THE EVOLUTION OF THE ALICE DETECTOR CONTROL SYSTEM

P. Chochula[1], A. Augustinus[1], P. M. Bond[1], A. Kurepin[1,2], M.Lechman[1,3], O. Pinazza[1,4]
[1]CERN, Geneva, Switzerland
[2]INR RAS, Moscow, Russia
[3]Institute of Physics, Slovak Academy of Sciences, Bratislava, Slovakia
[4]INFN Bologna, Italy

Abstract

The ALICE Detector Control System has provided its service since 2007. Its operation in the past years proved that the initial design of the system fulfilled all expectations and allowed the evolution of the detectors and operational requirements to follow. In order to minimize the impact of the human factor, many procedures have been optimized and new tools have been introduced in order to allow the operator to supervise about 1 000 000 parameters from a single console. In parallel with the preparation for new runs after the LHC shutdown a prototyping for system extensions which shall be ready in 2018 has started. New detectors will require new approaches to their control and configuration. The conditions data, currently collected after each run, will be provided continuously to a farm containing 100 000 CPU cores and tens of PB of storage. In this paper the DCS design, deployed technologies, and experience gained during the 7 years of operation will be described and the initial assumptions with the current setup will be compared. The current status of the developments for the upgraded system, which will be put into operation in less than 3 years from now, will also be described.

THE ALICE DCS DESIGN

The central ALICE detector consists of 19 subdetectors, built with different detection technologies and with largely different operational requirements, all supervised by a single operator. The architecture of ALICE Detector Control System (DCS) is based on standards adopted by the LHC experiments at CERN. The commercial SCADA system WINCC OA, extended by CERN JCOP and ALICE software frameworks, is configured as a large distributed system running on about 100 servers [1]. To guarantee the autonomous operation of each subdetector, the central distributed system is segmented into subdetector systems, each allowing for stable subdetector operation in isolation from the other subdetectors.

Wherever possible, the ALICE DCS is based on commercial hardware and software standards. OPC servers handle the communication between the WINCC OA and the controls devices. Nonstandard devices, like detector frontend electronics modules, are interfaced with WINCC OA using ALICE FED standard, a complex client-server mechanism based on the DIM communication protocol [2].

All controls tasks of the ALICE experiment could be fully satisfied with the pure WINCC OA system, however its operation would require a deep knowledge of many technical details. Using the SMI++ framework [3], installed on each WINCC OA system, the operation of controlled components is modelled as a finite state machine with well-defined behaviour. Figure 1 shows the schematics of main DCS components.

The various dependencies, like the need to configure a channel before it can be turned on, are encoded in the SMI++ logic. Using the mechanisms provided in the SMI++ framework,
ALICE is represented as a hierarchical tree, with the central DCS placed on top of the pyramid and individual channels on its bottom. Commands sent from the top objects are propagated to the child nodes: the subdetectors, subsystems (high voltage, cooling, frontend electronics), devices and their channels. Each object reports its current state to its parent object. Any part of the tree can be excluded from the hierarchy and operate independently.

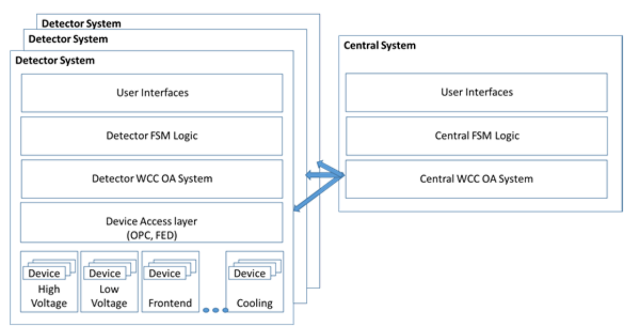

Figure 1: Architecture of the ALICE DCS.

OPERATIONAL EXPERIENCE

During eight years of ALICE operation, the initial architecture and system implementation proved their reliability and robustness. The DCS provided stable 24/7 services with small interruptions required mainly for the infrastructure modifications (cooling and ventilation upgrades, reorganization of the computer racks…). Even during the service breaks, the core systems related to safety remain operational.

With additions of new detector modules, the DCS is being continuously extended, profiting from scalability of the core architecture. In few cases new computers needed

ISBN 978-3-95450-148-9

to be added, to balance the load on busy subsystems. So far all detector requirements could be satisfied with the existing systems and tools, without the need for major changes in the original architecture. All system upgrades and patches followed the standard evolution of the hardware, software and operating systems.

Main modifications of the DCS are related to the experiment performance requirement. Procedures and operation tools are being continuously adjusted in order to implement data taking requirements and reflect operational experience. The original approach expected that the operator would interact directly with the state machines and move detectors to desired states as a response to operational conditions. This method, successfully implemented in the previous generations of high-energy physics experiments, soon reached its limits, due to the complexity of the controlled equipment and cross-dependencies between the subdetectors, infrastructure and external conditions.

For example, as a reaction to beam mode changes, the detectors need to adjust their settings. To protect gas detectors from damage, the high-voltage channels need to be ramped down to an intermediate state, assuring that an accidentally deposited charge will not damage the readout channels. Simple FSM mechanism does not assure, that all sensitive devices receive the command. Certain modules might be for example excluded from the central hierarchy for expert intervention and therefore ignore the commands sent by the central operator. A mechanism independent on FSM had to be implemented. Dedicated safety scripts regularly verify the status of physical hardware channels and calculate the safety condition of each detector based on readout values. Using this low-level bypass mechanism, the operator receives an overview of the experiment status based on the physical values instead of the logical state calculated by the FSM. The high-level tools allows to force the safety related commands even to modules, which are not under direct controls of the central operator.

The routine operations depend on a number of external conditions, which need to be verified before any command is sent to the subdetectors. For example a period of stable beam collisions is preceded by a complex procedure during which the particle bunches are injected into the LHC, accelerated, and adjusted. Each phase of this procedure represents a different set of risks for the detectors. While certain detectors can be set to nominal operational values during the whole procedure, some of them must remain at reduced voltage settings until the injection is completed. Operation of several subdetectors has to be postponed until the end of the particle acceleration phase and finally the operation of remaining detectors can be restored only after a phase of stable beam collisions has been reached. The status of LHC is not the only factor defining the actions to be executed by the operator. For example high radiation levels prevent the detector from nominal operations even if stable beam collisions were established.

All the operational rules are implemented in the high-level operational procedures. Instead of controlling the subdetectors directly, the operator only issues a command to reach a desired configuration. The role of the high level procedures is then to execute this request by controlling the FSM directly and taking all cross-dependencies and external conditions into account. Decoupling the operator from the low-level controls tasks significantly reduced the number of human errors as well as the time required for an execution of most complex actions.

The DCS operator has direct access to about 1 million of controlled and monitored parameters. The visualization and access to these values is established through graphical user panels, which are organized hierarchically on a single operator console. The user interface allows for easy navigation to any of the panels, using only single user interface.

The ALICE DCS is built in collaboration between the central coordination team and detector experts working at institutes remote to CERN. The information provided on the panels is largely biased by the deep expertise of the developers and is not necessary intuitive to operators, who have typically no prior experience with control systems. The central team has analysed the individual subdetector requirements and provided high-level tools and interfaces, which are presented to the operator. The tools invoke detector actions as needed.

The central team has issued a set of implementation rules and guidelines based on the experience with the previous generations of the control systems and supervised the developments carried on by detector groups. With growing experience as well as with the increasing complexity of the experiment requirements, the guidelines are being continuously refined. Big efforts are invested into implementing the new standards with a focus on uniformity across the whole DCS. For example, all popup messages aimed to bring the focus of the operator to a certain local problem were replaced with standardized alerts, displayed on a single screen. Each alert is accompanied by a set of instructions, which will guide the operator in the troubleshooting process.

The central DCS operator plays a key role in the operation. Originally, subdetector experts operated the systems and the central operator coordinated their actions. Currently the DCS operator is typically the only person in the ALICE run control centre linked to the controls tasks. To assure continuous 24/7 operation of the experiment, more than 150 operators were trained in 2015. The training procedure has evolved from a simple introductory presentation accompanied by hands-on experience to a formal process. It includes the lecture, hands on session using a simulator, compulsory training shifts during which the trainee operates the experiment under the supervision of an experienced shifter and then a final exam.

ISBN 978-3-95450-148-9

THE ALICE O2 PROJECT

A major experiment update is foreseen for the third phase of the LHC operation, starting in 2019. New detectors will be installed in ALICE and readout of detectors will be modernized. The present data acquisition mechanism based on triggered readout will be replaced by continuous data taking. The detectors will be connected through Alice readout links to a farm of 250 servers, the First Line Processors (FLP), which will assemble the detector information into consistent data sets. The acquired data will flow to 1500 Event Processing Nodes (EPN), installed in ALICE, performing the full data processing. This new approach merges the online and offline roles into one system, named O2.

The farms need to digest 1.09 TB/s of detector data, producing about 40 PB of data for physics analysis each year.

The new detectors will be integrated into the existing DCS using the well-established standards and procedures. Thanks to the scalability of the DCS architecture, this task does not present a major challenge. There are, however, two areas requiring new approaches: the conditions data flow and frontend electronics control.

The Conditions Data Flow in O2

The present DCS archives all acquired values in a ORACLE database, independently from data acquisition. After each period of data taking, called run, the DCS collects all archived conditions parameters and sends them to offline processing.

Depending on the duration of a run, the delay between the actual readout of a parameter and its transfer to offline can be in the order of several hours. This situation dramatically changes with introduction of the O2. The data is sent to FLPs in 20 ms time frames. Each time frame must be accompanied by a full set of condition parameters. The current estimates suggest, that about 100 000 parameters will need to be inserted to each O2 data frame. This requirement is of course out of scope of the DCS, where parameter are updated at a rate typically lower than 1 Hz. The present mechanism expecting the data first to be archived in ORACLE would add another delay in the processing. A new mechanism, named ADAPOS (Alice DAtapoint Server) is being developed.

A specialized software module, the Data Finder and Publisher (DFP), will equip all WINCC OA systems. Its role is to collect all conditions parameters provided by its host system and publish them to subscribers.

The data collector is a client part of ADAPOS. It connects to all DFPs and reads published values. Alternatively, the data collector can retrieve also data from the DCS ORACLE database as shown in Figure 2. The received parameters are passed to the module, which is maintaining a control process image. This memory resident object keeps information about all monitored parameters. The data collectors refreshes the process image each time when a new value arrives. The whole process image is mirrored on a dedicated DCS FLP. Synchronization between the ADAPOS and FLP is handled using a dedicated data transfer channel. The FLP will then retransmit the process image to the EPN, synchronized with the 20 ms data frames.

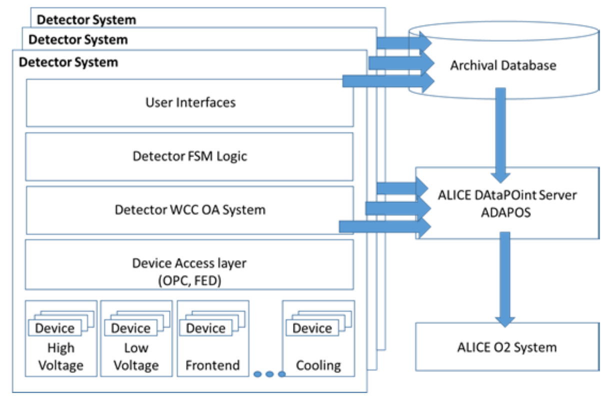

Figure 2: Architecture of the DCS-O2 interface.

The Frontend Control

Part of the DCS information is produced by the detector frontend modules. The DCS conditions data is inserted into dedicated packets, transmitted along with detector data to FLPs. The firmware of the receiver cards strips the DCS information off the data stream and publishes it to the DCS clients implemented in WINCC OA. Each client subscribes to a required subset of published values without the need to know the details on physical configuration of the frontends and FLPs. A common name service will handle the redirections of subscription requests. One of the already existing technologies supporting this mode of operation is DIM. First prototypes based on DIM are able to carry this task with sufficient contingency margins.

The configuration sent from DCS to the frontend modules is following a similar path – the WINCC OA based client sends a command to a server implemented on the FLP, which will insert it to a data frames sent to the frontend modules. This task requires synchronisation at many levels – the command may not interfere with regular data taking, conflicts sent by concurrent commands shall be resolved before the electronics is reconfigured, acknowledge signals must be sent after the commands are executed, etc. To handle this complexity, a concept of ALICE FED [2], known from the current operation, is being extended. Each detector will be equipped with one or more FRontEnd Device servers (FRED), which will listen to commands sent by WINCC systems and transmit them to the FLPs with the physical access to target devices. The responses will be then transmitted back to requestors either using the standard conditions data channels, or via the FRED servers.

ISBN 978-3-95450-148-9

CONCLUSIONS

The ALICE DCS followed the evolution of the experiment and provided a stable service over a period of more than eight years. Major efforts were put into automation and unification of the operational procedures. The system extensions followed the evolution of the ALICE detectors and profited largely from the scalability of the ALICE DCS architecture.

ALICE upgrade for the LHC RUN3 period puts new challenges on the DCS. The access to the frontend modules has to be redesigned and existing standards will be modified to cope with the new electronics. Major architectural change is required for the conditions data flow. Current batch processing of the conditions data, based on the values archived in the ORACLE database, will be replaced with a publishing mechanism, allowing the transmission of measured valued to the O2 facility in quasi-real time.

REFERENCES

[1] P.Chochula et al., Operational experience with the ALICE Detector Control System, Proceedings ICALEPCS 2013, San Francisco (2013)

[2] P.Chochula, L.Jirdén, A.Augustinus, Control and monitoring of the front-end electronics in ALICE, 9th Workshop on Electronics for LHC Experiments, Amsterdam 2003

[3] Franek, B, Gaspar, C: “SMI++ object oriented framework used for automation and error recovery in the LHC experiments”, J. Phys.: Conf. Ser. 219 (2010) 022031

ISBN 978-3-95450-148-9