Enhancing Data Acquisition and Fault Analysis for Large-Scale Facilities: A Case Study on the Laser-Based Synchronization System at the European X-Ray Free-Electron Laser

Arne Grünhagen[1,3], Maximilian Schütte[3], Annika Eichler[2,3], Marina Tropmann-Frick[1] and Görschwin Fey[2]

[1]*Hamburg University of Applied Sciences, HAW, Hamburg, Germany*
[2]*Hamburg University of Technology, TUHH, Germany*
[3]*Deutsches Elektronen-Synchrotron DESY, Germany*

Abstract

The laser-based synchronization system of the European X-Ray Free-Electron Laser is responsible for precisely synchronizing various components within the large scale facility. It comprises several embedded components that are directly connected to the accelerator control system. In this study, we introduce a data acquisition system, which is integrated into the control system and builds the base for data-driven root cause analysis and predictive maintenance. To optimize the data retrieval process, we extended the existing data acquisition system by a lightweight database system based on Apache Parquet. This extension significantly enhances the data readout speed by a factor of 2000 , enabling efficient processing of operation-critical data. Additionally, we present a user-friendly dashboard that visualizes the operation-critical data, allowing for intuitive monitoring and analysis. Moreover, an unsupervised fault detection pipeline is created, capable of identifying faults retrospectively. Further validation through extensive real-world testing and deployment in daily operations is essential to ascertain the reliability and effectiveness of the integrated solution. This study serves as a foundation for future research and development efforts in optimizing and automating data acquisition and fault diagnosis methodologies for large-scale complex facilities, enhancing their robust performance and ensuring operational reliability.

Keywords

Data Acquisition, Large-Scale Data Management, Dashboard, Anomaly Detection

1. Introduction

The 3.4 km long linear accelerator-driven European X-ray Free-Electron Laser (EuXFEL) [1] opens cutting-edge research possibilities in molecular and material science, as well as system biology. To achieve the temporal precision required for these groundbreaking measurements, most subsystems within the EuXFEL demand timing accuracy within the femtosecond range. These requirements are met through the utilization of the laser-based synchronization (LbSync)

LWDA'23: Lernen, Wissen, Daten, Analysen. October 09–11, 2023, Marburg, Germany
arne.gruenhagen@haw-hamburg.de (A. Grünhagen); maximilian.schuette@desy.de (M. Schütte); annika.eichler@desy.de (A. Eichler); marina.tropmann-frick@haw-hamburg.de (M. Tropmann-Frick); goerschwin.fey@tuhh.de (G. Fey)

system [2], responsible for synchronizing the crucial accelerator components and experimental processes. Even a minor performance degradation in the LbSync system can significantly impact the proper execution of experiments, consequently resulting in the ineffective utilization of valuable and costly beam time. For that reason, an instance of the Data Acquisition System (DAQ) [3], which is tightly coupled to the accelerator control system and is already in use at other EuXFEL subsystems, has been added to the LbSync system three years ago [4]. The DAQ system is primarily responsible for extracting and storing data from the control system. Therefore, the DAQ is necessary to perform subsequent tasks such as data-driven root cause analysis and predictive maintenance. Predictive maintenance by monitoring and analyzing system parameters to detect early signs of potential issues, mitigates the risk of major failures and optimizes overall reliability and availability of the LbSync system of the EuXFEL.
The currently operated DAQ system and especially the way of data storage entail a very slow data readout ($< 8\,\mathrm{kB/s}$). In order to bridge this gap, we developed an extension based on Apache Parquet [5] to augment the existing system. The extension is the basis for efficient data visualization and applications such as fault analysis and predictive maintenance.

In Section 2, we present relevant literature concerning database systems for similar large-scale facilities. Section 3 gives an overview of the architecture of the LbSync system and its DAQ system, including the challenges faced. In Section 4, we discuss the main contribution of this work which is the extension to the existing DAQ system that effectively addresses these challenges. Finally, Section 5 concludes this work, summarizing the key ideas and underlining the importance of the proposed solutions to increase the LbSync system availability.

2. Related Work

Several researchers have made contributions to the field of predictive maintenance. The authors of [6] focus on developments in predictive maintenance, algorithms, methods, and the challenges encountered during implementation of predictive maintenance. The study presented in [7] provides an overview of the four phases of predictive maintenance: data acquisition, fault diagnostics, fault prognostics, and maintenance decision-making. The authors of [8] specifically address predictive maintenance for articulated robots, presenting a data acquisition strategy tailored to this application. Collectively, these papers enhance our understanding of predictive maintenance techniques for complex systems.

The authors of [9] present a definition of large-scale data systems that are similiar to the LbSync system, discuss respective challenges, and present a systematic framework to decompose large-scale data systems into four sequential modules: data generation, data acquisition, data storage, and data analytics. These four modules form a big data value chain. The research of [10, 11, 12] explores the applications, implications, and integration of data analytics within industry using large-scale databases, shedding light on the significance of data-driven approaches in transforming industrial processes and decision-making.

The research of [13, 14] centers on the application of data-driven analytics and large-scale database systems to industrial equipment maintenance and provide a set of data and system requirements for implementing equipment maintenance applications in industrial environments. The authors of [15, 16] evaluate the performance of popular Apache frameworks like Hadoop, Spark, and Flink for managing and processing large datasets. The authors of [17] present a database management system based on Apache Parquet for managing massive amounts of data from Internet of Things (IoT) devices. It integrates data collection, storage, management, and analysis functions. Similiar to these studies, we also decided to use an Apache framework as the basis for the LbSync DAQ extension.

Based on the review of related work, we present the experience, challenges, and the current state of the LbSync DAQ system. Utilizing the findings of the existing work, we developed an extension to the existing LbSync DAQ system, improving the data retrieval process.

The following three research papers present dashboards to depict data from big data systems. The work of [18] visualizes data from smart cities, including live environmental data and diagnostic overviews of society, demographics, health, and education. The authors of [19] present a general dashboard solution for several real-time industry cases. In [20], the authors focus on dashboards to visualize energy consumption and statistics for large buildings. In this study, we adapted their general design and architecture ideas to build a dashboard that provides an intuitive impression of the current health status of the LbSync system.

3. System Design and Implementation

3.1. The Laser-Based Synchronization System

A schematic overview of the EuXFEL and the LbSync system is given in Figure 1. The main oscillator (MO) [21] acts as the primary source for the synchronization system, playing a critical role in generating and delivering a radio frequency (RF) reference signal. This signal is essential for driving the linear accelerator within EuXFEL machine. The LbSync system is the foundation to reduce arrival time fluctuations by means of advanced feedback controls, while also synchronizing experimental laser systems with the accelerator for time-resolved measurements with unprecedented accuracy.

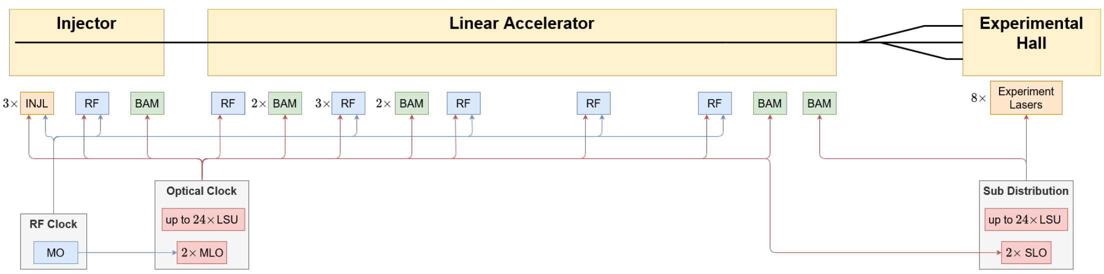

Figure 1: Schematic Overview of the Laser-Based Synchronization System of the European XFEL

The LbSync system consists of two redundant main laser oscillators (MLO) which are phase-locked to the 1.3 GHz RF signal of the RF MO, both emitting a laser pulse train with a pulse repetition rate of 216.667 MHz and a pulse duration of 200 fs. The phase of the MLO is actively stabilized using a PI controller in a phase-locked loop [22].

The pulse train from the MLO is split and transmitted to various fiber link stabilizing units (LSU) which actively stabilize their length. These fibers are employed to establish connections between the LSUs and the respective end-stations in the accelerator, such as laser synchronization setups, the RF re-synchronization units, and Bunch Arrival Time Monitors (BAM). Furthermore, a sub-distribution system is set up in the experimental hall 3.4 km away, in which another two laser oscillators are phase-locked to the LbSync system. This arrangement allows the sub-distribution laser oscillators (SLO) to replicate the MLO signal and to distribute the optical synchronization signals to additional end stations. This large setup ensures synchronization on a femtosecond level between the accelerator components and the various experimental lasers, facilitating advanced research at the EuXFEL.

All LbSync devices are driven by interconnected hardware running control algorithms. The devices generate data through different means: sensors attached to the devices, diagnostic data derived as a result of the control algorithms, and configuration parameters. The LbSync system is a dynamic system subject to ongoing maintenance, upgrades, and configuration changes. To obtain a comprehensive overview of the entire system, it is essential to capture monitoring, diagnostic, and configuration data from all LbSync components.

In total, we collect data from the following components as depicted in Table 1.

Table 1
Overview of LbSync Devices and Their Respective Data Channels

Device	< 10 Hz	10 Hz	> 10 Hz	total
2 Main Laser Oscillators (MLO)	2319	498	20	2837
2 Sub distribution Lasers Oscillators (SLO)	2313	509	24	2846
27 Link-Stabilizing Units (LSU)	14 378	2778	80	17 236
7 Bunch Arrival Time Monitors (BAM)	118	231	0	349
8 RF re-synchronization units	2898	576	54	3528
3 Injector Laser (INJL)	1155	240	12	1407
8 Experiment Lasers (EXPL)	7821	1641	81	9543
non device specific and environment	1499	976	0	2475
total	32 501	7449	271	40 221

3.2. Data Acquisition System

A pulse-synchronized data acquisition system has been seamlessly integrated into the accelerator Distributed Object-Oriented Control System (DOOCS) [23] of the EuXFEL. Its primary objective is to collect and store data from beam diagnostics and RF devices of the accelerator. Since the inception of the EuXFEL, multiple instances of this DAQ system have been deployed for various purposes and subsystems, playing a vital role in the operation of the facility.

The LbSync DAQ system aims to capture all available monitoring, diagnostics, and configuration data from the LbSync system over extended periods, distinguishing itself from existing DAQ systems that only store selected data channels for a limited duration. Table 1 gives an overview of how many data channels are planned for long-term storage and their respective data rate.

3.2.1. Data Flow

Figure 2 illustrates the data flow from the the physical systems to the final long term storage. In the LbSync DAQ system, multiple physical devices such as laser oscillators and link stabilizing units are connected to respective device hosts, which enables the exchange of monitoring, configuration, and diagnostics data specific to each device. Each device host executes at least one DAQ sender instance, responsible for transmitting its data to a centralized DAQ server through middle-layer services. The DAQ server collects data packages from all the DAQ senders, where each package comprises approximately 60 s of data. These packages are merged and stored in a single file using a proprietary raw file format, which is optimized for synchronized and compressed storage of data. This means that all data sharing the same bunch ID or timestamp respectively is grouped together in the storage, making it easier and more efficient to access and retrieve data for specific bunch IDs or timestamps. To ensure data integrity, an automated process transfers the raw files to a long-term storage system utilizing dCache [24] technology. This architecture allows for comprehensive data collection, organization, and storage of the diverse data generated by the interconnected physical devices.

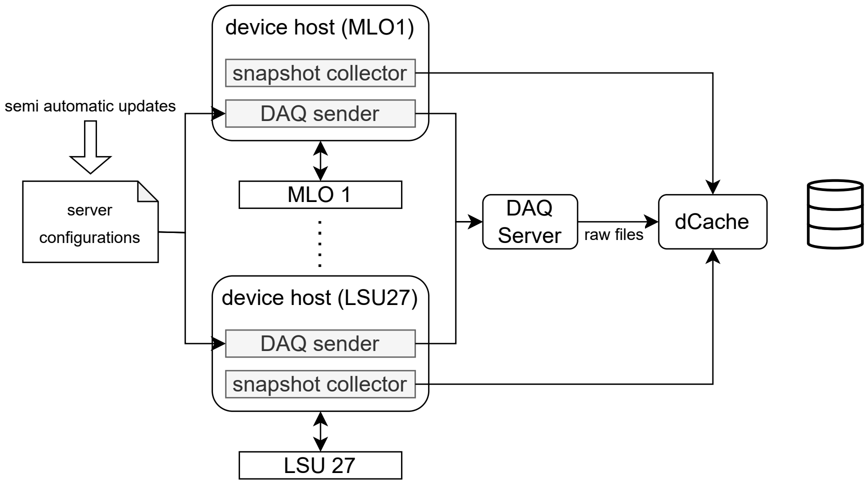

Figure 2: Data Stream of the Laser-Based Synchronization System at the European XFEL

The DAQ is designed to accommodate varying update rates for different data channels. These update rates range from infrequent updates up to 300 MHz. Slow channels primarily involve configuration properties, such as hostname or software versions, while fast channels encompass controller I/O or ADC/DAC data. Most of the DAQ channels have an update rate of 10 Hz, aligning with the accelerator's electron bunch rate. However, due to network limitations, the current infrastructure cannot handle the data throughput of the fast DAQ channels. As a result, the system currently skips most of the fast channels.

Fast data channels however can provide valuable insights into the overall performance of a system. To capture the most critical fast data channels, we utilize snapshot collectors running on the respective device hosts. Due to limited storage space on the device hosts, it is not feasible to record the fast channels continuously. In order to process all of the fast data channels, we implemented a strategy to collect operation-critical data for a fixed duration of 10 s every day at 5 am. These periodic snapshots complement the DAQ system by allowing us to obtain a daily overview of the system's health under the same conditions.

3.2.2. Configuration of the DAQ Instances

The configuration of DAQ senders on the device hosts plays a crucial role in determining which data channels are sent to the central DAQ server. As the device server's features continuously evolve, the available data channels undergo regular modifications, expansions, or removals. Consequently, the configuration of the DAQ senders must be consistently updated to reflect these changes. To ensure efficiency, the process of updating DAQ sender and DAQ server configurations is performed semi-automatically using a version-controlled central configuration database, enabling comprehensive tracking of all changes. A detailed description of the configuration update process is given in [4].

3.2.3. Status and Experience

Figure 3 provides a detailed representation of the specific time periods during which the DAQ system was actively recording data. It is important to note that the DAQ is intentionally not operational during the EuXFEL shutdown maintenance phase, which is dedicated to rolling out major changes and updates to the device servers. Following the completion of device server updates, the DAQ undergoes its own maintenance phase. Due to the timing of server rollouts at the end of the shutdown maintenance phase, the DAQ maintenance naturally coincides with the initial period of EuXFEL operation. Overall, the DAQ system operated for approximately 54.70 % of the total EuXFEL operation time, resulting in the collection and storage of approximately 160 TB of data. The downtimes of the DAQ can be explained by the fact that the DAQ system itself was being commissioned.

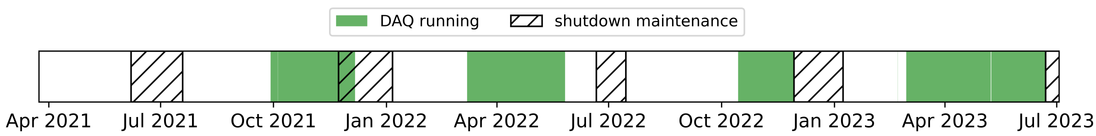

Figure 3: Availability of the DAQ Data

In contrast, Figure 4 provides an overview of the time periods during which the snapshot servers were operational. The snapshot servers were collecting data at 97.29 % of the EuXFEL operation days, which is significantly higher compared to the DAQ servers. This is primarily due to the minimal maintenance requirements of the snapshot servers, as the underlying DOOCS properties typically remain unchanged for extended periods. As a result, the snapshot servers reliably and consistently capture snapshots throughout the EuXFEL operation, resulting in an accumulation of 1.3 TB of data.

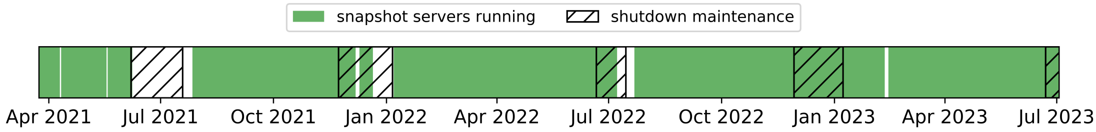

Figure 4: Availability of Snapshot Data

4. DAQ Extension

The current data acquisition (DAQ) system stores data as raw files, which poses limitations for efficient data retrieval. The primary objective of the LbSync DAQ is to quickly display and analyze operation critical data within specific time ranges and to support fault diagnosis and predictive maintenance applications. Both applications require a fast and reliable data readout process. The existing DAQ solution is not optimized for use cases where raw files have a large number of data channels. It has been primarily optimized for feedback applications and direct communication with device servers.

Visualizing a week's worth of operation critical data (90 different data channels) amounts to approximately 2.5 GB, taking more than 5 days to retrieve. Typically, the data readout speed is around 5 kB/s. This extended retrieval time is not acceptable for daily operations. We have extended the existing architecture (depicted in Figure 5) to overcome these challenges and incorporate improved data retrieval capabilities.

4.1. Data Processing

Figure 5 illustrates the data flow of the extension to the existing DAQ system. The DAQ system utilizes raw files which contain data from all data channels. To efficiently extract operation critical data, an automatic transformation script is executed as soon as a new raw file is created by the DAQ Server. This script extracts the operation relevant data, including the corresponding timestamps and bunch IDs. The step of extracting data from the raw files is the bottleneck to make the data accessible. We store this extracted data in Apache Parquet files, creating a separate Parquet file for each data channel and raw file. The organization of Parquet files is based on data channel names, resulting in a well-structured database. This solution satisfies the requirements of the dCache technology, which emphasizes write-once read many (WORM) files. The raw files are retained on the dCache system to preserve data from all channels for root cause analysis. Consequently, the extracted data is duplicated on the dCache. Additionally, the raw files are retained to allow potential future enhancements or modifications aimed at improving raw file management and handling.

Figure 6 illustrates the superior readout speeds achieved by this extension compared to the raw file solution. This results in a readout duration of $9.85 \cdot 10^{-3} \cdot n \cdot d$ sconds for the raw file solution and $5.448 \cdot 10^{-6} n \cdot d$ seconds for the Apache Parquet extension, where n corresponds to the number of extracted channels and d corresponds to the window length of requested data in seconds. Furthermore, the Apache Parquet extension not only enhances data readout speed but also supports access from multiple programming languages. Subsequently, we present two applications built on the Apache Parquet database: a dashboard and an unsupervised anomaly

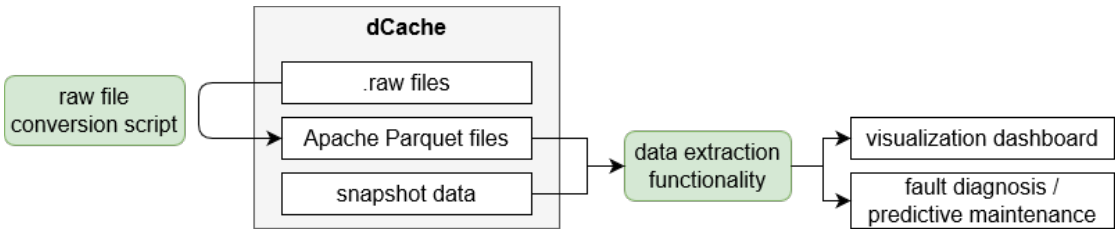

Figure 5: Schematic Overview of the Apache Parquet Extension

detection pipeline.

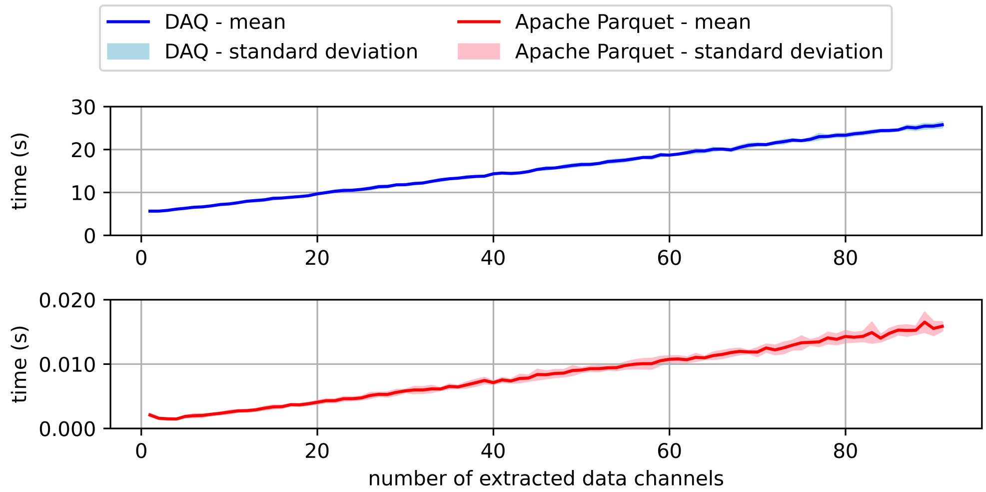

Figure 6: Readout Speed of Data Channels of 30 s of Data

4.2. Data Applications

4.2.1. Unsupervised Anomaly Detection

This subsection presents the efforts in developing an automatic fault detection system [25] that examines operation critical data obtained from the DAQ system and its subsequent integration into the daily operation of the LbSync system. To identify faulty or unusual system behavior, an unsupervised detection pipeline was designed for analyzing the time series data, as illustrated in Figure 7. The first step involves dividing the signals into 30 s segments. Subsequently, a compact set of features, similar to the proposed *EfficientFCParameters* features by tsfresh [26], is calculated for each signal. The features are then transformed using z-normalization. Next, an IsolationForest is fitted to the normalized features to generate an anomaly score for each 30 s window. If the anomaly score surpasses a pre-defined threshold, an automatic email will be sent to the system operators.

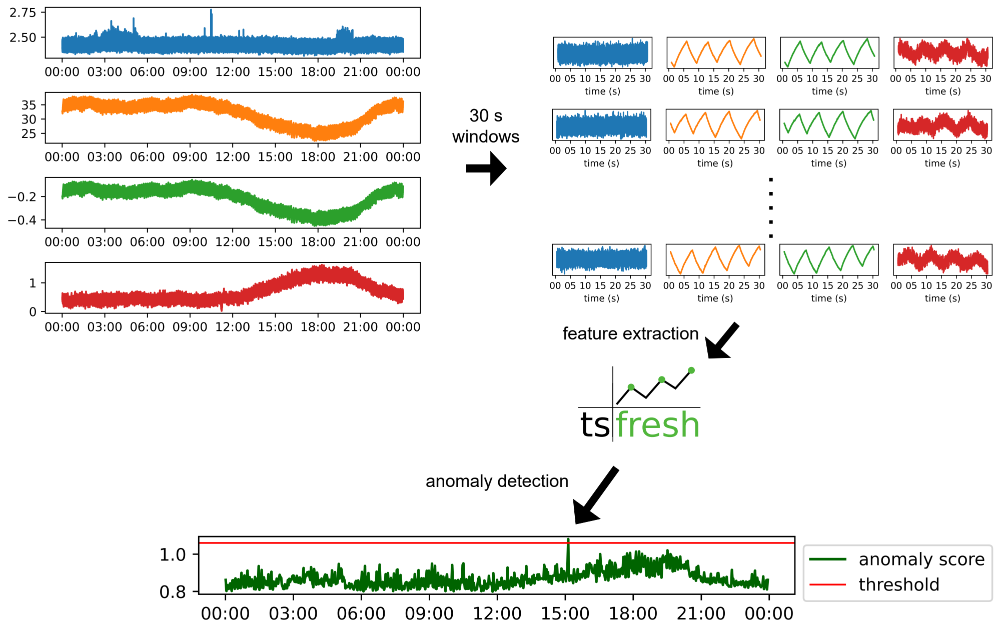

Figure 7: Unsupervised Anomaly Detection Pipeline

4.2.2. Dashboard

The visualization of data is facilitated through a Plotly [27] dashboard that is connected to the dCache long-term storage. The dashboard serves as a user interface (see Figure 8), allowing the selection of a specific LbSync device and day for analysis. It is designed to visualize both snapshot data and DAQ data in a uniquely combined view.

The snapshot data is presented in both the time and frequency domains, as both domains provide distinct information crucial for analyzing the health status of the systems. The dashboard application offers configurations for the frequency domain calculations such as window functions and sizes. Furthermore, operation critical LbSync data channels of the selected device are plotted for the selected day.

The operation critical signals are segmented into windows, and for each window, an anomaly score was calculated. These anomaly scores are further visualized on the dashboard, allowing for an intuitive representation of the detected anomalies in the time series data.

5. Conclusion

In our study, we presented the successful development of a comprehensive database system which collects signals from multiple sending servers running on device hosts interconnected with various physical devices. The semi-automatic configuration of these sending servers enabled the accumulation of 40 221 distinct data channels with different update rates. By implementing an extension based on Apache Parquet to the existing DAQ system, we significantly improved data retrieval speed for applications like history visualizations and fault analysis.

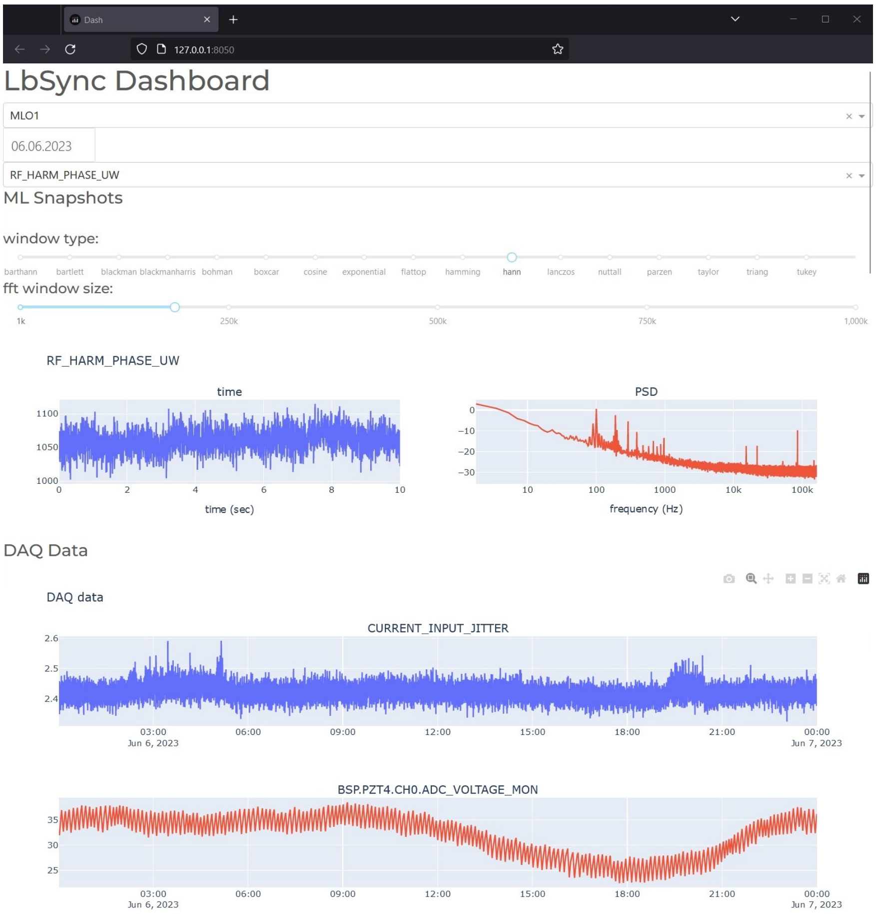

Figure 8: Lbsync Dashboard Showing Operation Critical Data of One Day

Moving forward, several improvements and next steps are identified. Firstly, incorporating root cause analysis and fault localization into the anomaly detection process would enhance effectiveness of the fault diagnosis. Functionalities should be added to the dashboard, allowing users to select any data channels and not only the predefined set of operation critical data. Furthermore, it should be noted that the dashboard and fault detection pipeline are not yet in daily use, so the design and layout of the dashboard components is not yet finalized. The reliability of the fault detection pipeline still requires validation through daily operation, ensuring its effectiveness and robustness.

Acknowledgments

We acknowledge the support by DASHH (Data Science in Hamburg - HELMHOLTZ Graduate School for the Structure of Matter) with the Grant-No. HIDSS-0002.

We acknowledge the support from DESY, a member of the Helmholtz Association HGF.

References

[1] E. Sobolev, S. Zolotarev, K. Giewekemeyer, J. Bielecki, K. Okamoto, H. K. N. Reddy, J. Andreasson, K. Ayyer, I. Barak, S. Bari, A. Barty, R. Bean, S. Bobkov, H. N. Chapman, G. Chojnowski, B. J. Daurer, K. Dörner, T. Ekeberg, L. Flückiger, O. Galzitskaya, L. Gelisio, S. Hauf, B. G. Hogue, D. A. Horke, A. Hosseinizadeh, V. Ilyin, C. Jung, C. Kim, Y. Kim, R. A. Kirian, H. Kirkwood, O. Kulyk, J. Küpper, R. Letrun, N. D. Loh, K. Lorenzen, M. Messerschmidt, K. Mühlig, A. Ourmazd, N. Raab, A. V. Rode, M. Rose, A. Round, T. Sato, R. Schubert, P. Schwander, J. A. Sellberg, M. Sikorski, A. Silenzi, C. Song, J. C. H. Spence, S. Stern, J. Sztuk-Dambietz, A. Teslyuk, N. Timneanu, M. Trebbin, C. Uetrecht, B. Weinhausen, G. J. Williams, P. L. Xavier, C. Xu, I. A. Vartanyants, V. S. Lamzin, A. Mancuso, F. R. N. C. Maia, Megahertz Single-Particle Imaging at the European XFEL, Communications Physics 3 (2020). doi:10.1038/s42005-020-0362-y.

[2] S. Schulz, et al., Few-Femtosecond Facility-Wide Synchronization of the European XFEL, in: 39th International Free Electron Laser Conference, 2019, p. WEB04. doi:10.18429/JACoW-FEL2019-WEB04.

[3] T. Wilksen, A. Aghababyan, L. Fröhlich, O. Hensler, R. Kammering, K. Rehlich, V. Rybnikov, et al., A Bunch-Synchronized Data Acquisition System for the European XFEL Accelerator, in: 16th International Conference on Accelerator and Large Experimental Physics Control Systems, PUBDB-2018-00678, Control-System, 2017.

[4] M. Schütte, A. Eichler, T. Lamb, V. Rybnikov, H. Schlarb, T. Wilksen, Subsystem Level Data Acquisition for the Optical Synchronization System at European XFEL, in: Proceedings of the 12th International Particle Accelerator Conference, IPAC2021, Campinas, SP, Brazil, 12th International Particle Accelerator Conference, Campinas (Brazil), 24 May 2021 - 28 May 2021, JACoW Publishing, Geneva, Switzerland, 2021, p. 3. URL: https://bib-pubdb1.desy.de/record/472171. doi:10.18429/JACOW-IPAC2021-TUPAB291.

[5] D. Vohra, Apache Parquet, Apress, Berkeley, CA, 2016, pp. 325–335. URL: https://doi.org/10.1007/978-1-4842-2199-0_8. doi:10.1007/978-1-4842-2199-0_8.

[6] E. Lughofer, M. Sayed-Mouchaweh, Predictive Maintenance in Dynamic Systems: Advanced Methods, Decision Support Tools and Real-World Applications, Springer, 2019.

[7] C. Nentwich, G. Reinhart, Towards Data Acquisition for Predictive Maintenance of Industrial Robots, Procedia CIRP 104 (2021) 62–67. doi:https://doi.org/10.1016/j.procir.2021.11.011, 54th CIRP CMS 2021 - Towards Digitalized Manufacturing 4.0.

[8] G. Xu, M. Liu, J. Wang, Y. Ma, J. Wang, F. Li, W. Shen, Data-Driven Fault Diagnostics and Prognostics for Predictive Maintenance: A Brief Overview, in: 2019 IEEE 15th International Conference on Automation Science and Engineering (CASE), 2019, pp. 103–108. doi:10.1109/COASE.2019.8843068.

[9] K. Zhou, T. Liu, L. Zhou, Industry 4.0: Towards Future Industrial Opportunities and Challenges, in: 2015 12th International Conference on Fuzzy Systems and Knowledge Discovery (FSKD), 2015, pp. 2147–2152. doi:10.1109/FSKD.2015.7382284.

[10] L. Duan, L. Da Xu, Data Analytics in Industry 4.0: A Survey, Information Systems Frontiers (2021) 1–17.

[11] L. Ardito, R. Cerchione, E. Mazzola, E. Raguseo, Industry 4.0 Transition: A Systematic Literature Review Combining the Absorptive Capacity Theory and the Data–

Information–Knowledge Hierarchy, Journal of Knowledge Management 26 (2022) 2222–2254.
[12] A. Sharma, H. Pandey, Big Data and Analytics in Industry 4.0, A roadmap to Industry 4.0: Smart production, sharp business and sustainable development (2020) 57–72.
[13] P. O'Donovan, K. Leahy, K. Bruton, D. T. O'Sullivan, An Industrial Big Data Pipeline for Data-Driven Analytics Maintenance Applications in Large-Scale Smart Manufacturing Facilities, Journal of big data 2 (2015) 1–26.
[14] A. Cakir, Ö. Akın, H. F. Deniz, A. Yılmaz, Enabling Real Time Big Data Solutions for Manufacturing at Scale, Journal of Big Data 9 (2022) 1–24.
[15] J. Veiga, R. R. Expósito, X. C. Pardo, G. L. Taboada, J. Tourifio, Performance Evaluation of Big Data Frameworks for Large-Scale Data Analytics, in: 2016 IEEE International Conference on Big Data (Big Data), 2016, pp. 424–431. doi:10.1109/BigData.2016.7840633.
[16] N. Ahmed, A. L. Barczak, T. Susnjak, M. A. Rashid, A Comprehensive Performance Analysis of Apache Hadoop and Apache Spark for Large Scale Data Sets Using Hibench, Journal of Big Data 7 (2020) 1–18.
[17] C. Wang, X. Huang, J. Qiao, T. Jiang, L. Rui, J. Zhang, R. Kang, J. Feinauer, K. A. McGrail, P. Wang, D. Luo, J. Yuan, J. Wang, J. Sun, Apache IoTDB: Time-Series Database for Internet of Things, Proc. VLDB Endow. 13 (2020) 2901–2904. URL: https://doi.org/10.14778/3415478.3415504. doi:10.14778/3415478.3415504.
[18] M. Farmanbar, C. Rong, Triangulum City Dashboard: An Interactive Data Analytic Platform for Visualizing Smart City Performance, Processes 8 (2020). URL: https://www.mdpi.com/2227-9717/8/2/250. doi:10.3390/pr8020250.
[19] M. Steinkamp, T. Mühlbauer, HyDash: A Dashboard for Real-Time Business Intelligence Based on the Hyper Main Memory Database System, in: Datenbanksysteme für Business, Technologie und Web (BTW) 2013 - Workshopband, Gesellschaft für Informatik e.V., Bonn, 2013, pp. 217–226.
[20] Y. Agarwal, T. Weng, R. K. Gupta, The Energy Dashboard: Improving the Visibility of Energy Consumption at a Campus-Wide Scale, in: Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, BuildSys '09, Association for Computing Machinery, New York, NY, USA, 2009, p. 55–60. doi:10.1145/1810279.1810292.
[21] L. Zembala, K. Czuba, B. Gąsowski, D. Sikora, J. Branlard, H. Schlarb, H. Weddig, Master Oscillator for the European XFEL, in: 5th International Particle Accelerator Conference, 2014, p. WEPRI116. doi:10.18429/JACoW-IPAC2014-WEPRI116.
[22] M. Heuer, Identification and control of the laser-based synchronization system for the European X-ray Free Electron Laser, doctoral thesis, Technische Universität Hamburg-Harburg, 2018. URL: http://tubdok.tub.tuhh.de/handle/11420/1706. doi:10.15480/882.1703.
[23] T. Wilksen, E. Sombrowski, M. Walla, S. Aytaç, W. D. Gerhardt, A. Petrosyan, J. Szczesny, L. Petrosyan, V. Petrosyan, J. Jaeger, O. Hensler, H. Kay, R. Kammering, C. Behrens, H.-D. Wu, P. Duval, P. Castro-Garcia, V. Rybnikov, G. Petrosyan, P. K. Bartkiewicz, W. Schütte, G. Schlesselmann, A. Aghababyan, K. Rehlich, M. Staack, J. Wilgen, S. H. Karstensen, L. Fröhlich, K. Hinsch, R. Bacher, The Control System for the Linear Accelerator at the European XFEL: Status and First Experiences, 2017.
[24] P. Fuhrmann, dCache Overview, Online, 2014. URL: https://www.dcache.org/manuals/

dcache-whitepaper-light.pdf.
[25] A. Grünhagen, J. Branlard, A. Eichler, G. Martino, G. Fey, M. Tropmann-Frick, Fault Analysis of the Beam Acceleration Control System at the European XFEL using Data Mining, in: 2021 IEEE 30th Asian Test Symposium (ATS), 2021, pp. 61–66. doi:10.1109/ATS52891.2021.00023.
[26] M. Christ, N. Braun, J. Neuffer, A. W. Kempa-Liehr, Time Series FeatuRe Extraction on basis of Scalable Hypothesis tests (tsfresh – A Python package), Neurocomputing 307 (2018) 72–77. URL: https://www.sciencedirect.com/science/article/pii/S0925231218304843. doi:https://doi.org/10.1016/j.neucom.2018.03.067.
[27] P. T. Inc., Collaborative Data Science, 2015. URL: https://plot.ly.