

REALIZATION OF A CONCEPT FOR SCHEDULING PARALLEL BEAMS IN THE SETTINGS MANAGEMENT SYSTEM FOR FAIR

H. Hüther, J. Fitzek, R. Müller, A. Schaller, GSI, Darmstadt, Germany

Abstract

Approaching the commissioning of CRYRING, the first accelerator to be operated using the new control system for FAIR (Facility for Antiproton and Ion Research), the new settings management system will also be deployed in a production environment for the first time.

A major development effort is ongoing to realize requirements necessary to support accelerator operations at FAIR. The focus is on the pattern concept which allows controlling the whole facility with its different parallel beams in an integrative way. Being able to utilize central parts of the new control system already at CRYRING, before the first FAIR accelerators are commissioned, facilitates an early proof of concept and testing possibilities.

Concurrently, refactorings and enhancements of the commonly used LSA (LHC Software Architecture) framework take place. At CERN, the interface to devices has been redesigned to enhance maintainability and diagnostics capabilities. At GSI, support for polynomials as a native datatype has been implemented, which will be used to represent accelerator settings as well as calibration curves.

Besides functional improvements, quality assurance measures are being taken to increase code quality in prospect of productive use.

COMMISSIONING OF CRYRING AT GSI

At the time of writing, the CRYRING heavy-ion storage ring, a Swedish in-kind contribution to the FAIR project, has been set up at GSI, with only few additional components still to be installed. The machine features an electron cooler, an RFQ linear accelerator and two injectors for different types of ions. See Fig. 1 for an overview of the CRYRING ring section.




Figure 1: Overview of the CRYRING storage ring as setup at GSI (injection lines not shown), W. Geithner, GSI.

From a controls perspective, its commissioning represents a major milestone, as CRYRING will be the first machine

solely operated via the new FAIR control system. Currently, the commissioning group still relies heavily on low-level control, but by the end of the year, they will have shifted their primary work place from the accelerator tunnel to the control room, utilizing the whole control system stack through the high-level applications provided.

TEST-BED FOR THE NEW FAIR CONTROL SYSTEM

Since CRYRING is equipped with its own injector line, it can continue operation even at times when the existing GSI accelerator chain is shut down for necessary FAIR upgrade and civil construction work, making it an ideal test-bed for the new control system. This way, it will contribute to validating concepts and technologies under real-world conditions [1], ensuring that the control system components work properly, individually and as a whole, and that business processes within and between involved parties are effective.

Although operating CRYRING does not imply the same requirements on the settings management system as the future FAIR facility will, core concepts necessary for highly flexible future operation scenarios can nevertheless be tested. As such, the beam-oriented approach to scheduling, designed for parallel beam operation at FAIR, will be utilized at CRYRING for the first time.

PARALLEL BEAM SCHEDULING CONCEPTS FOR FAIR

The designated operation modes of FAIR put demanding requirements on the new control system currently in development. To optimize the number of concurrent research programs, the facility will provide up to five beams in parallel with pulse-to-pulse switching between different particle types. Additionally, great flexibility shall be provided, allowing to change the parallel operation schemes on a daily basis.

Beam production chains and patterns are the central technical concepts within the new LSA-based settings management system to fulfill these requirements. Representing a major change in perspective, beam production chains establish a beam-oriented view on the facility, as compared to the accelerator-oriented view towards settings management dominant at GSI up to this point.

Beam production chains are defined using beam processes as atomic building blocks. Beam processes represent a specific procedure on the beam within one accelerator (e.g. injection, ramp, extraction). Within a beam production chain, the order of all beam processes necessary to provide the settings for producing a certain beam is described, from its

source up to its target, spanning all machines and transfer lines involved.

To be able to coordinate multiple beams traversing the facility in parallel, beam production chains are grouped into patterns. An example of typical parallel operation for the modularized start version of FAIR is given in Fig. 2.




Figure 2: Example for parallel beam operation showing scheduling of beam production chains into patterns. HESR accumulating the anti-protons is omitted.

STATUS OF THE FAIR SETTINGS MANAGEMENT SYSTEM

LSA at GSI has made significant progress since it was first deployed as a test instance in 2008 [2]. In close cooperation with CERN, a structural redesign aimed at modularization led to more efficient handling of change management and release processes while also fostering extensibility [3]. First successful tests with beam at the heavy-ion synchrotron SIS18 were performed in 2010 [4]. More advanced machine experiments were made possible in the subsequent years through implementation of GSI-specific features, like flexible beam process lengths and optics definitions which are relative regarding time [5].

Development is currently focused on CRYRING and implementing the pattern concept. Although still at a prototype state and major features for controlling the entire future FAIR facility to be implemented, LSA now supports all requirements initially needed for CRYRING commissioning. Machine physicists were able to calculate reference settings using the pattern scheduling mechanisms and interface tests using mock-up front-end device controllers are being conducted. Testing data supply with actual devices is expected to start in November.

Two of the most recent enhancements leading to this achievement shall be described in more detail here.

RECENT ENHANCEMENTS OF THE FRAMEWORK

Polynomial Data Types

Developers at GSI implemented a polynomial data type directly in the central value package of the LSA framework. Before polynomials were available, all functions data had to be provided as discrete functions.

There are a lot of benefits to using polynomial data types instead of discrete functions. Polynomials produce continuous values rather than x-y-pairs for potentially millions of steps. Hence less data has to be stored because polynomials are represented as an array of its coefficients only. All polynomials can have an interval describing for which range of x-values they are defined, so it is possible to build sequences of polynomials over different intervals to easily represent complex functions. Consequently, it is possible to provide a more accurate representation of a function than before.

Combining a polynomial with an interval is inevitable for settings management, but it also brings up the problem that it can be interpreted in two different ways: A bounded polynomial can be treated absolute or relative to its interval. Using the absolute interpretation, the polynomial $p[n]$ is evaluated as $p[n].interpolate(x)$ independent of the boundaries, whereas on the relative interpretation the polynomial $p[n]$ is evaluated as $p[n].interpolate(x - b)$, where b is the lower bound. Therefore all polynomials are treated as relative in the current implementation.

The magnet group at GSI provides calibration curves for the magnets in the central component database also as polynomials. Before the polynomials data type was introduced, these polynomials had to be rasterized into discrete functions before they could be used in LSA. Now, calibration curves can be imported in their native, unmodified format as received from the magnet group and handled as polynomials internally as well. The same applies to function settings calculated by LSA, which are sent to the hardware, e.g. the function generator for ramped devices [6], as polynomials.

Data Supply

The LSA subsystem responsible for data supply to devices received a major overhaul in 2014. The project was carried out as a joint effort between CERN and GSI, with one developer from Darmstadt staying in Geneva for several months. The core motivation, equally important for both parties, was to improve diagnostics capabilities of the system.

While data supply results were formerly presented in a flat data structure for the whole supply process, the new hierarchic result classes allow operators to trace errors down to individual parameters, e.g. in case a set value is rejected by a front-end device controller. Navigating the hierarchy level of data supply results is possible in a very convenient way, starting at the beam process that was sent, via device adapters handling value conversion, down to set calls to the middleware, showing exactly where the error occurred. These levels of aggregation can also be identified in the corresponding class diagram shown in Fig. 3.

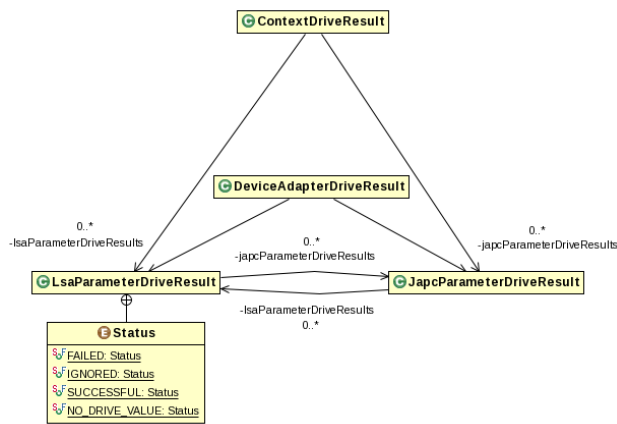


Figure 3: Class diagram showing the refactored LSA data supply results structure.

The enhanced diagnostics capabilities have already proved beneficial when supplying settings to large numbers of devices. As the need to handle error states can hardly be avoided during commissioning, they will also be a valuable tool for CRYRING.

Furthermore, the subsystem's code, which has been evolving for nearly a decade, received a major refactoring. Coding styles were unified, classes restructured and general maintainability and extensibility enhanced.

QUALITY ASSURANCE MEASURES

Both the implementation of the pattern concept as well as the aforementioned refactoring, being mission-critical core components of LSA, were carried out using a test-driven development approach to facilitate focus, correctness and test coverage. Beyond that, aiming to minimize potential errors within and between LSA components during commissioning of CRYRING, a variety of additional quality assurance measures have been implemented.

The machine model code, written by physicists, is being reviewed both technically by the LSA software development team and content-wise by other accelerator physics experts. During these sessions, to complement the traditional inspection approach, automatically generated reports utilizing a variety of static code analysis tools are consulted to support framework and machine model developers in writing correct and comprehensible code.

Besides methodology and code quality, the third pillar of quality assurance measures currently employed is integration testing. For CRYRING, there are roughly around 15 FESA (Front-End Software Architecture) classes to be run on front-end device controllers, either already implemented or currently under development. Each will be handling at least one, but usually several devices of the same type. To ensure that front-end device controller software works as expected from an LSA perspective, a test suite to perform automated integration testing has been set up. The tests typically operate on mock-up instances of FESA classes, which

mimic the behavior of the actual hardware they are designed to control. At a later stage, tests may also be performed on front-end device controllers connected to real machine components. To maximize efficiency, the underlying testing concept distinguishes and reuses sets of generic tests for all device classes, device-type-specific tests and device-specific tests.

OUTLOOK

As previously mentioned, there are still major features to be implemented in LSA in order to fully support flexible parallel beam operation envisioned for the future FAIR facility. Taking SIS18 booster mode as an example, scheduling a certain beam production chain multiple times as a sub-chain will be necessary to consistently model the successive injections into SIS100. Another challenging task will be to support multiple patterns scheduled at the same time, with one main pattern being executed repeatedly and one or more other patterns serving experiments that require beam on demand only.

In 2016, tests are planned for transferring beams from the storage ring ESR, which is part of the existing GSI facility, to CRYRING. This endeavor requires synchronization between the previous control system still being used at ESR and the new control system used to operate CRYRING. Once LSA has been equipped with additional features specifically targeted at storage ring operation at GSI, most importantly beam manipulation during the cycle for experimentation phases, this will open interesting possibilities e.g. for working with rare isotopes at very low energies at CRYRING [7].

REFERENCES

- [1] R. Bär et al., "News from the FAIR control system under development", PCaPAC'14, Karlsruhe, Germany, WPO004.
- [2] R. Müller, J. Fitzek, D. Ondreka, "Evaluating the LHC Software Architecture for Data Supply and Setting Management within the FAIR Control System", ICALEPCS'09, Kobe, Japan, THP012.
- [3] R. Müller et al., "Benefits, Drawbacks and Challenges during a collaborative Development of a Settings Management System for CERN and GSI", PCaPAC'14, Karlsruhe, Germany, TCO101.
- [4] J. Fitzek, R. Müller, D. Ondreka, "Settings Management within the FAIR Control System based on the CERN LSA Framework", PCaPAC'10, Saskatoon, Saskatchewan, Canada, WEPL008.
- [5] H. Hüther et al., "Progress and Challenges during the Development of the Settings Management System for FAIR", PCaPAC'14, Karlsruhe, Germany, WPO005.
- [6] S. Rauch, M. Thieme, "Managing multiple Function Generators for FAIR", PCaPAC'14, Karlsruhe, Germany, FPO017.
- [7] F. Herfurth et al. "The Low Energy Storage Ring CRYRING@ESR", COOL'13, Murren, Switzerland, THPM1HA01.