THE SRS CONTROL SYSTEM: 25 YEARS OF OPERATION AND DEVELOPMENT

B.G. Martlew*
CCLRC Daresbury Laboratory, Warrington, UK

Abstract

The SRS has recently celebrated 25 years as an operational synchrotron light source and will continue to receive central funding for another three years. Throughout this lifespan of nearly 30 years, both the accelerator and the control system have undergone a number of major upgrades as well as following a path of continuous development and improvement.

When the SRS control system was first conceived in the late 1970s, many of the technologies now taken for granted such as PC workstations, local area networks (LANs), intelligent instrumentation and graphical user interfaces (GUIs) were either in their infancy or undreamed of. It is to be expected that similar, if not more dramatic, changes will take place in technology over the next 30 years. The control system designer faces similar challenges today – the need to design systems that can expand to cope with the developments of the accelerator as well as being able to embrace new computing hardware and software with minimum disruption to operations. Also of vital importance is the need to consider long-term maintainability especially with ever decreasing product lifetimes in the modern electronics and computing industry.

This paper looks back at some of the changes that have taken place in the SRS control system since the opening of the accelerator and highlights some of the successes and failures from an evolutionary point of view. It also considers what lessons can be learned for the design of future accelerator control system hardware and software.

INTRODUCTION

In the 1970s early pioneering work on exploitation of synchrotron radiation via a parasitic beamline on the NINA 5 GeV electron synchrotron led to a proposal to build the world's first facility dedicated to the production of synchrotron radiation. By the time NINA closed in 1977, funding was already approved for the new Synchrotron Radiation Source - later to become know simply as the SRS – the first of the new 2nd generation light sources. This machine was to be centred round a 2 GeV electron storage ring [1]. Design of the SRS commenced in 1975 and first operation took place 5 years later in 1980.

The SRS accelerator complex consists of a 12 Mev linac, a 600 MeV booster synchrotron and the 2 GeV storage ring. Normal operation consists of the storage ring being filled with 600 MeV electrons from the booster synchrotron once every 24 hours. Once sufficient electrons have been injected, the energy of the storage ring is slowly ramped up to its normal working point of 2 GeV.

Experiments officially began in 1981 and by the end of that year four beamlines were operational supporting protein crystallography, EXAFS and topography. By 1982, nine experimental stations were in use with circulating beams of 350 mA at 2 GeV being regularly achieved in the storage ring. This year also saw the first use of single-bunch mode where only 1 of the 160 available RF buckets in the storage ring is filled with electrons permitting time resolved measurements to be taken.

The early beamlines on the SRS exclusively used dipole radiation from the main storage ring bending magnets. Increasing interest in insertion devices such as wigglers and undulators to increase flux at short wavelengths or to provide high brightness at specific spectral points led to the inclusion of such devices. Initially, the SRS gained a single 5T superconducting wiggler. This was later followed by a 6T superconducting wiggler and, more recently, by a number of fixed magnet undulators and multipole wigglers. The most recent addition to the collection of installed insertion devices is a variable polarisation undulator to the APPLE II design. The present line-up is: 2 superconducting wavelength shifters, 3 multipole wigglers and 1 variable polarisation wiggler.

* B.G.Martlew@dl.ac.uk

In 1987 a major upgrade of the storage ring optics was undertaken to meet the increasing demand from users for high brightness rather than raw flux. This entailed a complete re-build of the magnet lattice and took nearly 12 months to complete. Because of the reduction in beam size from this upgrade, the early 1990s saw considerable interest in the development of feedback techniques to improve beam stability.

In the late 1990s a proposal was developed for a new 3rd generation light source in the UK to ultimately replace the ageing SRS. This proposal ultimately became Diamond which is currently being constructed at the Rutherford Appleton Laboratory. Diamond will be a state-of-the-art 3rd generation light source based around a 3 GeV electron storage ring.

It has recently been announced that government funding for the SRS will cease at the end of 2008 bringing to a close nearly 30 years of world-class synchrotron radiation research on the SRS.

A HISTORY OF THE SRS CONTROL SYSTEM

It is interesting to note that in the initial SRS Design Study [1] written in 1975 a large part of the control system chapter was devoted to justifying the decision to use a fully computerised design. The main argument was one of economics rather than functionality; it was argued that the cost of dedicated wiring for hardware controls would be greater than that of a general purpose, expandable, computerised control system. So, on the basis of that seemingly obvious, to modern eyes at least, statement the SRS control system was born.

The original system was based on a two-level design (Figure 1) built using minicomputers from the Perkin-Elmer (Interdata) 7/16, 8/16 and 7/32 range [2]. The 7/16s and 8/16s were 16-bit CPUs with 64 Kb of magnetic core memory running Perkin-Elmer's propriety multi-tasking operating system, OS/16-MT. The 7/32 was a more powerful 32-bit design with 320 Kb of core memory together with two 10 Mb removable disc drives and a line-printer. The operating system was the 32-bit OS/32-MT. Access to the plant was through four CAMAC serial highways communicating with a number of geographically distributed CAMAC control stations. Additional in-house development had to be undertaken on the operating system to add support for networking, CAMAC and plant I/O – none of which were available from the manufacturer.

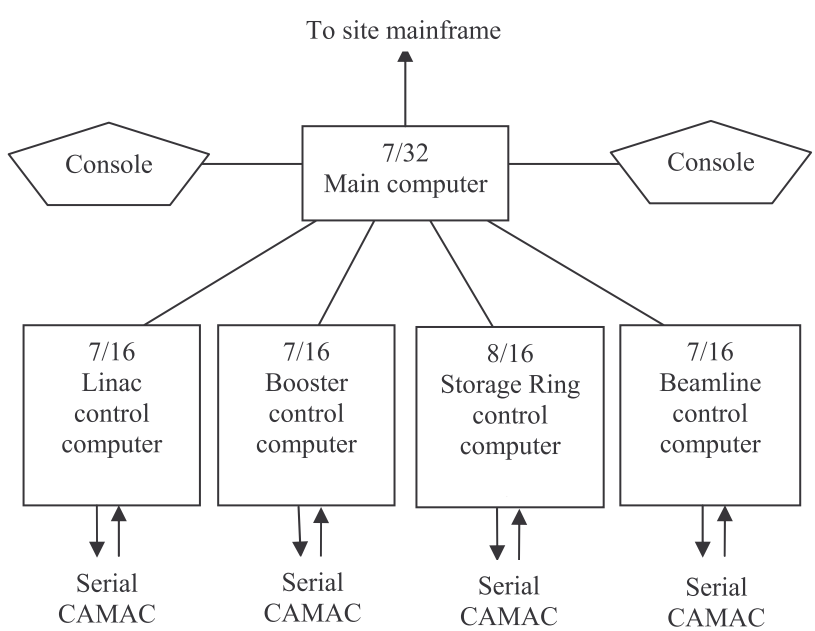

Figure 1: Original control system architecture.

The user interface hardware was driven from the 7/32 through a parallel CAMAC interface. It consisted of a text-mode colour display with keyboard and tracker ball, a vector graphics display and an optical control knob for fine adjustment of analogue parameters. Three of these combinations were provided in the main control room. Most user interaction took place using the “Control Desk” application which displayed a list of user selectable devices on the screen together with their principal properties such as analogue value, status and state of interlocks. This application was designed to be as flexible as possible by allowing the operator to select devices for display either from predefined pages or directly from the command line with wild-card support. Several other general purpose applications were also available for trend plotting, execution of script files and retrieval of archive data.

The method of interfacing plant to the control system was very carefully defined. It involved developing a custom status and interlock monitoring system that was capable of operating and maintaining plant integrity independently from the upper levels of the system. This system (Mk. I status system) was based on TTL logic and electromechanical relays and was a very successful and reliable part of the control system. However, with the wide spread availability of microprocessors in the mid 1980s it was upgraded using Motorola 6809 and CMOS logic. This became known as the Mk. II status system but maintained the same basic structure and philosophy of the earlier version [3]

The next significant event in the life of the control system came in 1985 when the 7/32 and 7/16 minicomputers were replaced by Concurrent 3200 series 32-bit processors. These machines were software compatible with the old processors and required no changes to the overall control system architecture. They introduced performance improvements and simplified support by removing the need to maintain the 16-bit version of the operating system.

Following the upgrade to the SRS magnet lattice described above, new systems were developed to provide enhanced monitoring and feedback of the electron beam position in the storage ring. This project saw the first use of VME and Ethernet LANs in the SRS control system. It has been described in some detail elsewhere [4, 5, 6].

By the start of the 1990s it was becoming clear that a major upgrade of the control system hardware and architecture would soon be required. The existing combination of minicomputers and CAMAC serial highways could no longer provide the performance that was increasingly being expected from the control system. The user interface was still command line driven and Concurrent was starting to warn that the long-term supportability of the 3200 processors could only be guaranteed until 2000. All these factors prompted an investigation of alternatives leading to a decision to adopt a modified form of the ISOLDE control system from CERN [7]. This upgrade was a very complex operation involving complete redesign of all application software using Visual Basic and Windows NT as well as migrating the overall system architecture to the client/server paradigm. Work started in 1992/3 with first operational use on the main magnet power supplies taking place in 1995 [8]. Between 1995 and 2002 sub-systems were gradually transferred from the Concurrent minicomputers to the distributed system starting with the beamline control system followed by the linac, booster, storage ring and, finally, monitoring of the personnel safety system (Figure 2).

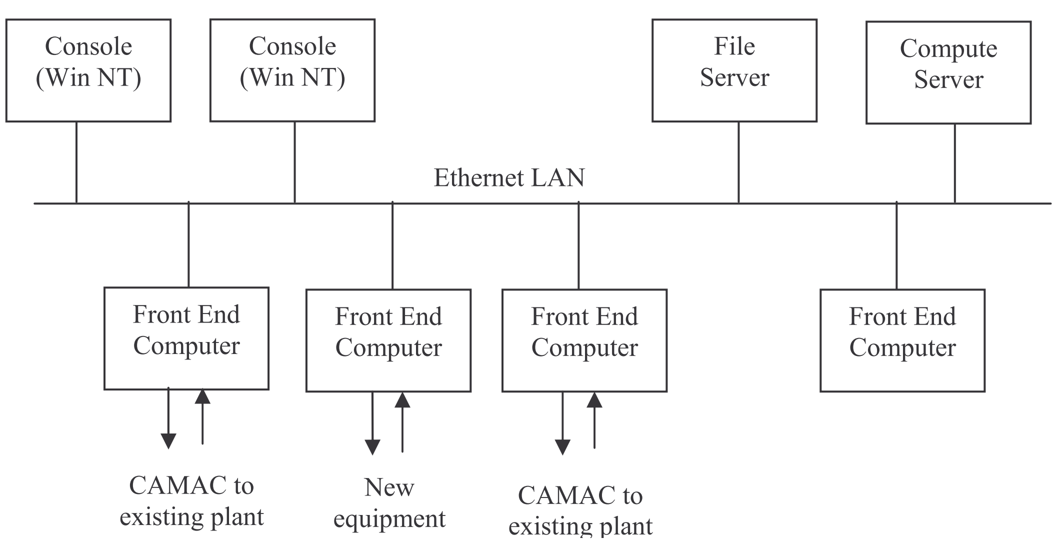

Figure 2: Client/Server architecture of the 1990s

After the introduction of distributed intelligence the next development was the design of standardised control racks containing 3U VME front end computers using Motorola 68040 processors and running the OS-9 real-time operating system. This arrangement allowed control electronics, status, interlocks and instrumentation all to be housed in a single rack requiring only power, network and field cabling to be supplied externally. Coupled with this was the evolution of the status and interlock system to a Mk. III variant utilising microcontrollers communicating over CANbus [9, 10]. This approach has the advantage that the rack can be designed, built and tested in the development laboratory and only integration and commissioning is necessary after final installation in the accelerator.

In the last few years the main changes have focussed on upgrading a number of obsolete MS-DOS front end computers which have been replaced by industrial PCs running Linux [11]. Also, a new RF klystron power supply with integrated EPICS controls has been connected to the control system. This was achieved by writing a gateway process that converts internal SRS control system requests into Channel Access for forwarding to the power supply [12].

THE GOOD…

Accelerator control systems commonly have a lifetime of several decades. The SRS is typical in this respect with a total expected life of nearly 30 years. During this period it is certain that several major upgrades of system hardware will be required and, hopefully, these changes will be possible with minimal disruption to operation of the accelerator. Despite the opportunity for change, many aspects of the control system are fixed by the original design philosophy and can significantly influence the ease with which subsequent upgrades can be undertaken and minimize the operational disruption. A number of good design decisions were taken back in the 1970s which helped simplify the development of the SRS Control System. These are briefly described below.

Clear and Consistent Naming System

Naming conventions for accelerator systems are often overlooked but are a very important tool in designing, developing and operating a successful control system. With typical installations containing thousands of physical devices and tens or hundreds of thousands of individual I/O signals it is clearly impossible to remember the function and location of each one. The original SRS naming convention used a very simple but effective scheme: <A><T>.<DDDD>.<XX> where: <A> is a single letter indicating major machine area, <T> is a single letter indicating technical area, <DDDD> is a four letter device type and <NN> is a two digit serial number. So, a typical device name is SV.PIRG.03 which translates to **S**torage Ring, **V**acuum, **PIR**ani **G**auge, device number **03**. This scheme, although limited to only 10 characters by memory limitations in the early minicomputers, has proved to be flexible and effective and with some enhancements is still in use today.

Clear Plant Interface Guidelines

Plant interface guidelines define two things: 1. where the control system stops and the plant starts and, 2. the protocols and signal levels to be used to communicate between the control system and the plant. The SRS attempted to address both questions by defining a very precise concept of how a typical device operates and is interfaced to the control system. It was assumed that every device could potentially have some, or all, of a set of standard interfaces: +/- 10V analogue input, +/- 10V analogue output, a number of distinct status modes (On, Off, etc.) and up to 16 discrete interlock signals. Conformance to this standard was enhanced by designing specialised hardware to manage the status control and interlock handling functions. This approach also helped enforce an objected-oriented style to the control system with every device supporting a standard set of interfaces and properties, thus allowing generic application software to be written (see below). In recent years, considerable flexibility has been required in order to support intelligent devices, serial interfaces etc. but the same general principles have been maintained thus ensuring compatibility with existing software.

Generic Application Software

Following on from the object-oriented hardware interface described above is the ability to produce generic application software. In the early days, the SRS control system offered very limited graphics capabilities and memory was at a premium therefore it was decided to design a limited set of general purpose applications for operations. The primary application was Control Desk, a user configurable, tabular display of device values, status and interlocks. This single program has remained the core of the operations toolkit to the present day. It has been ported, developed and enhanced several times during its life, but would still be immediately recognizable by an operator from 25 years ago. Many synoptic displays are now available but are little used in the control room; the operators preferring the familiarity and flexibility of Control Desk. Another long-lived application is SRS History, a general archival and retrieval facility. Every two minutes the state of the entire accelerator is monitored, processed and archived to a database. The two minute scan was chosen 25 years ago as a compromise to ensure that sufficient information is retained without incurring excessive data rates and volumes. Even though these are no longer serious limitations, the two minute scan is still retained because it has proved to be perfectly adequate for post-mortem analysis of most faults.

Other generic programs that have changed little over the 25 year history of the SRS are trend charting and sequence scripting.

…THE BAD…

No system is without its failures. This section highlights some of the areas of the SRS Control System that have not been as successful as hoped and also documents an omission that has ultimately limited the overall performance of the system.

Unreliable Alarm system

During the life of the SRS there have been a couple of attempts to implement a comprehensive alarm handling system. Both systems have failed to gain the confidence of operators and engineers for a number of reasons:

- Unreliability of the underlying signals from the plant. An alarm system is only as good as the data it monitors and intermittent or noisy signals can produce spurious alarms that reduce the confidence in the overall system.
- Lack of alarm hierarchy. Unless alarm messages are prioritised and arranged in a clear hierarchy it is all too easy for important alarms to be lost amongst a forest of unimportant and unrelated warnings.

Lack of asynchronous communications

The system adopted from ISOLDE in the 1990s supported both synchronous and asynchronous communication. The asynchronous mechanism used Dynamic Data Exchange (DDE) as the client API and was found to be difficult to use. As a consequence, it was decided on the SRS to drop support for this aspect of the software. In hindsight, this was a mistake as it has limited the performance of the control system and resulted in a number of complex and application-dependant attempts to simulate asynchronous communications in client software.

…AND THE UGLY

Finally, a couple of parts of the system that would, in hindsight, have been done differently. Both the examples mentioned here have no impact on the operational success of the system but have created support difficulties.

Fragmented databases

Configuration databases for the SRS control system are very fragmented with separate text databases for every front end computer and a different binary format database for use by client software. Similarly, there are numerous application configuration databases in a variety of formats. Although there is no performance penalty with this approach there are serious problems with support and maintenance.

Isolated steering system architecture

The beam steering upgrade was designed and developed during a period of major change in the SRS control system and, as a consequence, it is not well integrated into the main control system. This has resulted in a rather clumsy design that does not allow the full versatility of the control system to be fully utilised on this sub-system. The main problems are poor update rates when monitoring the system and the inability to make even simple modifications without committing significant development effort.

THE FUTURE

Central funding for the SRS is due to cease at the end of 2008. In the remaining 3 years of operation the number of supported experimental stations will be gradually reduced from the present 22 FTE (Full Time Equivalent) stations to 10 FTE stations on 31st December 2008. In conjunction with this run-down, capital developments and major upgrades will also be significantly reduced. It is expected that work on the control system will be restricted to necessary maintenance and administration together with minor developments where appropriate. No further major developments are expected.

ACKNOWLEDGMENTS

The author wishes to thank the many people, unfortunately too numerous to list here, who have contributed to the success of the SRS Control System over the last quarter of a century. Without their expertise and dedication none of this work would have been possible. However, special thanks go to D.E Poole who was responsible for much of the pioneering design work in the early days of the project.

REFERENCES

[1] "Design Study for a Dedicated Source of Synchrotron Radiation", Daresbury Laboratory Internal Document DL/SRF/R2, 1975
[2] D.E. Poole, W.R. Rawlinson, V.R. Atkins, "The Control System for the Daresbury Synchrotron Radiation Source", Europhysics Conference on Computing in Accelerator Design and Operation, Berlin, Germany, September 1983.
[3] B.G. Martlew, D.G. Peters, D.E. Poole, "A Microprocessor-based Status Control and Interlock Monitoring System for the SRS", Particle Accelerator Conference 1987, Washington DC, 1987.
[4] W.R. Rawlinson, B.G. Martlew, M.J. Pugh, M. McCarthy, "Development of a VME-based Control System for the SRS Orbit Feedback Project", EPAC '94, London, June 1994.
[5] J.B. Fitzgerald, B.G. Martlew, P.D. Quinn, "Operational Performance of the Local Vertical Servo System on the SRS at Daresbury", EPAC '96, Sitges, Spain, June 1996.
[6] J.B. Fitzgerald, B.G. Martlew, C. Morrell, S.L. Smith, "The Daresbury Global Horizontal Servo Feedback System: Status and Outlook", EPAC '96, Sitges, Spain, June 1996.
[7] I. Dellose, A. Pace, "The ISOLDE Control System", ICALEPCS '93, Berlin, October 1993.
[8] B.G. Martlew, M.J. Pugh, W.R. Rawlinson, "Planned Upgrades to the SRS Control System", EPAC '94, London, June 1994.
[9] B. Corker, M.T. Heron, B.G. Martlew, W.R, Rawlinson, "A New Control System for Beamline 5D at Daresbury", EPAC '98, Stockholm, 1998.
[10] J.R. Alexander, B. Corker, S.V. Davis, M.T. Heron, A. Oates, "A CAN based Status and Interlock Protection System for the SRS", ICALEPCS '99, Trieste, October 1999.
[11] B.G. Martlew, B. Corker, G, Cox, P.W. Heath, M.T. Heron, A. Oates, W.R. Rawlinson, C.D. Sharp, "Upgrade of the Photon Beamline Control System on the SRS at Daresbury", ICALEPCS '01, San Jose, November 2001.
[12] A.J. Moss, R.J. Smith, S.A. Griffiths, "Upgrade of the SRS Klystron Power Supply", PAC '01, Chicago, June 2001.