Operational Experience from a Large EPICS-Based Accelerator Facility *

Daniel J. Ciarlette and Rodney Gerig
Advanced Photon Source, Argonne National Laboratory
9700 South Cass Avenue, Argonne, Illinois 60439

I. Abstract

The Advanced Photon Source (APS) at Argonne National Laboratory is a third-generation x-ray light source which uses the Experimental Physics and Industrial Control System (EPICS) to operate its linear accelerator, positron accumulator ring, booster synchrotron, and storage ring equipment. EPICS has been used at the APS since the beginning of installation and commissioning. Currently, EPICS controls approximately 100 VME crates containing over 100,000 process variables. With this complexity, the APS has had to review some of the methods originally employed and make changes as necessary. In addition, due to commissioning and operational needs, higher-level operator software needed to be created. EPICS has been flexible enough to allow this.

II. INTRODUCTION

EPICS has moved from the test stand to large accelerators. In making this transition, EPICS has grown from a few input/output controllers (IOCs) and operator interfaces (OPIs) into a major control system. The APS is one of the first large EPICS installations to use EPICS from the initial R&D phase of the project to the operational phase. This growth has magnified some of EPICS' strengths and weaknesses, forcing APS to review and rethink some of EPICS' original features. Some features that will be highlighted here are:

- Adding IOCs and OPIs at will
- Problems with IOC reboots
- Process variables
- Use of MEDM screens
- Using Channel Access Security
- Using the channel access libraries with user software
- BURT files and utilities
- Layering software on top of EPICS.

III. DESCRIPTION OF THE APS

The APS consists of an injector complex located within a storage ring (Figure 1). The injector complex consists of two linear accelerators and two circular machines called the positron accumulator ring (PAR) and the booster synchrotron. Each machine consists of a number of subsystems such as: power supplies, radio frequency, diagnostics, and vacuum. Many of the subsystems employ the use of EPICS subnetworks such as Allen Bradley, VXI, BitBus, and GPIB. The five machines and subsystems are all controlled by EPICS.

The EPICS control system is the means by which the APS controls all of the equipment to produce a 7-GeV electron beam in the storage ring. To do this, electrons are produced by an electron gun and accelerated through the first linear accelerator to 200 MeV and hit a target, producing positrons. These are accelerated to 400-450 MeV in the second linac. The positrons are then transported to the PAR where they are bunched and compressed. Then the bunches are transported through the PAR-to-booster (PTB) transport line to the booster which accelerates the bunches to an energy of 7 GeV. After the positron bunches have achieved an energy of 7 GeV (in about 0.25 sec), they are transported by the booster-to-storage ring (BTS) transport line into the storage ring. Here, the positrons circulate around the storage ring as a stored beam at an energy of 7 GeV.

IV. APS EPICS HARDWARE

The Advanced Photon Source's EPICS IOCs consist of VME and VXI crates with Motorola 167 and 162 CPU single-board computers. The OPIs are Sun SPARCstation-10 workstations with 128 Mb of random access memory. The control system network consists of a fiber optic ethernet connected in a star configuration (Figure 2) [1]. Portions of the fiber optic ethernet system use redundant fibers and transceivers connected to ethernet hubs that support redundant fibers. The workstations are connected to the file server computer and network hub via an FDDI network [2].

* Work Supported by U.S. DOE Office of Basic Energy Sciences under Contract No. W-31-109-ENG-38.

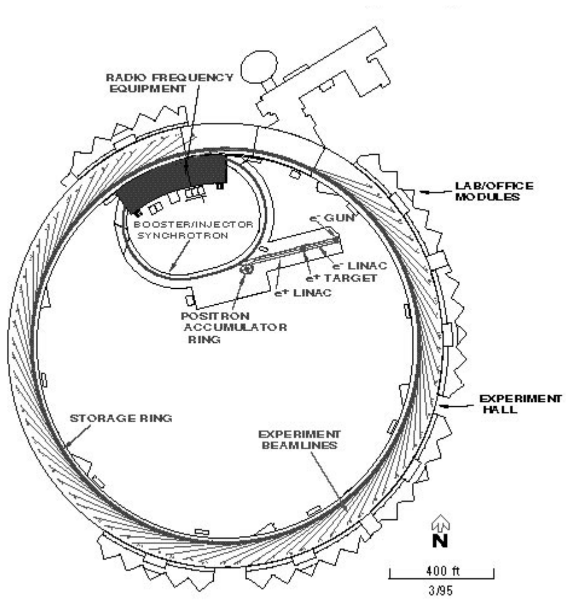

Figure 1. APS beam acceleration and storage complex

V. IOCS AND OPIS

A feature of EPICS that has proven to be useful is the distributed nature of the system. Since there is no main controlling computer behind EPICS, each IOC and OPI operates independently. This distribution allows IOCs and OPIs to be added or removed without interruption of service to other IOCs and OPIs. In fact, entire machines at APS have had the EPICS software upgraded on the IOCs without any interruption of services on the other machines. This significantly enhances the reliability of the total system and is a great benefit to operations. For example, a board being replaced in a linac IOC will not affect any beam stored in the storage ring. Also, since each OPI is generic, if an OPI problem occurs, the operator can easily move to another OPI to continue operation. Distribution allows parts of the total system to be changed easily without affecting other EPICS computers.

While rebooting an IOC does not disturb other IOCs, it can have a pronounced effect on the functioning of the accelerator. APS has had problems restoring process variables after the IOC reboots. During the design and commissioning phase, an IOC reboot did not present much of a problem. Now, during the operational phase, an IOC reboot can result in a large amount of work. For instance, an IOC reboot may cause beam to be lost because default values were loaded into the database rather than the last settings. If a subsystem such as rf or power supplies has a problem, the time to recovery increases. The magnets associated with a power supply IOC may require restandardizing or degaussing before known values can be input. The rf systems may need to be restarted again. In addition, due to the loss of beam, reinjection from the injector machines would be required. Therefore, applications developers are trying to minimize the effect that an IOC reboot has on a working system.

Some work in minimizing IOC reboot problems has been done at APS. After a reboot of a PAR, booster, or storage ring rf IOC, the IOC examines a backup and restore tool (BURT) snapshot file and loads in the required constants. This eliminates a large annoyance that previously occurred. In the past an IOC reboot, especially an undetected IOC reboot, would cause calibration constants and setpoints to be reinitialized to the database defaults. These defaults are not necessarily always correct. By having the IOC reload these from a BURT file with the correct constants, the limitations of the database are resolved. Another approach to this problem was taken with the linac IOCs. The IOCs controlling the linac will read the setpoint values from the hardware itself. The IOC does this by examining subnetworks, such as Allen-Bradley and GPIB that the IOC uses to control the linac and interrogates them to find the last written settings. The IOC then continues operation with the interrogated settings as the current settings. The rf and linac IOCs us these two different approaches to solve the reboot problem.

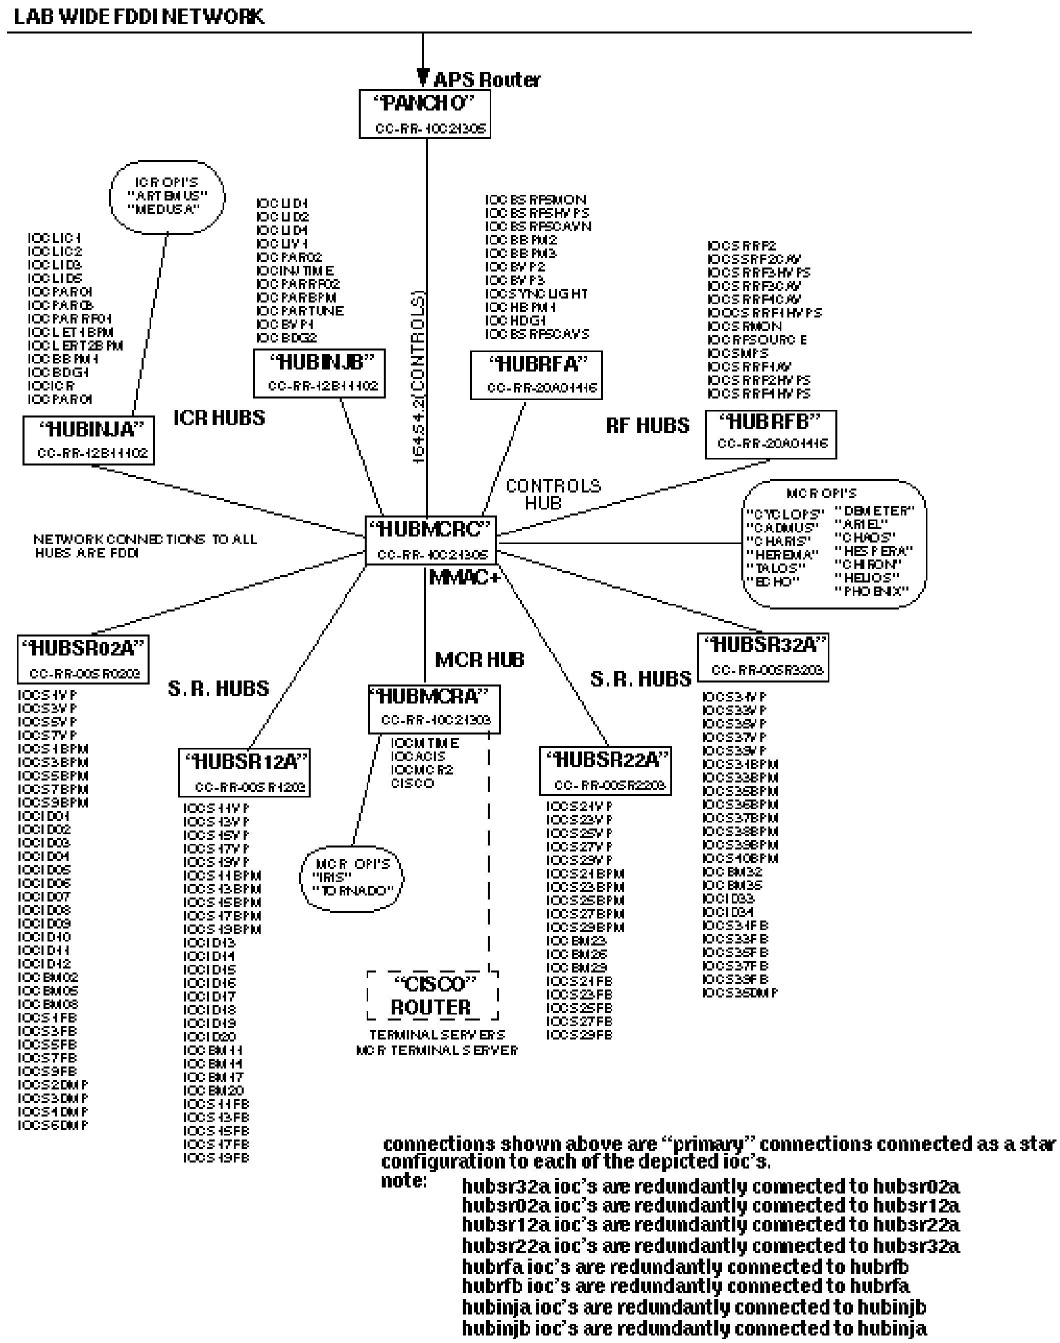

Figure 2. APS controls network topology

VI. PROCESS VARIABLES

EPICS uses records to store the data that is read and written. Each input or output channel used in the control system is contained in a record. There are currently over 100,000 process variables in use at the APS.

While each EPICS database record is unique, the process variable name can only be 28 characters long. This presents a problem when trying to both uniquely and descriptively identify a database record. At the APS a naming convention, or nomenclature, was introduced for each machine to identify each piece of equipment and database record. Unfortunately, even the best of nomenclatures cannot easily describe every signal in a control system

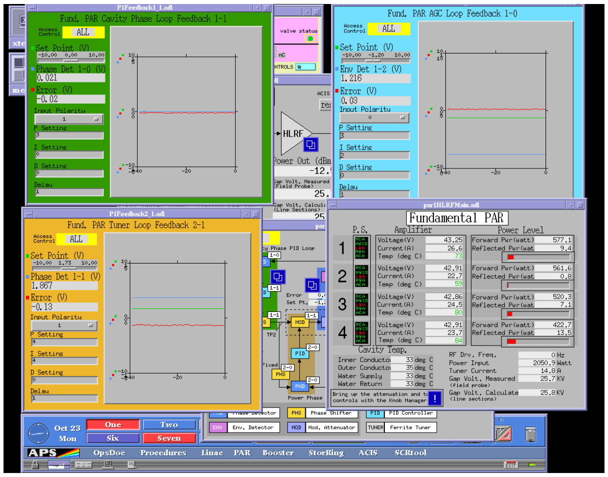

Figure 3. A cluttered set of PAR rf MEDM screens

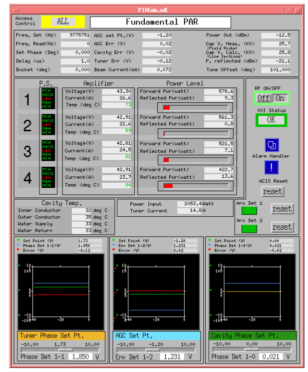

Figure 4. A redesigned PAR rf MEDM screen

within 28 characters. To help to alleviate this problem, the DESC (description) field of each record contains an additional 28 characters of information about the record. Care should be taken to utilize the DESC field wisely to assist others who are using the record.

While the DESC field allows an application developer to add a 28-character description to the record, it does not always give enough information. Sometimes the fields within a record also need descriptions and longer descriptions would be useful as well. A process variable database would be very helpful in alleviating these restrictions. A process variable database on the control system file server that is keyed by the process variable name could contain a better description and more information such as: documentation references, drawings, schematics, and on-line manuals. Ideally, this would be a source of information that would enhance the knowledge of operators interested in a particular area and would aid in troubleshooting when problems arise.

VII. OPERATIONAL SCREENS

APS uses the Motif-based Editor Display/Manager (MEDM) to graphically display control system process variables. During design and commissioning of the APS accelerator subsystems, engineering screens were created to satisfy the needs of engineers and scientists who needed as much information as possible at their disposal to commission the equipment. However, MEDM screens containing large amounts of information cause problems during the operational phase.

Now that the systems are commissioned and operating together, the abundance of information being displayed requires the use of too many MEDM screens. This results in screen clutter (Figure 3). The standard operational setup for the PAR and linac, for instance, requires that ten virtual screens be created for each machine in order to operate them on separate OPIs. Operating both machines on one OPI would cause undesirable delays in response, meaning that vital information may not be seen when other virtual screens are being viewed. Finally, the display of unneeded process variables would generate more channel access traffic than necessary. Therefore, a simplified set of screens with less channel access traffic was needed to operate the APS. Figure 4 is a redesigned version of Figure 3 that was created by the Controls Group to eliminate most of the unneeded information and to reduce screen clutter.

The Operations Group has begun work on creating operational screens. The process requires the following steps:

- Determine how the operators use parts of the current engineering screens and how frequently.
- Analyze this information with the appropriate groups and Controls Group application developers.
- Design and create new screens when needed.
- Review the new screens with the operators, appropriate subsystem groups, and Controls Group application developers.

Since this process requires an overall understanding of most of the engineering screens at the APS, this endeavor has only been carried out recently. This has given the Operations Group a comprehensive understanding of how each set of engineering screens function.

The operations screens alone will not solve all of the APS screen usage problems. The engineering screens will still need to be used by the subsystem engineers and operations group personnel when a problem arises that needs more information than the day-to-day operations screens contain. The operations screens therefore should have related display buttons to easily display the engineering screens. By designing a set of standard operational screens that only contain the information needed to operate the APS, the operator overload will be reduced and the machine load will stay within responsive limits while still maintaining the ability to quickly view everything needed for operations.

VIII. SDDS BASED TOOLS

APS has adopted the Self Describing Data Set (SDDS) file format [3]. While SDDS is not directly part of the control system, a significant fraction of the data that is read and written by the control system during commissioning and operations is in SDDS format. All of the commissioning backup and restore tool (BURT) and experiment data files are in SDDS format. In addition, the Operations Group BURT files use the SDDS file format as well. The SDDS suite of programs also includes channel access support. Some of the SDDS-based programs use by the Operations Group and physicists are sddsexperiment [4], sddssnapshot, and the Operations Group backup and restore tool graphical user interface (GUI) named scrtool [5]. EPICS' "open communication" feature allows others to write layered and integrated applications easily.

IX. BURT

BURT is the EPICS backup and restore tool. BURT's backup facility takes a list of process variables and creates a process variable/value pair in a snapshot by reading the values from the running IOCs. Conversely, BURT's restore function takes a list of process variable/value pairs and restores them to the IOCs.

The Operations Group found that BURT was lacking some important features. Rather than redesign BURT to add selective restores and a GUI only used by operations, the choice was made to add a another layer of software that processed SDDS BURT snapshots. The result of this work was the save, compare, restore program named scrtool. Without any changes to the control system, a totally new application was able to save snapshots, restore portions of snapshots, and compare snapshots with process-variable-defined tolerances. By utilitizing BURT, the scrtool can take advantage of new features added to BURT, bug fixes, and new EPICS releases. This software used a layered approach with both EPICS and SDDS tools to accomplish what was needed.

X. CHANNEL ACCESS SECURITY

During the design and commissioning period, EPICS did not provide security via channel access. This did not prove to be a problem since most of the activity was located in our linac gallery and injection control room and most of the work was tightly coordinated. Now, during the operational phase we are finding that users are geographically distant and more of them are logging in to monitor their subsystems. This raises the possibility of people accidentally changing things.

Channel access security is being used initially with the booster and storage ring rf systems. If everything goes well, it will be implemented in all of the APS at a later date. The main feature of channel access security is to restrict what people can do at each workstation. Channel access controls which groups of users can access the process variable database while limiting the computer from which a user can access the data. When fully implemented at APS, channel access will allow the more secure system which will be required when we are in a user service mode of operation.

XI. WHAT IS RUNNING WHERE

While channel access allows EPICS to control who can access the database from various workstations, it does not solve all access problems. One situation that has not been a big problem at APS, but is a real threat, is not knowing what programs are running on which computers. EPICS, being a distributed control system, does not have a central repository of information about what is running on a single computer. This problem is exacerbated when a program, such as a feedback program running in the background of an OPI, contends simultaneously with another version of the same program. There are many solutions to a problem such as this. Two of these solutions are:

- Use process variables to grant permission for a program to run/not run.
- Use a license manager to allow programs to run/not run similar to commercial software.

The latter solution, a “channel access license manager,” would require the program to check out a license in order to use the control system. The license manager would have the intelligence to:

- allow one or more copies of programs to run, depending on how the license manager is configured or if the program checked out the license, e.g., checking out a license as “exclusive,”
- determine if the machine is in a mode in which the program is allowed to run,
- log license usage of each process,
- report what is running and what programs ran previously,
- allow secure users to be able to terminate or suspend processes through the license manager,

Some of these uses may be mediated by the use of process variables, and some may not. Regardless of the implementation, some feature such as this is needed to enable operations to monitor and log the processes that are running or have been run which could affect operations.

XII. DOWNTIME

One important metric for a systems’s operations is the amount of downtime attributed to it. During the period of January 1, 1995 through July 31, 1995, the Controls Group accounted for 31.3 hours of downtime, only a portion of which was attributed to EPICS. Most of the EPICS-related downtime was due to OPIs logging themselves out, IOCs hanging up, network links between OPIs and IOCs failing, and the occasional engineer not telling operations when they were going to work on an IOC. EPICS has proven to be a very reliable control system.

XIII. CONCLUSION

EPICS was ready and running at the APS before the commissioning stage started; therefore, when changes or new features needed to be added, the Controls Group was able to respond quickly. In addition, APS has gone through several EPICS upgrades with virtually no problems. EPICS has proven to be stable and ready on schedule.

While EPICS has proven to be a versatile control system, it is not without its limitations. Some of these limitations stem from its distributed nature, while others are manpower limitations that all EPICS installations face. By identifying these limitations, APS and other collaboration members may be able to provide solutions. One of EPICS’ great strengths is the collaborative effort; in time, solutions to these deficiencies will be solved by the APS EPICS developers and others. This will allow the Operations Group to have an even more useful control system.

XIV. REFERENCES

[1] G. Gunderson, private communication

[2] K. V. Sidorowicz and W. P. McDowell, “The APS Control System Network,” these proceedings.

[3] M. Borland and L. Emery, “The Self Describing Data Set File Protocol and Program Toolkit,” these proceedings.

[4] M. Borland, L. Emery, and N. Sereno, “Doing Accelerator Physics Using SDDS, UNIX, and EPICS,” these proceedings.

[5] D. Ciarlette, private communication