ISBN: 978-3-95450-214-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2021-TUPAB291

SUBSYSTEM LEVEL DATA ACQUISITION FOR THE OPTICAL SYNCHRONIZATION SYSTEM AT EUROPEAN XFEL*

M. Schütte†, A. Eichler, T. Lamb, V. Rybnikov, H. Schlarb, T. Wilksen‡,
Deutsches Elektronen-Synchrotron (DESY), Hamburg, Germany

Abstract

The optical synchronization system for the European X-Ray Free Electron Laser provides sub-10 femtosecond timing precision [1] for the accelerator subsystems and experiments. This is achieved by phase locking a mode-locked laser oscillator to the main RF reference and distributing the optical pulse train carrying the time information via actively propagation-time stabilized optical fibers to multiple end-stations. Making up roughly one percent of the entire European XFEL, it is the first subsystem to receive a large-scale data acquisition system [2] for storing not just hand-selected information, but in fact all diagnostic, monitoring and configuration data relevant to the optical synchronization available from the distributed control system infrastructure. A minimum of 100 TB per year may be stored in a persistent archive for long-term health monitoring and data mining whereas excess data is stored in a short-term ring buffer for high-resolution fault analysis and feature extraction algorithm development. This paper describes scale, challenges and first experiences from the optical synchronization data acquisition system.

INTRODUCTION

Data Acquisition (DAQ), storage and evaluation are increasingly relevant topics at experimental physics facilities such as particle accelerators and colliders. For this reason, a shot-synchronized, bunch-resolved data acquisition system [2] has been added to the design of the accelerator control system Distributed Object-Oriented Control System (DOOCS) [3] at the European X-Ray Free Electron Laser (European XFEL) [4], with the explicit task to collect and store data from the accelerator's beam diagnostics and RF devices. Since starting operation in 2017, several such DAQ instances have been deployed for different purposes and subsystems at European XFEL and middle-layer diagnostic and feedback services running on the online DAQ data streams have become essential to the successful operation of the machine [2]. Over the years since its inception during the development of the Tesla Test Facility (TTF), the DAQ system has continuously evolved to keep up with increasing requirements created by the newer facilities Free Electron Laser in Hamburg (FLASH) and European XFEL and has proven to be a scalable and adaptable solution. These requirements do not only stem from a higher number of RF stations and diagnostic devices, but also from the increased complexity of said subsystems as well as completely new subsystems needed to achieve sufficiently good beam parameters and synchronization for the experiments. One of these newer subsystems is the optical synchronization system [5] which provides the necessary synchronization performance between the facility's main RF oscillator and the injector laser, the RF stations and the experiments. So far, little data is available on the long term performance of the optical synchronization system and the health of its components, for example the used mode-locked lasers, which are known to be sensitive to configuration and environmental influences. Despite being tightly integrated with the DOOCS control system, no data from the optical synchronization system used to be stored in any of the DAQ instances at European XFEL. This has been a hindrance in the past, as observed anomalies often could not be analyzed thoroughly afterwards. To address this issue and in an effort to further the use of intelligent process control and automated algorithms at DESY and European XFEL, the optical synchronization system has been selected to be the first subsystem to be equipped with a subsystem level DAQ system, which differs from pre-existing DAQ instances by the intent to collect all monitoring and configuration data available in the control system for long time periods as opposed to storing selected data channels for just a few days in which data may be tagged for archiving before being discarded per default. The following sections will in further detail discuss the requirements for this new DAQ instance as well as the encountered challenges and design choices made to address them.

ARCHITECTURE

Data Flow

Figure 1 illustrates the flow of data from the physical process and the front-end device server, which transforms the raw ADC data and contributes additional configuration data, to the DAQ server and the long term storage archive dCache [6]. For a detailed discussion on the inner architecture of the DAQ system, the reader is referred to [2]. With the shown multi-level architecture, inevitable trade-offs between local processing power, network communication channel bandwidth, latency, data resolution and storage space can be addressed with great flexibility. Data streams that would exceed the Ethernet uplink bandwidth of the front-end may directly be preprocessed locally on the local FPGAs and CPU. Higher level algorithms may be offloaded to the DAQ server, on which middle-layer services provide feedback information to the front-ends in a low-latency manner. These services are able to aggregate information from the different front-ends as is often needed in supervisory algorithms. For

* The authors acknowledge support from DESY (Hamburg, Germany), a member of the Helmholtz Association HGF.
† maximilian.schuette@desy.de
‡ tim.wilksen@desy.de

ISBN: 978-3-95450-214-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2021-TUPAB291

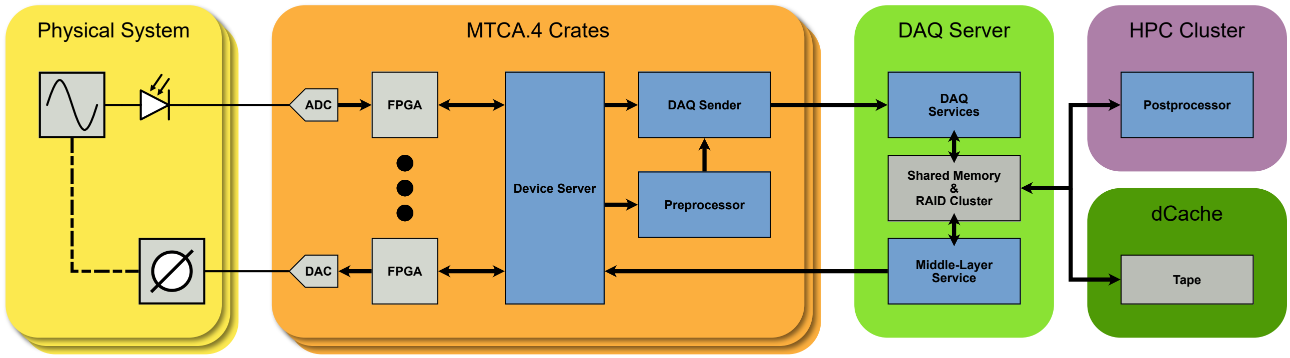

Figure 1: Data acquisition flow overview.

research and development, data may be transferred temporarily to a High Performance Computing (HPC) cluster, where the data may further be merged with data from experiments to train models or perform other big data operations. Batch operations scheduled on the HPC cluster provide ways to offload computationally expensive feature extraction and reduction tasks, optimizing the storage efficiency on the long-term archive.

Configuration Management

Several new challenges arose from the requirement to be able to replicate the state of the optical synchronization system as closely as possible from the archived data over long time periods. Whereas existing DAQ instances at European XFEL are concerned with collecting core data channels e.g. for beam diagnostics, which allowed the DAQ configuration to be managed by hand, this is no longer feasible with the increased number of data channels needed to replicate the full state of a given subsystem. Adding to the configuration management complexity is the fact that data channels made available by the device server continuously change as features implemented in the server get modified, expanded or removed. It is therefore not sufficient to build the required configuration once, but instead it needs to be continuously updated as the system evolves, reinforcing the need for an automated solution. This is addressed using a version controlled central configuration database from which the decentralized DAQ configuration is derived, illustrated in Fig. 2. A suite of scripts ensures that the database contents are semantically valid and consistent with the data channels available on the control system, tracks changes in the set of data channels provided by the front-end server so that the configuration can be updated to reflect these changes and generates and distributes the aforementioned decentralized configuration files. This minimizes the work for the maintainer, who merely has to add metadata labels to the newly found data channels, assign them to individual DAQ data streams as required and ensure that the processing power and communication bandwidth limits on each participating node are not exceeded. To avoid duplicate metadata labels in the database for data channels with identical purpose but different address, e.g. the arrival time measured at each of the Bunch Arrival Time Monitors (BAMs) along the beamline, metadata labels are associated with the data channels by means of Regular Expression (REGEX) patterns. Each unmasked live data channel visible in the DOOCS control system must be matched by exactly one such metadata template and this is checked during the semantic verification check. As a result, the number of templates to be maintained is drastically reduced, consistency between equivalent data channels is maximized and errors in the hand-crafted metadata are less likely to be missed.

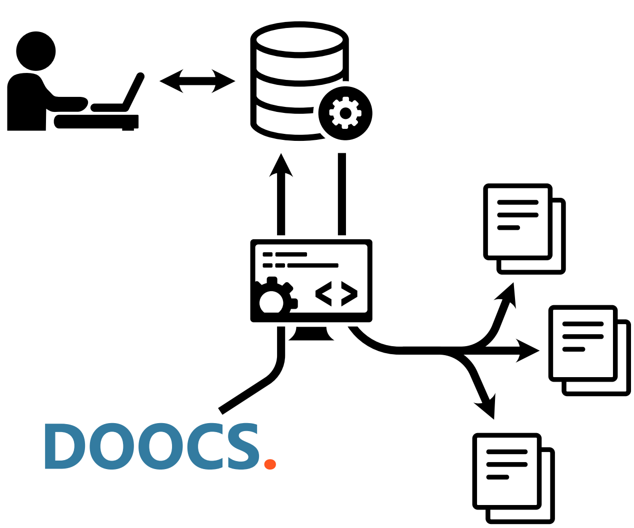

Figure 2: Centralized configuration database workflow.

STATUS AND EXPERIENCES

Conceptualization of the subsystem level DAQ for the optical synchronization system at European XFEL has started in mid 2019. First scalability and proof-of-concept tests were conducted in a lab environment in early 2020. Here, data-rates of several $\mathrm{Gbit\,s^{-1}}$ have been demonstrated. A first configuration covering the core systems of the optical synchronization system has gone into operation after the European XFEL summer maintenance shutdown of 2020 and configuration for additional remaining auxiliary and attached systems was added in Q1 of 2021. The current configuration covers over 115000 available data channels

ISBN: 978-3-95450-214-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2021-TUPAB291

from the front-end servers on 30 MicroTCA.4 crates related to the synchronization, which are grouped by and labeled with over 3800 metadata templates. Of these, nearly 40000 data channels are marked for acquisition in the central configuration database. Over 98 % of the estimated data-rate of up to 3.2 Gbit s^{-1} are accounted for by less than 1 % of the data channels, mostly highly-timed resolved controller input and output data. These channels are arguably not required to sufficiently replicate the system state, and therefore do not need to be stored long term. However, for research and post processing, they will be held available in a file system ring buffer, requiring approximately 35 TB per day ring-buffer length. Momentarily, they are excluded from the configuration until their impact on the network infrastructure can be thoroughly evaluated in the summer maintenance shutdown of 2021. The long-term storage requirement of the remaining data is estimated at 75 TB per year. This includes data format overheads, which at an estimated factor two is especially significant for the major share of scalar valued data channels and is sought to be reduced in the future.

Data from the optical synchronization DAQ system collected during a longitudinal arrival time stability machine study in October 2020 has been been successfully used to identify effects of frequency instabilities in the main timing reference on the beam arrival time at European XFEL by correlating data from the actively length stabilized fiber links used to transmit the optical timing signal and arrival time data from bunch and photon arrival time monitors.

A machine learning based mode-locked laser health detection algorithm currently in development in collaboration with the Helmholtz AI Consultant Team at Helmholtz Zentrum Dresden Rossendorf [7] is also intended to run as middle-layer service on the acquired data once production ready. Access to the data from within the facility is primarily made available by means of a Python API. For sharing data with collaborators outside of DESY / European XFEL, data may be converted to a HDF5 file, which can be read by all major programming languages used in data processing today.

CONCLUSION

The subsystem level DAQ system for the optical synchronization system at European XFEL has successfully demonstrated the scalability of the DOOCS based shot-synchronous and bunch-resolved DAQ system at European XFEL to a high number of data channels in conjunction with high data-rates of several Gbit s^{-1}. A version controlled centralized configuration database and supporting tools ensure that the DAQ system is well prepared for future changes in the monitored subsystem and that data collected by the system is interpretable even after long time periods and therefore usable for long term system performance evaluation. Extension of this methodology to further subsystems at European XFEL is in preparation.

REFERENCES

[1] S. Schulz *et al.*, "Few-Femtosecond Facility-Wide Synchronization of the European XFEL", in *Proc. FEL'19*, Hamburg, Germany, Aug. 2019, pp. 318–321. doi:10.18429/JACoW-FEL2019-WEB04

[2] T. Wilksen *et al.*, "A bunch-synchronized data acquisition system for the european xfel accelerator", in *Proc. ICALEPCS'17*, Barcelona, Spain, Oct. 2017, pp. 958–961. doi:10.18429/JACoW-ICALEPCS2017-TUPHA210

[3] T. Wilksen *et al.*, "The control system for the linear accelerator at the european xfel: Status and first experiences", in *Proc. ICALEPCS'17*, Barcelona, Spain, Oct. 2017, pp. 1–5. doi:10.18429/JACoW-ICALEPCS2017-MOAPL01

[4] R. Abela *et al.*, *XFEL: The European X-Ray Free-Electron Laser - Technical Design Report*. Hamburg, Germany: DESY, 2006.

[5] T. Lamb *et al.*, "Large-Scale Optical Synchronization System of the European XFEL with Femtosecond Precision", in *Proc. IPAC'19*, Melbourne, Australia, May 2019, pp. 3835–3838. doi:10.18429/JACOW-IPAC2019-THPRB018

[6] P. Fuhrmann, dCache, the Overview, https://www.dcache.org/old/manuals/dcache-whitepaper-light.pd

[7] Helmholtz AI Consultant Team for Matter Research, https://www.hzdr.de/db/Cms?pOid=60710&pNid=0