DID WE GET WHAT WE AIMED FOR 10 YEARS AGO?

A. Augustinus, P. Chochula, L. S. Jirdén, M. Lechman, P. Rosinský, CERN, Geneva, Switzerland
O. Pinazza, INFN Sezione di Bologna, Bologna, Italy and CERN
G. De Cataldo, INFN Sezione di Bari, Bari, Italy and CERN
A. N. Kurepin, INR INR RAS – Institute for Nuclear Research of the Russian Academy of Sciences, Moscow, Russia
A. Moreno, Universidad Politécnica de Madrid, ETSI Industriales, Madrid, Spain

Abstract

The ALICE Detector Control System (DCS) is in charge of control and operation of one of the large high energy physics experiments at CERN in Geneva. The DCS design which started in 2000 was partly inspired by the control systems of the previous generation of HEP experiments at the LEP accelerator at CERN. However, the scale of the LHC experiments and the use of modern, "intelligent" hardware and the harsh operational environment led to an innovative system design. The overall architecture has been largely based on commercial products like PVSS SCADA system and OPC servers extended by frameworks. Windows has been chosen as the operating system platform for the core systems and Linux for the front-end devices. The concept of finite state machines has been deeply integrated into the system design and the design principles have been optimized and adapted to the expected operational needs. The ALICE DCS was designed, prototyped, and developed at a time when no experience with systems of similar scale and complexity existed. At the time of its implementation the detector hardware was not yet available and tests were performed only with partial detector installations. In this paper we analyse how well the original requirements and expectations set ten years ago comply with the real experiment needs after two years of operation. We provide an overview of system performance, reliability and scalability. Based on this experience we assess the need for future system enhancements to take place during the LHC technical stop in 2013.

INTRODUCTION

The design of the ALICE DCS started late 2000. In its initial phase the project was able to profit from already advanced developments in other LHC experiments and from tools and guidelines provided by the Joint Controls Project (JCOP) [1]. The existing concepts were compared with ALICE detector needs and from this the first ALICE DCS architecture was designed [2]. The first presentation given to the ALICE Technical Board in 2001 defined the roadmap for the whole DCS. In the following chapters we will review the presented key concepts and compare them with the current system implementation.

ALICE DCS SYSTEM CONTEXT

The architecture of the ALICE DCS is strictly hierarchical. The top level is formed by the central DCS which coordinates the individual detector systems. Each detector system is then divided into sub-systems which group devices with similar functionality. Subsystems are then further partitioned into devices, modules, and channels according to individual detector architectures. Each component of this hierarchy is modelled as a finite state machine (FSM) with standardized states and recognized commands. The commands are propagated through the hierarchy from parent to child, while the states are reported by children back to parents. The global status of the DCS takes always into account the states of all children in the hierarchy, with exception of any which have been masked (excluded) by the experts. This approach has been implemented using the SMI++ toolkit, which proved to be an extremely powerful, flexible and reliable component.

The core of the ALICE DCS is based on a commercial SCADA system – PVSS II [3] – which is extended by frameworks developed at CERN. Individual PVSS systems are grouped by detector and supervised by central DCS system.

The whole system is configured using data stored in the central ORACLE database. Up to 6GB of configuration data is uploaded to detectors before a physics run can be started.

Acquired data is compared with predefined operational limits and automatic or operator driven actions are taken in case of an anomaly. A subset of the acquired data is stored in the ORACLE database for further use in offline analysis or by detector experts.

A large amount of data is being exchanged between the central DCS and external systems such as electricity, magnet control, gas systems, cooling, etc. Data exchange between ALICE and the LHC is provided within the DCS context and currently represents one of the major data processing challenges.

Finally, the DCS communicates with all components of the ALICE online and offline system. The information exchange is mostly based on the FSM states and commands, but a large amount of data also flows via dedicated data publishers and file exchange servers.

COHERENT AND HOMOGENOUS SYSTEM

The DCS covers a large number and variety of devices, systems and components, which are developed by various groups in parallel. The role of the central team is the coordination and monitoring of these developments. Standards and recommendations are issued and reviewed regularly. Whenever possible, common solutions are recommended and deployed both in the hardware and software domains [4].

In the original concept, the DCS backend, consisting currently of about 170 computers, was based on the Windows platform, with a very limited number of Linux installations. With the growing complexity of the overall DCS systems, the number of Linux services increased to about 30% of all installed systems. Another increase in Linux installations is related to the front-end boards, which contain embedded operating systems. In the full configuration there will be ~800 of such boards installed in ALICE, making Linux the main OS platform.

In the ALICE DCS the operating system flavours and software versions are defined centrally. The developers are then requested to follow these requirements. Due to continuous system evolution, the number of required software packages is increasing and there are several justified deviations from the standards. The central team carefully monitors the exceptions and updates standards as needed. The software upgrades are usually carried out during the longer LHC breaks in the winter period.

FLEXIBLE AND SCALABLE SYSTEM

The control system has been operational since the first device installations. It follows detector evolution with the goal to cover the lifetime of the experiment. The system architecture has to be flexible to accommodate future detector developments and operational procedures. At the same time, the system needs to be scalable to cope with the experiment's growth.

Keeping these requirements in mind, the DCS has been designed to accommodate possible future extensions. The first architectural principle was to build a modular system and avoid a monolithic architecture. Separate and independent control systems have been built for each detector, avoiding cross-detector dependencies. The individual detector systems consist of one or several sub-systems which control devices with similar functionality. All detector systems are then integrated into the global controls system, using the distribution features of PVSS.

On the lowest level, the device access has been strictly separated from the control tasks, which are implemented in PVSS. Modules communicating with devices recognize simple commands, but do not implement any control logic. On the supervision level, the operator is presented with a set of tools which allow for experiment operation. The tools are executed on computers separated from the machines in charge of the device control. The clear advantage of the modular approach is the possibility to distribute the individual parts across several computers. This allows for better resource sharing and assures scalability. Separation of operator tasks from the controls functionality helps to prevent a possible critical system overload, triggered, for example, by an excessive request – like conditions data retrieval for a long time period which could overload the operator machine, but will leave the controls functionality intact.

OPERATIONAL MODES AND CONCURRENT OPERATION

One main difference between DCS and other online systems is that the DCS needs to remain operational during all phases of the experiment. Full DCS functionality is expected also during the periods without data taking – shutdown periods, upgrades, etc. The requirements during the different periods vary; the DCS is designed to cope with them, successfully providing its services 365 days a year.

When experiment conditions allow for it, the detectors are able to get some autonomy in the operation of their hardware. Using the SMI tools, whole detectors, or their parts, can be excluded from the central DCS and released to local operators. Such systems do not receive commands from the central operator, but the alerts are still propagated and followed in a standard way.

The operational experience revealed one potentially weak point of this approach: the devices under the local control do not report their status to the central operator and ignore commands. This might be dangerous during the critical phases of the experiment – for example during the magnet ramp or injection of particles to the LHC. In such periods, the detector settings, such as high voltage or frontend configuration, must be compatible with the intended operation. For these reasons a new technique has been developed and deployed in parallel to the standard SMI tools. A set of software probes controls all critical settings and read back values, independent of the FSM status. Thanks to this, the safety of excluded devices can be assessed and taken into account regardless of their ownership within the hierarchy and assure the correct execution of DCS procedures.

The original hierarchical approach and the FSM standardization across all detectors as defined 10 years ago proved to be valid and successful. However, with the knowledge gained during the operation with LHC, many extensions were implemented and the whole system became too complex to be efficiently exploited by the shift crew. A set of tools were then developed to group several functionalities and to allow the operator to execute complex tasks without the in-depth knowledge of the underlying FSM processes.

USER FRIENDLY AND INTUITIVE OPERATION

During normal operation, only a small shift crew of 2-6 people control the whole experiment from the central workspace. Since the operators are not necessarily detector experts, special attention is given to the presentation of the system.

A dedicated component, ALICE DCS UI, has been developed and deployed in all detector systems. It allows for standardized presentation of the systems to operators; it provides a uniform look and feel and is easy to use. Experience has shown that the operators have a tendency to cover the monitor screens with many windows. The philosophy adopted in ALICE has therefore been to aggregate the information and to provide a summary to the operator who then can decide to browse the hierarchy for more details.

All critical operations which require operator action are executed from dedicated interfaces. These contain all the buttons and indicators gathered from other components which are required for the current task. The operator can therefore fully focus on the action, without the need for browsing the system to get information.

Although lots of attention has been given to intuitive operation, human errors were responsible for the majority of the operational incidents, which led to data taking delays. After careful analysis of all events, the tools were redesigned and simplified, with most of the actions automated. In 2010 we started to deploy expert systems which monitor the critical operations and inform the operator about all anomalies, with clear troubleshooting instructions. These instructions are displayed directly on the interface related to the executed task along with all indicators and buttons required for the problem resolution. Unlike the alert system, which typically informs about exceeded thresholds for a monitored value, the expert system proactively follows the operation and is able to detect events such as an incorrect sequence of commands that did not yet trigger an anomaly.

AVAILABLE, SAFE AND RELIABLE SYSTEM

Whereas the safety of the personnel is the task of the CERN Safety System, ensuring the integrity of the detector equipment is largely the task of the DCS. The control system allows for hardwired or software actions in case of hazardous situations. The system must be reliable and available; where needed, the equipment is running on safe power.

The current operational performance proves that the ALICE DCS reached this goal. While during the start-up phase of the experiment the shift crew required daily assistance of the central team, in 2011 we reached a stable state where the shift crew does not need to consult the expert for several weeks in a row. Besides the stability of the deployed software and hardware, a large contribution to the reliability and smooth operation comes from comprehensive training provided to the shifters.

SYSTEM MAINTAINABILITY

The complexity of the ALICE experiments clearly exceeds the capabilities of a small central team. The expertise related to individual detector operation is maintained in the various institutes that developed the detectors. By using well defined interfaces and standards it is possible to integrate individual developments into the overall DCS, however long term support and maintainability becomes a major concern.

The overall culture of the academic environment expects exploration and deployment of latest tools and technologies. A small central team is not able to certify all the proposed components and is therefore forced to insist on conservative solutions, providing the required functionality.

In the lifetime of a large experiment such as ALICE, the migration of experts is an inevitable fact which has to be taken into account. It is not uncommon that the major developments are carried out by graduate or PhD. students who leave and continue their career in other fields. The developed systems are transferred to new colleagues. The transfer procedure usually does not explain all the background and context, and the developers tend to focus on the solution without taking into account the overall system architecture. The role of the central team is to monitor this evolution and ensure that the takeover will respect the agreed rules and principles.

The key to a maintainable system lays in the standardization. The ALICE DCS project team imposed rules and common solutions whenever possible. From the very beginning, the detector requirements were carefully reviewed and standard solutions were proposed. Thanks to this approach, despite the large differences in detector architectures, we were able to limit the variety of devices used; only three brands of power supplies are used in ALICE, for example. The uniformity of the hardware has clear benefits for the support provided by the central team.

Operational experience gained in past years has also proven the success of the deployment of common solutions also in other subsystems, such as gas or cooling. One exception to this model is however the front-end and readout electronics (FERO). By the time of the DCS design, the individual FERO architectures were already significantly advanced, with most of the modules already produced. The DCS had to cope with a large variety of architectures based on completely different solutions. A perfect example is the deployment of field buses. While in the power system we were able to restrict this layer to CANbus and Ethernet, the FERO access is based on JTAG, RS-232, VME, CANbus, Ethernet and a number of non-standard solutions developed in the institutes. To cope with this diversity the concept of Front-end Device (FED) was developed [5].

The FED architecture is inspired by the commercial OPC technology. The low level layer of the FED software is responsible for the communication with the FERO. Its development requires deep expertise provided by detector designers, but once deployed it remains a relatively stable component, with changes linked mostly to hardware updates. The upper FED layer is based on the standard communication protocol DIM. Similar to OPC, it is implemented as a server that reacts to standardized commands sent by the clients and reports back the status. Several devices require a middle layer which translates the general commands into device specific actions. In most cases, the functionality of this layer focuses on sequencing and synchronizing commands transmitted directly to device channels, formatting and compression of data and low level error handling. The FED clients are implemented in PVSS. The FED concept provides a hardware abstraction layer, which allows for communication with the different FERO architectures in a unified way.

In the past few years, the FED concept has been used for other non-standard devices and subsystems in ALICE. Despite the obtained unified operation, FED remains one of the major challenges in ALICE DCS due to the complex functionality of its low-level layers.

OUTLOOK FOR SYSTEM EVOLUTION

During the long LHC technical stop planned for 2013, the DCS will profit from the gained experience and the system will be modernized.

As a first step, the standards are already being reviewed and new solutions will be proposed and made available to the detectors. All exceptions accumulated during the past years of operation will be removed and the system uniformity will be restored.

The plan is also to update all software components and to deploy the latest operating systems and tools. This is very delicate process as many commercial components do not follow the same evolution. Validation of the new systems has therefore already started in order to give sufficient time for finding satisfactory solutions.

The design of the detector hardware, including the computer interfaces, was launched more than 10 years ago. Stable operation of the experiment does not allow for regular upgrades. As a consequence, the DCS needs to operate a large fraction of obsolete hardware which in turn triggers a necessity to maintain a stock of spares. In some cases – like PCI interfaces, it is not possible to support the existing solution long-term. Therefore, new standards are being tested and prepared for deployment during the long technical stop.

CONCLUSIONS

The ALICE DCS project was launched 10 years ago. The construction, testing and operation phase of the system proved that the original assumptions and design principles were correct and efficient. The hierarchical approach and deployment of standards and common solutions contributed to the successful DCS operation.

The experience with the production system revealed many additional aspects that were not originally anticipated. One of the most visible examples is the front-end electronics with all its complexity and diversity. The DCS had to adopt already developed systems and build abstraction layers to allow for smooth integration and operation.

Interaction with systems external to DCS, increased data flow, and synchronization issues led to a need to create new unforeseen tools and procedures.

Comparing the original plans with the existing system it can be concluded that the system became more complex than anticipated, but the design principles have been well validated by operational experience. The main goal – to have a stable, reliable and safe system – has been reached. So, we got what we aimed for ten years ago.

REFERENCES

[1] Holme, O. et al., "The JCOP framework," proceedings ICALEPCS 2005, Geneva, Switzerland, 2005.

[2] ALICE Collaboration,Technical Design Report of the Trigger, Data Acquisition, High-Level Trigger and Control System, CERN/LHCC/2003-062.

[3] ETM website http://www.etm.at/index_e.asp (the product was recently rebranded as "Simatic WinCC Open Architecture").

[4] Augustinus A. et al., "The wonderland of operating ALICE experiment", this Proceedings.

[5] Chochula P. et al., "Control and Monitoring of the Frontend Electronics in ALICE", Proceedings 9th Workshop on Electronics for LHC Experiments, LECC 2003, Amsterdam, The Netherlands.