VxWorks system call “memPartInfoShow()” resulted in IOC crashes. This bug had been previously uncovered by the RHIC controls group, who had developed a fix. The BNL/SNS controls group was able to recognize the bug, and offer the bug fix, and vastly improve SNS IOC reliability.

BNL was also able to provide a nearly turn-key timing system for SNS, including event generation, distribution, and decoding. If a timing system had been developed specifically for SNS it would likely have had a vastly different architecture than the timing system used for RHIC. However, the RHIC timing was recognized as good enough to meet SNS requirements, and had been in operational use sufficiently long to be free of uncertainty about reliability. Obviating the need for a timing system development effort relieved both budget and schedule pressures.

Intangibles

Asking if software is properly tested, robust, and documented for both users and maintainers has been compared to asking “How long is a piece of string?” That is, these characteristics are difficult to quantify in an objective, measurable manner, consistent with project management principles. The level of robustness required for demonstrating control system operation while commissioning a single piece of accelerator equipment in a lab environment is vastly different than the level of robustness required to simultaneously operate dozens, or even hundreds of such devices 24 hours per day, 7 days per week. Robustness is typically much easier to implement initially, than to re-fit later. Therefore the effort or cost of ensuring appropriate levels of robustness is typically borne by the partner laboratory developing the control system support. The benefit is chiefly enjoyed by the global controls group which has responsibility for accelerator operations. Since characteristics such as robustness are difficult to measure, there could be a temptation to cut corners. This makes it imperative that the partner laboratories have an understanding of what characteristics are important for accelerator operation and where engineering margins must be maintained, and a philosophy of maintaining high levels of quality in all aspects of design and construction. One way to achieve that is to staff "satellite" controls groups with veterans in the field of accelerator controls, especially those in more senior and leadership roles, as was done in the case of the BNL/SNS controls group.

When someone is asked to take responsibility for systems created by others, the "not invented here (NIH) syndrome" [6] can manifest itself. Smart, creative people can be tempted to substantially modify, or even completely redesign a system because they believe they can produce a superior product, even if the delivered product meets all requirements. It can be difficult to differentiate between a case of NIH syndrome and a case where redesign is truly necessary. If unchecked, this practice can lead to large duplications of effort, increased costs to the collaboration, and even ill will. The BNL experience with system handoff was generally positive. Most systems developed at BNL were delivered to ORNL, installed by the ORNL controls group, and were deployed for operational use with only minor refinements. The RHIC timing system, despite having a distinctly alien pedigree, was accepted and deployed without modification, and has performed quite adequately. However, the software for one beam diagnostic system was substantially redesigned after delivery to ORNL.

Conclusion

Collaborations on large experimental facilities may become more common in the future. Such collaborations may include partitioning control system effort among participants. The benefits of such collaborations can be substantial. However, avoiding common pitfalls associated with collaborations requires careful attention during the planning and partitioning of work, and constant vigilance as work is executed.

Special attention should be paid to the early adoption of global standards, effective use of communication channels, mechanisms for capturing and retaining institutional knowledge, and techniques for avoiding the "not invented here" syndrome.

References

[1] Frederick P, Brooks Jr., *The mythical man-month* (1975 Addison-Wesley Publishing Company Inc.).
[2] Thomas L. Friedman, *The World Is Flat: A Brief History of the Twenty-first Century* (Farrar, Straus and Giroux 2005)
[3] D. Gurd et al, *First Experience with Handover and Commissioning of the SNS Control System* (ICALEPCS 2003, Gyeongju, Korea)
[4] Paul Graham, *Hackers and Painters: Big Ideas from the Computer Age* (O'Reilly May 2004)
[5] James Surowiecki, *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations* (Doubleday May 25, 2004)
[6] http://www.wikipedia.org/wiki/Not_Invented_Here