

FAIR Commissioning & Control Working Group

Notes from the meeting held on 3rd June 2015

e-mail distribution: [FAIR-C2WG-ALL at GSI.de](mailto:FAIR-C2WG-ALL@GSI.de), [participants list](#)

Agenda:

- FAIR Accelerator Operation Paradigms – Personnel Perspective ([jump below](#)), S. Reimann
- Data Archiving & Post-Mortem – First Iteration ([jump below](#)), R.J. Steinhagen

1. FAIR Accelerator Operation Paradigms – Personnel Perspective, S. Reimann

S. Reimann provided a summary of the past operation of the GSI accelerator facility and its implications on the operation of FAIR from a personnel point of view (see [slides](#) for details).

Until 2012 (2014), a dedicated pool of 23 operators (part of GSI's LOAO group) handled the 24h/7 operation of GSI's accelerator facility, with three operators covering one shift that primarily handed UNILAC and SIS18. Operation of ESR is largely driven and operated by its experimental users and only to a limited extent by GSI's regular operation group. In addition, one operation coordinator, four machine experts (UNILAC, SIS18, ESR & FRS) and 28 equipment experts (controls, RF, magnets, vacuum, ...) are 'on call' in case of problems or major mode of operation changes.

The total number of required operators per one operator 'on shift' ("shift working places" or shift slots) is governed by GSI's operation and collective bargaining agreements ("Betriebsvereinbarung" and "TvöD", see [slides](#) for details). Including realistic availabilities, these imply that a standard early-late-night shift pattern needs to be serviced with a minimum of 7.4 persons per operator on shift (e.g. 3 operators on shift x 7.4 persons/'shift slot' ≥ 22.2 FTE \rightarrow 23 FTE total in pool). In 2015, a legal minimum of two operators on shift will provide limited operation of UNILAC, which corresponds to a required total pool of at least 15 operators. SIS18 re-commissioning in 2017 and SIS100 commissioning are expected to require at least four operators on shift, corresponding to a pool of 30 operators, or a pool of 37 if specialised operators for SIS100's and Super-FRS's cryo infrastructure are included.

S. Reimann distinguished the following phases of operation:

1. Regular Operation: covering the (initial/re-) commissioning, set-up of new beams, and ongoing machine developments (typically during daytime) that are typically performed by a mix of operators, machine experts, on-call experts, and experiments tightly linked with accelerators (total of about 24 persons in the control room).
2. Monitoring & Adjustment: performed by a skeleton crew of 5-6 operators typically during nights and weekends. This mode assumes a high level of automation of accelerator operation and that the machine set-up has been already performed.

And mixtures between the above in case of problems.

The various possible operation paradigms were discussed in order to address the future challenges of FAIR operation at hand: massive parallel operation, new operator skill requirements, limited operator resources. The extremities are covered by:

- One operator per machine (OPM): this scheme focuses on optimising the accelerator individually and is similar to present operation at GSI. The advantages are better skilled operators, causing less operational errors and faster beam set-up for a specific accelerator (only). The disadvantages cover [but are not limited to] reduced interface efficiencies of transferring beams across accelerators domains, limited possibility of setting up multiple experiments in parallel, and limited flexibility of shift planning (an operator can only be replaced by another with the same expertise). This scheme requires a much larger pool of operators (59 compared to presently 23 persons, excluding cryo-operators), increasing the annual operation costs of FAIR, and cannot avoid potentially idle resources when not all accelerators or experiments are being operated.
- One operator per experiment (OPE): this scheme focuses on the optimisation of the beam production chain across accelerators to the experiments and is the proposed control and operation strategy for FAIR. The advantages of this scheme are a more efficient set-up and interface across accelerators and to the experiments, reduced number of required personnel (30-37 compared to presently 23 persons, excluding cryo-operators), the operator being an expert and more highly motivated to deliver the required beam parameter (“my experiment”), and more redundancy thus flexibility with respect to shift planning. Some of the disadvantages to be addressed are better and continuing training requirements for operators, requirement of more common tools and automation of standard processes across accelerators, and adapted console scheme.

S. Reimann stressed that hybrid options between these two extremities are possible: e.g. that the more experienced shift-leaders/operators that can cover a broader range of accelerator domains are paired with operators that are machine-type specialists (e.g. LINACS, ring accelerators). While this hybrid (close to the one-operator-per-experiment paradigm) is favoured, he pointed out that the following aspects would need to be addressed: larger control room, specialized operator training, stable long-term beam time scheduling, digital control room, generic operation software, consistent look and feel high-level automation, restricted parallel access on machine settings. If we cannot be assured on these points, a substantial increase of the required personnel for standard operation is needed.

S. Reimann further proposed some additional measures to aid future operation:

- add option for shift work to new contracts (operators on demand)
- add software skills as a requirement for operators (creating automation procedures)
- new operators should also have some theoretical background in physics
- development of a dedicated operator training concept
 - periodic lectures (similar to [AXEL@CERN](#))

- hands-on training (2016 CRYRING, from 2017 on dedicated training shifts)
- operator exchange with other facilities

Discussion:

Actions:

- **FC²WG-all:** discuss and evaluate the presented proposal for its viability
- **M. Winkler (MPL), S. Pietri:** discuss which part of Super-FRS operation should, could, needs to be done by OP → by August (S. Pietri dicit)
- **L. Groening (MC), P. Gerhard:** discuss if UNILAC operation should be separated from the rest of FAIR-operations and follow a distinct different paradigm
- **R. Steinhagen (MPL), S. Reimann:**
 - organise a FC²WG follow-up to reiterate on the presented operational paradigm (target: FC2WG meeting early autumn)
 - formalise and submit the FC²WG proposal to the GSI/FAIR management

2. Data Archiving & Post-Mortem – First Iteration, R.J. Steinhagen

R. Steinhagen discussed (see [slides](#)) the various use-cases and requirements and the planned *Archiving* and *Post-Mortem Systems* for FAIR. The aim of these is the central collection and storage of pertinent accelerator data to permit a quantitative analysis of the accelerator performance and its proper function. He further distinguishes:

- 'Data Archiving' as a continuous, periodic and slow data acquisition (aka. 'Logging') that is used to reconstruct and re-analyse machine or beam condition, from
- 'Post-Mortem' as a fast, transient data recording, that is executed asynchronously in response to a 'post-mortem' event typically containing high frequency data, used to reveal the causes of emergency beam aborts, possible equipment damage and to validate the correct functioning of the machine protection systems.

Some of the main use-case examples include the quantitative comparison of present and past machine performances and stability in order to guide the ongoing optimisation of the machine from time scales of cycle-to-cycle up to covering several month to years of operation. An important application of gathering and storing, for example, beam data are cycle-to-cycle feedbacks that use the recorded information to automatically re-steer the machine for recurring systematic drifts based on these data. Similar schemes used at CERN have been shown. Similar semi-automated cycle-to-cycle feedback systems will be also deployed for FAIR that control (in order of priority): orbit (SIS18 recommissioning in 2017), trajectory (injection/extraction steering), tune and chromaticity.

Another important application of archiving and post-mortem is the quantitative analysis of 'near misses' or rare events that could be used to mitigate these events and/or to propose new counter measures intercepting these at an earlier (/safer) stage based on the observed actual frequency of the individual events. Some examples included but are not limited to: analysis of near-misses not caught by post-mortem or the machine protection systems, but that are visible (e.g. as slow drifts) in the archiving system; or establishing reliable reference and tolerance thresholds that optimise machine availability based on past operational experiences (see [slides](#) for details).

R. Steinhagen emphasised that the specific data sources that are required or useful for a later analysis are not always known in advance. Signals considered initially as 'noise' may contain pertinent information after detailed analysis ("somebody's 'noise' is somebody else's 'signal'" principle). This particularly applies for machine development studies during which, all relevant data may not necessarily have been recorded and that often reveal unexpected effects during the post-analysis of the studies [N.B one of the main reasons why MDs are performed]. In addition, the analysis of rare events can be particularly difficult, due to not being easily reproduced or because of a potential history of 'near misses' that might not have been detected by the post-mortem trigger infrastructure. A repetition of these types of study for the sake of uncovering additionally required information would waste beam time and create unnecessary costs.

It was proposed to aim to monitor and archive a wide range of accelerator-related signals continuously for FAIR to open the possibility of providing evidence for difficult operational, optimisation and machine protection scenarios. The non-exhaustive list of devices covers beam instrumentation, feedbacks, beam absorbers and collimators, cryogenics, kickers, power converters and quench protection systems, radio-frequency cavities, radiation monitors, timing system and interlocks, vacuum, access system and other general FAIR infrastructure.

R. Steinhagen provided a first estimate for the data bandwidth and total amount of data required for a full coverage of the FAIR accelerator facility: the required sampling rates are dependent on the expected normal accelerator time constants (control theory: sampling typically 20 to 40 times larger than the natural bandwidth), e.g. 1 kHz for SIS18, 100 Hz for SIS100, and 10 Hz for CR and HESR. The required sampling for HEBT and Super-FRS needs to be further evaluated since the time constants vary largely depend on whether the beam is extracted in a single shot (e.g. sampling ~0.1 ...1 Hz) or whether a slow resonant extraction is used (e.g. sampling 10 .. 1k Hz). The resulting corresponding total data bandwidth has been estimated to be about 10 MByte/s and data volume of about 5 TByte/week (see [slides](#) no 16 for details).

These data rates would need to be reduced by a factor of 25 to 50 to achieve a viable long-term storage in the order of 5 to 10 TByte/year. It is thus proposed to keep the high-frequency data streams in the short-term (couple of weeks up to a (few) month(s)), and to continuously reduce and store a reduced low-bandwidth copy of the data in the long-term (time scale: several years). The required data reduction could be implemented through a variety and combination of options:

- storing some variables only in the short-term storage which is acceptable for some (i.e. 'expert-only') variables, as they lose importance after a few weeks (N.B. individual definition of short- vs. long-term storage can be changed at any time if deemed necessary)
- dropping BI data for periods without beam (e.g. during shut-down, cycles without beam). This would require a reliable beam current information.
- 'Machine-Mode' or 'Beam mode' dependent data rates (initial proposal for modes to be discussed in the next FC²WG meeting).
- Beam-process dependent data rates: no-beam, injection, ramp, flat-top/(fast) extraction, slow-extraction, coasting, ramp-down
- broad rule of keeping 1 kHz data for a few weeks and then to indiscriminately down-sample it to 1 Hz. Option 1: keeping high-frequency data on request for machine development purposes. Option 2: beam mode dependence: 'high' (1 kHz) during set-up/intensity ramp up, 'low' (10/100 Hz) during production runs
- On-change data reduction: this concept is easy for boolean or integer type values (e.g. 'on','off', statuses etc.), but non-trivial for floating point values. A proposed possible compromise would be to:
 - compute and store the minimum, mean, standard deviation, median and maximum trace of a given variable for a given cycle and to keep a full snap-shot every e.g. one hour
 - store full data trace for transients, e.g. if the actual value exceeds more than a predefined number of standard deviations from the mean trace (see [slides](#) no 19 for a graphical visualisation).

As notable exceptions: beam profile data at injection or at extraction should be excluded from the data reduction. The archiving should keep the full data rate for the last minute before a post-mortem event in response to a beam-abort or quench. The latter would be used to detect and analyse drifts or precursors just prior to the post-mortem event.

While the required expected write-bandwidth is only about 10 MByte/s, many more users are expected and thus sufficient read-bandwidth margins need to be foreseen to support multiple accesses to extract, use and analyse the stored data. The total number of archiving clients needs to be established but probably ranges from 10 to 100. As an option, the available read-bandwidth could be reduced gradually as a function of time (e.g. high-bandwidth for a couple of month, low-bandwidth after a year).

The archiving system should provide additional pre-filter options in order to optimise the amount of data to be extracted for an analysis. This includes, for example,

- time selections covering multiple windows, intervals, or cycles, as well as
- more refined options that limit the data based on whether
 - there was actually beam in the machine during a given cycle or beam processes
 - retrieve data between two pre-defined recurring events
 - cycle or beam production chain selection for a given beam type (e.g. U28+-only, U92+-for-CBM, 'protons for p-production', ..), and
- other meta-information (filter conditions)
 - Beam parameter: ion-species, beam target (experiment), actual intensity, actual beam transmission, targeted intensity range, ...
 - Machine parameter: ramp-rate, min./max. rigidity, injection/extraction energy, slow/fast extraction, cycle/store length
 - Cycles with(-out) post-mortem events
 - Operational ranges: OP year, between technical stops, OP week, OP day, etc.

Next Steps and Actions:

- **FC²WG-all:**
 - discuss and evaluate are there other important use-cases
 - meta-Information (filter conditions): more input/confirmation from potential users is needed (please contact: R. Steinhagen, L. Hechler or J. Fitzek)
 - Is the proposed data reduction paradigm acceptable
- **All MPLs & Eq. Groups (BI, CO, RF, ...):**
 - Get better data rate/volume estimates: detailed list of devices, variables per device class (1st iteration), Post-mortem data → complete [Excel sheet](#) (send to R. Steinhagen)
 - check logging time constants – particularly for smaller machines (CR, HESR, ...)
- **CSCO: L. Hechler, J. Fitzek, S. Jülicher:**
 - first conceptual proposal for Archiving System architecture and required HW
- **R. Bär, L. Hechler, J. Fitzek, R. Steinhagen:**
 - Archiving System specification for circulation/approval by Q4-2015

Discussion:

R. Bär confirmed that the Archiving System requirements needs to be collected during summer 2015, implement during 2016, and deployed (in an early prototype version) for SIS18 recommissioning in 2017.

S. Petri affirmed that he would cover and fill-in the requirements deriving from Super-FRS operation, but asks who would cover other generic or specific equipment such as vacuum system. R. Steinhagen replied that both the vacuum group and other experiments or MPLs should provide the data and that the filtering and handling of duplicates would be handled off-line.

M. Steck asked whether and where the reference values for the equipment are being stored and whether changes that bypass LSA (the new settings management system) are being tracked as well. R. Steinhagen iterated that LSA preferably stores and handles reference settings since it not only intrinsically provides a history of all past references and their changes but also that these references settings could be easily restored to earlier values. Nevertheless, as also indicated in the Excel sheet, the Archiving System could easily store also reference values alongside the actual measured variables. J. Fitzek confirmed that LSA is designed to keep the history even for two or more year-old reference data. D. Ondreka pointed out that there are some limits for keeping and restoring the reference settings via LSA, e.g. if the machine modelling (parameter propagation, incorporation rules, etc.) has been changed in the mean time.

R. Bär outlined the scope of that CSCO will define the technical constraints, and Cosylab will, in response, propose a design based on these requirements. The cut-off of which requirements are included that would be implemented by Cosylab and which need to be implemented by CSCO remains to be clarified.

C. Kleffner inquired whether there would be any constraints on the frequency. R. Bär clarified that there are some limits of old devices which need to be checked w.r.t. realisability of high acquisition rates, but both R. Bär and R. Steinhagen confirmed that for non post-mortem type data a rate of 1 kHz is being given as a reasonable upper limit.

R. Steinhagen mentioned that a tentative list of devices to be logged already exists and that this had been used to give the afore-mentioned bandwidth and data volume estimates (initial estimate ~10 MB/s, tens of TB short-term and similarly long-term storage). However, this list needs to be checked with requirements coming from other archiving users and data providers (BI, RF, ...) to ensure that not too many devices are missed for the first iteration.

Questions about the flexibility of the system have been raised and about what could be added at a later stage. R. Steinhagen confirmed that the targeted system will be designed to be very flexible and that new devices or variables could be added at a later stage. The purpose of the first survey is to get a better estimate of the required total data bandwidth and storage space. A better estimate is necessary for the specification and the layout of the system.

R. Bär elaborated further that the data could be stored mode-dependent, e.g. different data rates during shut-down. M. Steck asked whether the archiving would run all year. R. Bär confirmed this.

O. Kester asked whether all machines shall be covered by this Archiving System. R. Steinhagen: Yes. A. Reiter asked whether the data will be also accessible online for operators. R. Steinhagen

confirmed and stressed that this would be a primary requirement for the system alongside the data extraction interfaces for post-analysis.

A. Reiter asked how much of the logged data is expected to be actually used. R. Steinhagen estimated that based on previous LHC experience, nearly everything has been used at some point of time. If not immediately after being stored, then some period after for post-operation performance analysis. He stressed the importance of a user API and that any user is invited to use the data for their analysis. The aim of the system is to facilitate and permit a quantitative analysis of the machine operation and performance.

R. Bär pointed out that while an early version is targeted to be rolled out for the SIS18 recommissioning, that it would not be possible to implement all of the mentioned pre-analysis tools for the start-up but that these need to be rolled-out gradually.

R. Steinhagen emphasised that the purpose of having a large data coverage is that one could implement analysis scripts at a later stage that use past-operation data. O. Kester confirmed this assessment that even if no analysis tool exists, that one should start recording data!

The next meeting is planned for: Wednesday 17 June 2015, 15:00-17:00 (SE 1.124c)

Reported by J. Fitzek, S. Reimann, R. J. Steinhagen