ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-WECPL02

ROADMAP TO 100 Hz DAQ AT SwissFEL: EXPERIENCES AND LESSONS LEARNED

T. Celcer†, S.G. Ebner, F. Märki[1], A. Babic, L. Sala
Paul Scherrer Institut, 5232 Villigen PSI, Switzerland

Abstract

Providing a reliable and performant Data Acquisition System (DAQ) at Free Electron Lasers (FELs) is a challenging and complex task due to the inherent characteristics of a pulsed machine and consequent need of beam synchronous shot-to-shot DAQ, which enables correlation of collected data associated with each FEL pulse.

We will focus on experiences gathered during the process of moving towards 100 Hz operation at SwissFEL from the perspective of beam synchronous DAQ. Given the scarce resources and challenging deadlines, many efforts went into managing conflicting stakeholder expectations and priorities and into allocation of time for operation support and maintenance tasks on one side and time for design and development tasks on the other side. The technical challenges we encountered have shown a great importance of having proper requirements in the early phase, a well thought system design concept (which considers all subsystems in the DAQ chain), and a well-defined test framework for validation of recorded beam synchronous data

INTRODUCTION

SwissFEL Project Timeline

The SwissFEL free electron laser facility [1] is the newest accelerator at the Paul Scherrer Institute (PSI). The overall length of SwissFEL measures 720 m and consists of a 270 m long common accelerator line (injector, linac 1 and linac 2), followed by a switchyard, splitting it into two lines. The first, straight line, being part of the first construction phase, represents the third part of the linac and the hard X-Ray branch Aramis. Aramis currently serves two experimental stations, namely Alvra and Bernina. The project officially started in 2013 and was declared finished at the end of 2017.

In the currently ongoing second construction phase, a parallel, soft X-Ray branch Athos [2] will be added, delivering photons to two additional experimental stations, Maloja and Furka. The project timeline is 2018 – 2020 and first lasing is planned for December 2019. First pilot experiment in Maloja experiment station is planned for April 2020, while first pilot experiment in Furka is planned one year later, in April 2021. In addition, the project of building a third Aramis experimental station Cristallina is slowly starting. SwissFEL layout is shown in Figure 1.

† tine.celcer@psi.ch
[1] now at Institute for Data Science, School of Engineering, University of Applied Sciences NorthWestern Switzerland, 5210 Windisch, Switzerland

DAQ Challenges at FELs

In order to better understand the challenges of a FEL DAQ system, one first needs to understand the main differences between the continuous beam facility like a synchrotron and the pulsed machine like a FEL.

At PSI, three other large user facilities, operated before SwissFEL (SLS, SINQ and SμS), are all continuous beam sources. If we take SLS as a comparison, it delivers a stable, continuous beam to 18 user beamlines (or experimental stations). Due to the beam stability and continuity, the data obtained at the experimental station at a certain point in time does not depend (in large extent) on the machine parameter. This greatly simplifies the data acquisition process, as the timing requirements are quite relaxed and there is no need to correlate the experimental data with machine data (except for some asynchronous, slowly changing parameters).

The inherent characteristic of a pulsed machine is that the data readout has to be triggered synchronous to the electron or photon beam in order to collect useful data. As the beam (consisting of individual electron bunches) properties can fluctuate considerably from bunch to bunch, the data collected at the experimental station depends significantly on physical parameters of every bunch (pulse) on its path along the accelerator line, and therefore needs to be tagged appropriately. The BS DAQ system needs to store very large data sets and provide the ability to correlate the collected experimental data associated with each FEL pulse. For a 100 Hz repetition rate, this means that data (scalars, waveforms and images) all along the accelerator line have to be reliably collected, processed, tagged and recorded inside a 10 ms time slot. Therefore, the accelerator part of the FEL facility is much more involved in the experiment than in a synchrotron, even more so as usually a single bunch train delivers photons to only a single experimental station.

In addition, the users would like live (or close to live) analysis of beam synchronous data, which further increases the overall complexity as not only beam synchronous data recording but also beam synchronous data retrieval is needed.

All this presents an important shift in required technical skills of personnel working in the Control system section. A significant increase for experienced software developers and versatile control system specialists with higher knowledge level in real-time (embedded) systems arises.

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-WECPL02

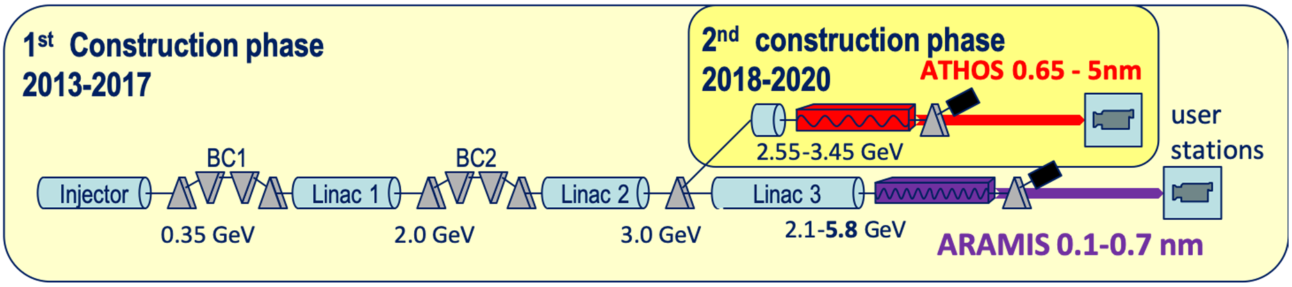

Figure 1: SwissFEL layout (phase 1 and phase 2).

BS DAQ System Design

The entire BS DAQ chain consists of several subsystems [3]. The first and most central building block is the timing system [4]. It is responsible for the reliable and precise distribution of timing triggers and timing information (global time stamp, pulse ID) to individual devices (sources). The sources are responsible for data readout, for tagging the data with the corresponding pulse ID and forwarding it as messages (using ZMQ protocol [5]) to the dispatching layer (each message consists of several channels). Most of the sources are EPICS IOCs [6], while some sources are implemented as real-time applications, therefore entirely bypassing the EPICS layer. Depending on the timing event sources subscribe to, they may readout and send the data with different frequencies (not necessarily with the beam repetition rate).

The dispatching layer is responsible for collection, synchronization and dispatching of BS channels. If requested by a client, it can directly dispatch a synchronized stream of a subset of BS channels, coming from different sources to that client. Besides, it forwards all channels to the buffering layer responsible for recording. Scalars and waveforms are recorded in the data buffer (currently 13 servers with local SSDs attached), while images are recorded in the image buffer (2 servers, attached to the GPFS cluster). The data remains in the buffering layer for a determined period of time (depending on the data type), during which it can be retrieved via a REST API. The dispatching and buffering layers are separated only logically, physically they are hosted on the same machines. The general layout of the entire beam synchronous DAQ system is depicted in Fig 2.

Large experimental pixel detectors are not integrated in the described DAQ system due to large volume of data they produce. For example, the biggest Jungfrau (JF) 16M detector produces 32 MB of raw data (32 modules, 0.5 Mpixels/module, 16 bit), which results in 3.2 GB/s when used at 100 Hz repetition rate. All versions of Jungfrau detectors write beam synchronous data (tagged with the pulse ID) directly to the experimental storage [7].

SwissFEL DAQ in Numbers

In SwissFEL there are currently over 12.000 beam synchronous channels recorded by the data buffer (scalars and waveforms) and approximately 60 cameras configured to record images into the image buffer (they never run all simultaneously).

Most cameras are operating with the beam repetition rate, which is currently 25 Hz, while a large portion of data buffer channels already delivers data with 100 Hz (e.g. all RF channels and most of diagnostic channels).

The total network throughput in the data buffer is approx. 65 MB/s per server, which gives a total of little less than 1 GB/s for 13 servers. The total incoming network traffic for the two image buffer servers is just below 2 GB/s and we are writing to GPFS storage with roughly 800 MB/s, with limit being 5 GB/s.

In addition, we are recording almost 360.000 asynchronous channels using the archiver appliance [8].

ROADMAP TO 100 Hz OPERATION

As already mentioned, it is foreseen that SwissFEL operates with the repetition rate of 100 Hz. In the initial project plan, the 50 Hz operation milestone was set very optimistically for July 2017 and the 100 Hz one by the end of 2017 (end of project). Although various beam synchronous sources could individually already cope with higher rates, the initial repetition rate was set to 10 Hz for two main reasons. Firstly, the BS data integrity was not validated (verification of correct data readout and tagging) and secondly large portion of individual sources as well as the system as a whole were not yet prepared to reliably handle the vast amount of data. Furthermore, insufficient experimental storage capabilities were preventing the use of large JF detectors at higher rates (currently still the case for the biggest detectors).

First pilot experiment, which took place in December 2017, at the end of the project phase, was therefore performed with the repetition rate of 10 Hz. The same rate was kept during first pilot experiments in 2018. In May 2018 the beam repetition rate was increased to 25 Hz. This was also the available repetition rate at the start of official user operation in 2019. Some user experiments in 2019 were performed already with 50 Hz, although relying mostly on the beam synchronous data from the main JF pixel detector (smaller version), which bypasses the BS DAQ backend infrastructure and writes data directly to the experimental storage. Nevertheless, as of today, SwissFEL officially still runs at 25 Hz.

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-WECPL02

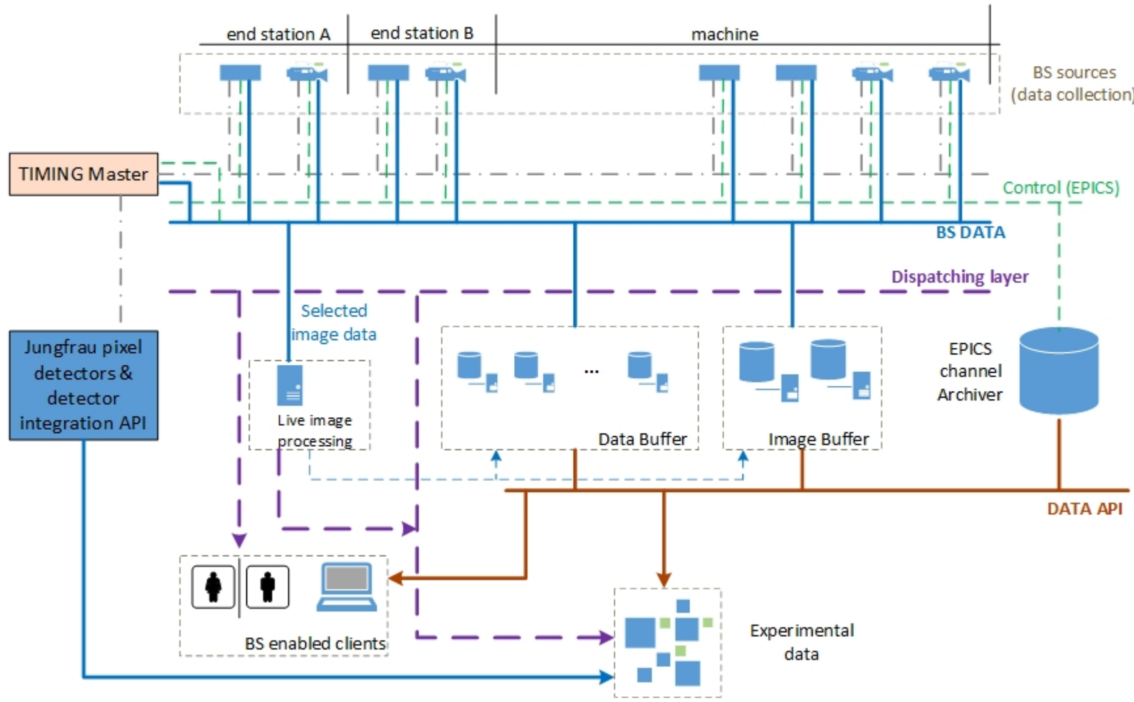

Figure 2: SwissFEL beam synchronous DAQ system layout.

Due to slow progression, increasing pressure from the users to move to higher repetition rates, limited resources and lack of prioritization combined with (too) high workload, we quickly saw a need for a coordinated discussion with relevant stakeholders. As the result of these discussions a document, describing the roadmap to reach 100 Hz operation, was presented in September 2018. It predicted the switch to permanent 50 Hz operation in the second half of 2019 and potential final move to 100 Hz in beginning of 2020 (pending experiences with 50 Hz operation). All important milestones and dates are depicted in Fig. 3.

EXPERIENCES AND LESSONS LEARNED

Many experiences with the BS DAQ system at SwisFEL were gathered in the past few years and many valuable lessons were learned, both from technical as well as from project and stakeholder management perspective. In this section some of the most valuable insights will be presented.

Requirements – Plan for Change

The first condition for a successful project is having clear requirements. For the beam synchronous DAQ system in an FEL this is not easy (if not impossible) to accomplish. The main problem is that the most demanding stakeholder – end station scientist – is just not yet there when requirements are collected. Hence, first input will come from the accelerator expert groups (diagnostics, beam dynamics, RF & LLRF, etc.), which have a different set of needs and priorities. One can make certain assumptions about the needs of experiments based on experiences of other FEL facilities, which can be used for the design of the DAQ system, but one has to be prepared for changes (which will certainly increase the complexity of the system), especially in the data aggregation and data retrieval part.

From the project management perspective, adopting certain agile principles during the development of the DAQ backend system will facilitate the team to better react to those changes and new feature requests. This makes it easier to change directions while still delivering a working product.

Planning Phase – System Design

Regardless of the resource situation and project time constraints, it is important to spend a lot of time on the initial system design phase. Having solid DAQ system bases will help the stability and ability to maintain and further develop the system. Based on our experiences, topics that play a very important role in the design phase are the following:

Network Topology is something that has to be planned in advance. Things like: network separations, firewall locations or which connections will require high throughput, will determine the usability of the DAQ system. In order to make correct decisions, one has to know the main dataflow path and which beam synchronous sources will be important (information usually not available in the design phase). The main lesson is: data flows mainly from the machine towards the end station. Therefore, avoid bottlenecks on any path towards the end users.

Data Retrieval Pay attention not only to data collection/recording but also to data retrieval (especially for image data) and shot-to-shot image processing. Topics like data retention policy, where will image data processing happen, which servers will be used, how they are placed

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-WECPL02

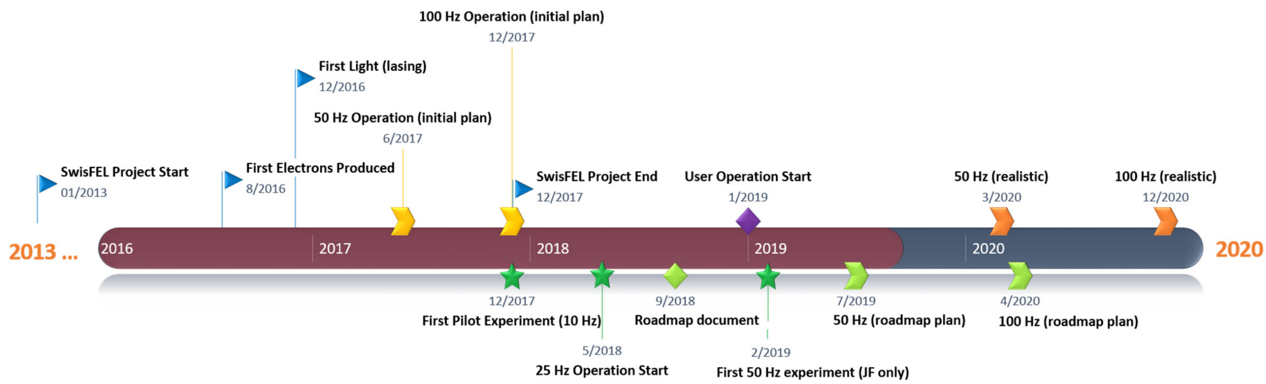

Figure 3: SwissFEL project and DAQ timeline.

inside the network topology, how much data can users (or intend to) request simultaneously and in what way will users access this data (live or delayed), should be addressed. Another important point from the user perspective is the ability to merge beam synchronous data (tagged with pulse IDs) with archived asynchronous data (tagged with local timestamps).

Integrated Design The reliability of the BS DAQ system depends on the reliability of individual subsystems in the entire chain (timing, individual sources and data recording/retrieval backend). Different teams are usually responsible for the development of different subsystems. It is important to avoid the silo-driven, non-collaborative development where each subsystem is designed and optimized without thinking about other links in the DAQ chain. Special care needs to be put into a collaborative design, definition of interfaces and optimization of the system as a whole. There is no benefit if the timing system itself is very precise, but delivers the timing data to beam synchronous sources in a way, which prevents optimal data collection and tagging. Or, for example, if the data backend and its protocol is designed in a way to simplify its architecture, but at the same time prevents sources to efficiently forward the data.

System Implementation Phase

During the implementation phase, the BS DAQ system could be logically split into two main parts. The first part are the data sources coupled with the timing system and the second part the data aggregation and data dispatching backend.

Data Collection

- **Optimize sources carefully and run data validation tests.** A lot of effort has to be spent on a reliable data collection. Non-optimal timing interface, timing jitter (jitter between subsequent pulses), bad implementation of EPICS IOCs, or non-optimal configuration of the embedded systems are most common causes for misbehaving sources. A single unreliable source, which cannot deliver beam synchronous data inside a given time slot (either delivers it with a delay, or skips pulse IDs) can create many problems in the DAQ system backend. However, such sources are easy to identify with proper diagnostics tools. A different problem are sources, which seemingly deliver valid data, where in fact the data is tagged with a wrong pulse ID (usually shifted for one pulse). Such problem will go undetected by diagnostic tools, and will result in false data correlation. Therefore, it is important to invest into proper HW-based data validation tests for individual sources and select a responsible person to perform these tests (usually the device responsible or the device control system integrator).
- **Consider bypassing the control system.** There will be many different beam synchronous sources using different type of hardware with different amount of data to be read-out. For some it will be more challenging to achieve the real-time requirements than for others. For time critical sources, it is worth considering bypassing the control system used in the facility (EPICS in our case) and performing beam synchronous data tagging at a lower level.

Data Aggregation and Dispatching

- **Detector/Camera Data Reduction** Provide tools for reducing data volumes (compression, lossy algorithms) from the very beginning, as it takes time to introduce them at a later stage, and convince users to spend time validating them. Try to implement the image processing routines already at the source (image acquisition servers) to mitigate the amount of data sent over the network (where applicable).
- **Diagnostic tools are important.** The user does not perceive the DAQ system as a collection of subsystems but as one unified system. The only important point is that they can obtain their data reliably. It is important to understand that they will always see the problems in the last link (data aggregation and data retrieval system). Therefore, it is very important to place a special care on the DAQ system diagnostic tools from the very beginning in order to be able to collect useful and valuable statistics. Having a general system overview (ideally centralized), containing information about corrupt sources, status of the sources (live/idle, connected/disconnected), status of the timing triggers/events and status of the DAQ servers, among others, will significantly decrease the troubleshooting time and simplify the search for the cause of a problem.

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-WECPL02

Resource Management and Operational Aspects

Skill Set The necessary skill set of control system engineers shifts when dealing with an FEL compared to a continuous light source. Efficient integration of beam synchronous devices in SwissFEL proved to be a much more complex task than the integration of devices into the control system in other facilities at PSI.

In addition, the need for experienced high-level software developers increases significantly due to the challenges of the BS DAQ backend. We learned that we were significantly understaffed in that area and unfortunately were not able to bring more resources to the team due to ongoing reorganization at PSI. Creating a solid team from the beginning is very important, as adding temporary personnel or new members to the team at a later stage is not an efficient solution.

Separation of Development and Operation Work This of course heavily depends on available resources. Ideally, you should try to keep the developers away from day-to-day operational aspects, 1st level user support and maintenance. If you have enough resources available, split the development and operational resources, but make them work closely together. As already mentioned in the previous section, this was/is not the case at SwissFEL. Developers spend a considerable amount of their working hours on operational aspects instead of software development.

Avoid Going into Operation Too Early Software commissioning time is crucial (especially when resources are scarce), but it is almost always neglected. Partially because it comes last and the pressure to keep the deadlines for operation start is already high due to accumulated delays during the project phase, and partially because it is simply underestimated during the (already optimistic) project planning phase. At SwissFEL hardly any time was allocated for the DAQ system commissioning before the start of pilot experiments in 2018. If too little time is dedicated to commission this crucial part of the FEL, it is highly probable that solid foundations will not be laid, and users will consequently suffer a lot during the early operation phase. In addition, once the FEL is in user operation, the pressure multiplies, available time for development and testing decreases, and experiment-specific requests have to be addressed as well. Lack of commissioning time combined with lack of resources proved to be the main reason for delays in reaching the 100 Hz milestone from the BS DAQ perspective.

Stability Over Performance - Avoid Workarounds as Much as Possible Do not increase the performance (e.g. go to a higher repetition rate or introduce new features) before the existing system is stable. This is, however, not so easy to accomplish, because users will push for performance even though they in principle understand the need for prioritizing stability. They also have their deadlines and pressure for providing scientific results. In such situation, it proves difficult to avoid workarounds or quick fixes, which (partially) provide users what they need but create much bigger problems in the long term. Aim to minimize such actions as much as possible, and rather concentrate on long-term, stable development. The more short-term alternative solutions you provide, the more problems you will have to converge back to a stable and maintainable DAQ system.

Managing Stakeholders

Effective stakeholder management is crucial if you want to achieve the goal of delivering a reliable DAQ system in due time.

Keep End Users Engaged In the beginning of this section we said that adopting some agile principles will help the team to react faster to change requests (usually coming from the side of the beamline scientists). Working closely with users is one of main agile principles and involving them in the development process is crucial. Hold regular meetings with users, plan regular tests with them, actively seek their feedback and enquire about their short- and mid-term plans.

Educate Users Educating users about the nature of the BS DAQ system and best usage practices is also vital to ensure stable operation of the DAQ system. Users will tend to request as much data at their disposal as possible. Not necessarily because they need all this data, but just because it might come useful. If they understand system limitations (data volume, throughput) and that requesting too many data at once can bring down the system, it is much more likely that they will use the system with more care.

Manage User Expectations When operation starts, users will be under a lot of pressure to produce scientific output, and the BS DAQ system is crucial for this. Accelerator expert groups will have different set of priorities, which will also involve the DAQ system. The DAQ team will be faced with competing requests from different stakeholders. Some will push for stability improvements, others will push for increasing repetition rates, and some will push for adding new features. In addition, the DAQ team will have their own priority tasks concerning the core DAQ system. When resources are scarce and time is limited, it will be impossible to deliver everything in due time. Expectation management is vital if you want to avoid high levels of frustration among involved parties. However, no matter how successful you are in managing user expectations, you can only alleviate this problem to a certain extent, but cannot avoid it completely.

Force High Management to Set Priorities If agreement on priorities cannot be reached in direct interaction with requesters themselves (usually the case), prioritization has to be requested from the board of higher-level management, representing all involved organizational units. As the high-level managers cannot evaluate the importance of tasks by themselves, this request has to contain clear information about task types, time estimates, benefits and consequences if the task is performed or not. In addition, it should be accompanied with the proposal of priorities from the perspective of the DAQ team. Pushing high-level management to take these decisions will also force them to take responsibility and be more aware of the status of the project.

ISBN: 978-3-95450-209-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2019-WECPL02

Depending on the organizational structure, this task might be more or less challenging. At PSI, Control system section belongs to Large Research Facilities division (GFA), the same as operations section and different expert groups. Photonics groups (end stations and photon diagnostics) belong to Photon Science research division (PSD) and the Science IT section, which works closely with controls section in DAQ related topics, belongs to logistics division (LOG). Consequently, up to third level management (out of four levels) had to be involved to reach a certain level of agreement on priorities (still not entirely successful).

CONCLUSION

The described BS DAQ system has been used in daily operation of the SwissFEL facility, which started with official user operation in 2019. In a great deal due to lack of resources and lack of commissioning time, it is not in its final state yet. It shows certain reliability issues, especially with beam synchronous image retrieval, which is increasingly used at the experimental stations. Significant efforts are put into solving existing problems, while at the same time keeping the support of running user experiments at a high level. In order to first achieve the desired reliability, the operation is currently kept at 25 Hz, with plan to move to 50 Hz operation in the following months, and to 100 Hz operation some time in 2020.

ACKNOWLEDGEMENT

The authors would like to take this opportunity to acknowledge the great effort of many people that contributed their pieces into the SwissFEL beam synchronous DAQ system puzzle.

REFERENCES

[1] C. Milne *et al.*, "SwissFEL: The Swiss x-ray free electron laser", *Applied Sciences (Switzerland)*, vol. 7, no. 7, art. no. 720, 2017. doi:10.3390/app707072

[2] R. Abela *et al.*, "The SwissFEL soft x-ray free-electron laser beamline: Athos", *J. Synchrotron Radiat.,* vol. 26, no. 4, pp. 1073-1084, 2019. doi:10.1107/S160057751900392

[3] S. G. Ebner *et al.,* "SwissFEL - beam synchronous data acquisition - the first year", *in Proc. 16th Int. Conf. on Accelerator and Large Experimental Physics Control Systems (ICALEPCS'17)*, Barcelona, Spain, Oct. 2017, pp. 276-279. doi:10.18429/JACoW-ICALEPCS2017-TUCPA06

[4] B. Kalantari and R. Biffiger, "SwissFEL timing system: First operational experience", *in Proc. 16th Int. Conf. on Accelerator and Large Experimental Physics Control Systems (ICALEPCS'17),* Barcelona, Spain, Oct. 2017, pp. 232-237. doi:10.18429/JACoW-ICALEPCS2017-TUCPL04

[5] ZeroMQ, http://zeromq.org

[6] EPICS, http://www.aps.anl.gov/epics/

[7] A. Mozzanica *et al.*, "Characterization results of the JUNGFRAU full scale readout ASIC", *J. Instrum.,* vol. 11, no. 2, p. C02047, 2016. doi:0.1088/1748-0221/11/02/C02047

[8] M. Shankar, L. Li, M. Davidsaver, and M. Konrad, "The EPICS archiver appliance", *in Proc.15th Int. Conf. on Accelerator and Large Experimental Physics Control Systems (ICALEPCS'15)*, Melbourne, Australia, Oct. 2015, pp. 761-764. doi:10.18429/JACoW-ICALEPCS2015-WEPGF030