

LA-UR-15-27874

Approved for public release; distribution is unlimited.

Title: Replacing the Engine In Your Car While You Are Still Driving It

Author(s): Bjorklund, Eric A.

Intended for: ICALPECS 2015, 2015-10-19/2015-10-23 (Melbourne, Australia)

Issued: 2015-11-04 (rev.2)

Disclaimer:

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

REPLACING THE ENGINE IN YOUR CAR WHILE YOU ARE STILL DRIVING IT*

E.Bjorklund[#], LANL, Los Alamos, NM 87545, USA

Abstract

Replacing your accelerator's timing system with a completely different architecture is not something that happens very often. Perhaps even rarer is the requirement that the replacement not interfere with the accelerator's normal operational cycle.

In 2011, The Los Alamos Neutron Science Center (LANSCE) began the purchasing and installation phase of a nine-year rolling upgrade project which will eventually result in the complete replacement of the low-level RF system, the timing system, the industrial I/O system, the beam-synchronized data acquisition system, the fast-protect reporting system, and much of the diagnostic equipment [1]. These projects are mostly independent of each other, with their own installation schedules, priorities, and time-lines. All of them, however, must interface with the timing system.

INTRODUCTION

LANSCE had its beginning in 1972 as an 800 MeV "meson factory" [2]. Since then it has expanded its missions to include such diverse projects as pion treatment for inoperable cancers, spallation neutrons, ultra-cold neutrons, medical isotope production, and proton radiography.

In preparation for the new MaRIE project [3], we are undertaking an ambitious overhaul of a large part of the facility while still trying to maintain a viable user program.

This paper will focus mostly on the timing system replacement project, its conversion from a home-built, centralized, discrete signal distribution system, to a commercial event-driven system from Micro Research Finland [4]. We will explore some of the challenges faced by having to interface with both the old and new equipment until the upgrade is completed.

PROJECT STRATEGY AND SCHEDULE

The installation/operations schedule can be compared to driving through mountainous terrain on a road with many peaks and valleys. When you start down a valley, you shut down your engine, replace as much of it as you can, then try to get it running again before you have to start up the next peak. At the bottom of each valley there is a relatively flat stretch of road representing the "startup period" – during which you mostly coast while you discover how your changes affected the machine's operation (for good or for ill).

The current installation and operation schedule is

shown below in Figure 1. The green blocks represent the operating periods, the red blocks represent the installation and maintenance periods, and the yellow blocks represent the startup periods. The durations of the operation, maintenance, and startup periods vary as the project progresses. The first three years of the schedule call for longer operational periods (seven to nine months), shorter upgrade periods (four months), and shorter startup periods (one month). The middle three years – during which the most complex upgrades take place – have longer maintenance periods (four to five months), longer startup periods (three months), and shorter operational periods (three to four months). During the last three years, things theoretically get easier and we go back to longer operations, shorter maintenance, and shorter startup periods.



Figure 1: Installation And Operation Schedule

Budget Schedule

Our controls group adopted a budgeting strategy of purchasing all the equipment for each project at once. With some exceptions, for each fiscal year a different project had its own "year of profligate spending". The first year it was the network installation. The second year it was the industrial I/O system. Next was timing, etc. There were a number of reasons for adopting this strategy.

One technical reason was uniformity. All of the equipment for a sub-system would be from the same vendor, with the same firmware level, and therefore have a uniform interface to the controls software.

A scheduling reason was flexibility. Once you have had "your year" your system doesn't have to worry about having enough equipment to meet other projects' sometimes unpredictable schedules. This, of course, implies that those sub-systems most depended on by other sub-systems (for example, network and timing) will need to be financed earlier.

One financial reason was an uncertain funding profile (as described in [1]). Purchasing everything at once

*Work supported by US DOE under contract DE-AC52-06NA25396
bjo@lanl.gov

guarantees that at least some part of the accelerator will be improved if funding dries up the next year.

OBSERVATIONS AND RECOMMENDATIONS

As of this writing (September 2015), we are currently coming to the end of the fifth year startup period. This is approximately the mid-point of our upgrade project. We are now far enough along in the project to a) make us think we have at least some idea about what we are doing, and b) provide us with a little hindsight into what worked well for us and what didn't. With this in mind, we offer a few of our general observations and recommendations – with specific examples from the timing system upgrade project.

Observation 1:

You Can't Replace The Whole System At Once

In fact, in most cases you can't even replace a whole subsystem at once.

Admittedly, this is a pretty obvious observation. After all, we did schedule nine years for the project! But the implications of this observation can sometimes be less obvious. Upgrade tasks need to be broken up into sub-tasks that are small enough to be accomplished during the scheduled outage periods. Interfaces to other projects need to be considered along with the other projects' installation schedules. Long lead-time equipment needs to be budgeted and acquired in time to meet the installation schedules.

The timing system upgrade installation is complicated by the fact that it interfaces with so many other systems – each of which have their own installation schedules. We were, however, able to come up with the following general plan:

- First, install the event link distribution infrastructure. Here we were able to use the same fiber-optic cables as the network distribution. So that part was easy.
- Install the new timing pattern generator and use it to generate timing signals for a small number of other upgrade project installations. This was probably the most difficult part of the project.
- Interface the entire machine protection system to the new timing system. This may well be the most ambitious part of the project. Once done, however, it will make life easier for machine protection, industrial I/O, and the beam-synchronized data systems.
- Provide timing for other projects based on their installation schedules.

Observation 2:

Some Compatibility Must Be Maintained Between The Old And New Systems

This follows from Observation 1. If you can't replace an entire system in one outage, then you will have to run your machine with both the new and old systems working

in parallel. For us, this reality prompted the following question:

"Can one accelerator be served by two timing masters?"

In our case, the initial answer was "yes", but the ultimate answer was "no".

Early on in the upgrade project, the only equipment requiring the new timing system was the new wire-scanner system. The new wire scanners required only one timing output signal that could be switched between one of five different beam gates. The Micro Research Finland system has the ability to sample up to eight signals and replay them across the event link. This allowed us to "slave" the new timing system to the old timing system by simply digitizing the desired gates. When we started installing the low-level RF systems, however, we needed to provide more than just eight gates, so slaving through sampling was no longer an option.

We tried running the old and new timing masters in parallel, but in the end we could not keep the AC zero-crossing circuitries synchronized and the jitter between the two systems was unacceptable.

What we finally ended up doing was constructing a big 15-slot VME system with ten 16-gate event receiver modules and programmed it to generate all 96 of the original timing gates (plus a few ancillary gates). We were then able to connect the gates replicated by the new timing system to the old system's distribution network.

Instead of slaving the old timing system to the new timing system, we slaved the old distribution network to the new timing system.

Recommendation 1:

Always Have A Way To Fall Back

We have found it prudent to keep the old equipment around for at least a year while we work out the kinks in the new equipment. The first year we installed the new timing system, we ran both systems in parallel – each system producing exactly the same timing gates. As we mentioned above, the jitter between the two systems was unacceptable, so only the new system was connected to the distribution network. However, if for some reason the new system failed, we could easily switch back to the old system by simply relocating four ribbon cables.

Sometimes you may have to fall back even if your equipment is working perfectly. Within a week after installing the new timing system we got a request to put the old system back because of continuous and unexplained machine protection faults. This posed a problem for us because the new low-level RF systems (also installed that year) needed timing features that were only available from the new timing system. Fortunately, we were able to resolve the problem, but if we hadn't we would have been required to roll back to both the old timing system and the old low-level RF system.

Sometimes, the fall back does not have to be to the old system. One useful strategy we have found for the timing system was to have a separate, redundant, set of hardware

we could switch over to whenever we needed to do maintenance or a software update on the system.

Observation 3: You Will Be Surprised

One thing you will be surprised at is how long old technology can keep running! It can continue running long after its designers have retired, long after the original implementers have left, and certainly long past the time that any spares are still available. Obviously this equipment is ripe for replacement, if only someone could remember how it works!

We are now the second and third generation of engineers and programmers to work on this accelerator. Frequently, when asked why something is done a particular way, the only answer we can give is “STOLA” – which stands for “Sacred Tradition whose Origins are Lost in Antiquity”. During an upgrade, however, the antiquities resurface and the origins are revealed.

Sometimes we start out “knowing” how the equipment works only to discover hidden design “features” when we try to replicate its functionality. Even more insidious are the undocumented inter-system dependencies waiting to be uncovered.

After we got the machine protection problem sorted out, we went back to work on the low-level RF. Once we reached the point where we were ready to try sending beam, we suddenly started getting machine protection faults again. What we found was that the “I’m OK” signal the low-level RF system sends to the machine protection system was being derived from a timing signal generated locally by the new timing system. However, the masking gate used by the machine protection system to determine when to look for the “I’m OK” signal was coming from the old distribution system. Even though both these gates originated in the new timing system, the skew between the old distribution system and the locally generated gates was enough to cause the fault.

The lesson learned here was that all the gates going to a given system should come from the same source. In fact, it might be best if all the gates in a given geographical area came from the same source.

Recommendation 2: Sympathy For The Operations Staff

Accelerators are complicated machines and change is hard. Even a change that simplifies operation will initially make operation more difficult simply because it is different.

With the new timing system, we changed from a gate-oriented system to an event-oriented system. There are a lot of things an event system can do better than a gate-oriented system. There are also a lot of things that a gate-oriented system can do better than an event system. The gain of new capabilities from a new system is often

eclipsed by the loss of accustomed capabilities from the old system. The difference can sometimes be dramatic. In one instance, the change to the new timing system completely altered the way an entire section of the accelerator behaved because of a change in where the beam chopping occurred! These types of changes can seriously impede a startup period if the operations staff is not kept in the loop.

Below are a few suggestions for keeping the operations staff up to speed:

- Training sessions on what has changed during the last maintenance period.
- Involve operations personnel in design reviews – especially regarding the operator interface.
- Involve operations personnel in the installation of the new systems.

Training sessions are certainly important, and many facilities (ours included) will have a “changes meeting” shortly before a startup period begins. A single “changes meeting” is certainly the most efficient way of communicating what’s new, but it can sometimes be quite lengthy which does not contribute to retention.

Involving operations personnel in the design and installation activities can be very productive, both for the operators and the systems engineers. Operators tend to have a more global perspective than the system engineers and can help spot those inter-system dependencies that a system engineer might miss.

CONCLUSION

As mentioned above, we are now around the halfway point in our upgrade project. There will no doubt be many more observations to make and recommendations to share in the coming years. In the meantime, we hope these observations have been useful – especially if you are contemplating a similarly ambitious upgrade project. If you are, however, you might want to check back with us at the 2019 ICALEPCS.

REFERENCES

- [1] M. Pieck et al., “LANSCE Control System Upgrade Status and Challenges”, MOPGF161, these proceedings, ICALEPCS’15, Melbourne, Australia (2015).
- [2] S.C. Schaller, “A LAMPF Controls Retrospective”, *Proceedings of the Third International Conference on Accelerator and Large Experimental Physics Control Systems* (Elsevier, 1994), p. 516
- [3] M. Pieck et al., “MaRIE – Instrumentation and Control System Design Status and Options”, MOPGF162, these proceedings, ICALEPCS’15, Melbourne, Australia (2015).
- [4] Micro Research Finland website: <http://www.mrf.fi>