

MACHINE LEARNING ASSISTED CAVITY QUENCH IDENTIFICATION AT THE EUROPEAN XFEL

J. Branlard[†], A. Eichler, J. Timm, N. Walker
DESY, Hamburg, Germany

Abstract

A server-based quench detection system is used since the beginning of operation at the European XFEL (2017) to stop driving superconducting cavities if they experience a quench. While this approach effectively detects quenches, it also generates false positives, tripping the accelerating station when failures other than quenches occur. Using the post-mortem data snapshots generated for every trip, an additional signal (referred to as residual) is systematically computed based on the standard cavity model. Following an initial training on a subset of such residuals previously tagged as “quench” / “non-quench”, two independent machine learning engines analyse routinely the trip snapshots and their residuals to identify if a trip was indeed triggered by a quench or has another root cause. The outcome of the analysis is automatically appended to the data snapshots and distributed to a team of experts. This constitutes a fully deployed example of machine-learning-assisted failure classification to identify quenches, supporting experts in their daily routine of monitoring and documenting the accelerator uptime and availability.

INTRODUCTION

The European X-ray free electron laser (EuXFEL) is a user facility delivering ultra-short hard and soft X-ray flashes with the highest brilliance worldwide, through three undulator lines and serving up to six experiments. It is based on a 10 Hz pulsed 17.5 GeV superconducting linac, commissioned since 2017. Large accelerators require a high level of automation, in particular for trip detection, classification and recovery. *XTLReport* is a tool developed over the last years [1] to track the linac uptime and categorize trips according to their root cause. In its current implementation, once an hour, the tool monitors for each RF station a total of 50 hardware interlock histories available in the control system, coming from the following subsystems: klystron, modulator, coupler, quadrupoles, cryogenics, vacuum and machine protection system. On top of these hardware interlock, software interlock properties (for example coming from the quench detection system or the finite state machine) are also being monitored. Often a trip fires several interlocks, requiring experts to look at the post-mortem time signals to reconstruct the chain of events that lead to the trip. However, the accelerator conditions at the time of the trip and the sequence of interlock often provide a unique signature allowing for trip classification, so that once identified, this process can be automated. *XTLReport* has gone through several iterations, each time assigning a newly found sequence of interlocks to a new root cause, hence providing live up/downtime accounting and

root cause analysis, or leaving the root cause unknown if it cannot (yet) be identified. The tool also fetches DAQ data of the tripped station, before (20 sec) and after (5 sec) the event, and bundles it into a data snapshot saved for post-mortem analysis. A classic example of such a trip is when the quench detection server (QDS) detects a quench and trips the station by stopping the RF to minimize the impact of the excessive cryo load induced by the quench on the cryogenic system. The QDS decision algorithm follows the standard approach of computing the cavity loaded quality factor (Q_L) for every pulse and comparing it to a running average [2]. Sudden drops in Q_L are interpreted as quench and the RF is stopped. While this approach proved extremely successful to minimize down time when quenches occur, false positives (i.e. “fake” quenches”) have also been observed. As a part of operations supervision, an expert routinely reviews the last trips, analyzes the RF waveforms and concludes if a quench was real or not. An example of a real and a fake quench, compared to a nominal pulse is given in Fig. 1. This plot illustrates the differences observed during the decay part of the RF pulse, where Q_L is computed by the QDS.

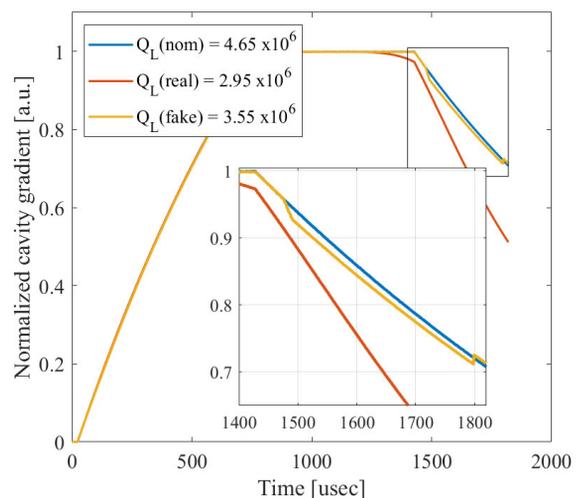


Figure 1: Normalized cavity gradient amplitude for a nominal, fake- and real-quenched pulse. The inset zooms on the decay section at the end of the pulse, when the RF is switched off.

As illustrated in Fig. 1, a real quench behavior corresponds to a steeper slope during the decay. However, other phenomena can occur during the pulse, interfering with the Q_L computation and fooling the QDS into detecting fake quenches. In the fake-quench example of Fig. 1, one can see a step-down in the decay around 1500 usec followed by a step-up around 1800 usec. This behavior (suspected to come from a failure in the digitization chain) is not a

[†]julien.branlard@desy.de

quench but still yields a lower Q_L and will trigger the QDS. More details about mechanisms that can fool the QDS and fake-versus-real quench statistics are given in [3]. The corresponding Q_L values computed over the first 100 usec of the decay are given in the legend of Fig. 1. The threshold for quench detection is set to 5×10^5 (i.e. 10% of nominal value), which corresponds approximately to a three order of magnitude drop of the unloaded quality factor. In both cases (fake and real quench), the threshold is exceeded. Empirical threshold adjustment to minimize the number of false positives was pursued, along with techniques to improve the robustness of the Q_L computation (such as extending the area of the decay over which Q_L is computed). These measures helped in minimizing the number of false positives. One should also mention that the QDS implementation at the EuXFEL is software-based (as opposed to a hardware- or firmware-based implementation). Due to possible delays between server communication, it can happen that the RF is not systematically stopped on the next pulse following a quench, but several pulses later (up to five in extreme cases).

The main motivation for this work was to develop a classification algorithm based on machine learning techniques to decide if the quench was real or not, looking at the available post-mortem data sets. Purely data-based approaches have been pursued in other labs [4]. The idea here is to exploit the known cavity model [5], hence requiring less training data. The approach is briefly presented in the following section; its benefits are illustrated in the results section and future work is discussed in the conclusion.

UPDATED APPROACH TO QUENCH IDENTIFICATION

The QDS makes use of the cavity probe signal to detect quenches. However, the cavity forward and reflected RF waveforms are also available and comply with the standard electrical cavity model [5]. One can make use of this additional information to get more insight on the nature of the trip, by exploiting differences between the detected cavity probe and the expected trace (or virtual probe) computed using the cavity model and the measured forward power. This idea was explored in [6], computing a figure of merit for each pulse, referred to as residual, which is an indicator of model deviation, defined as:

$$r(t) = \frac{-\dot{V}_{P,I}(t) + \omega_{1/2} \left(-V_{P,I}(t) + 2V_{F,I}(t) - V_{B,I}(t) \right)}{V_{P,Q}(t)} \quad (1)$$

$$\frac{\dot{V}_{P,Q}(t) + \omega_{1/2} \left(V_{P,Q}(t) - 2V_{F,Q}(t) + V_{B,Q}(t) \right)}{V_{P,I}(t)}$$

where the subscripts P , F , B , I and Q stand for probe, forward, beam, in-phase and quadrature voltages while $\omega_{1/2}$ is the cavity half-bandwidth. A generalized likelihood ratio (GLR) is then computed to quantify the probability that the measured residual indicates a fault. The complete derivation of residual and GLR is available in [3]. One benefit of using the GLR as metric is that it intensifies small

differences between probe and model, while being robust against changes due to standard accelerator operation such as cavity detuning. The three residuals corresponding to the three traces of Fig. 1 are shown in Fig. 2.

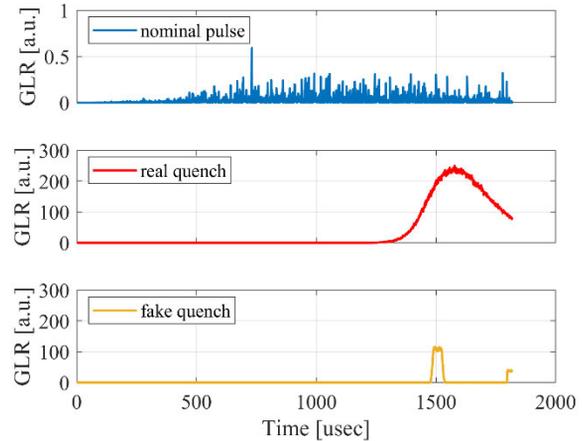


Figure 2: The Generalized Likelihood Ratio (GLR) proved to be a useful metric to categorize trips. The shape and magnitude of the GLRs differs greatly between trips.

The next step consisted in automating the trip review to decide if a quench was real or not. A training set of 453 trips were reviewed by an expert and tagged as “real” or “false”, to indicate if the trip was indeed a quench or was due to some other fault. An unsupervised classification based on k-means square was then used to define classes for the datasets tagged as quenches. A class threshold was defined using the quenched and non-quenched trips of the training data set. More details on the machine learning algorithm are given in a paper currently in preparation. The evaluation of new trips consists then of computing the distance to the class centre point. The trip is tagged as non-quench if this distance exceeds the threshold, as quench otherwise. This analysis is run daily as a cron job for every new trip snapshot [7]. The program computes the residual and GLR for all pulses of the snapshot and decides if the trip was a quench based on the GLR distance to the class centre point. A picture to highlight the result (Fig. 3) is generated:

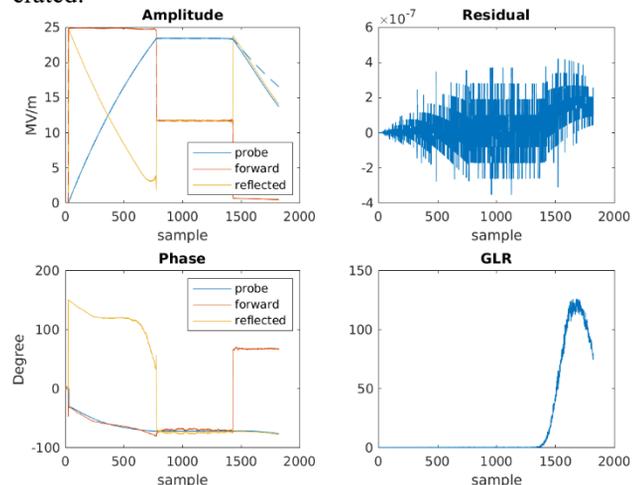


Figure 3: Graphical outcome of the GLR analysis.

On the left, the RF waveforms: forward, reflected and probe for the last nominal pulse (solid) and for the faulty pulse (dashed lines) in amplitude (top) and phase (bottom). The top right shows the time domain residual trace for that faulty pulse and the lower right plot shows the corresponding GLR. The GLR waveforms (one per cavity per RF pulse) are appended to the trip snapshot file. The outcome is summarized into a table and emailed as a daily summary to operation experts.

RESULTS

In this section, the following definition of accuracy a is used:

$$a = \frac{TP + TN}{TP + TN + FP + FN} \quad (2)$$

where TP is a true positive (i.e. the algorithm accurately detected a quench), TN is a true negative (the algorithm accurately recognized that a trip was not a quench), FP is false positive (i.e. a “fake” quench) and FN is a false negative (the algorithm failed to identify a real quench).

The QDS has been running since the first day of operation of the European XFEL (Jan. 2017) but the GLR post-mortem analysis is run daily only since September 2021. In this contribution, statistics are derived for the period Sept. 22nd 2021 to June 8th 2022, corresponding to 195 days of nominal RF operation (removing machine shutdown, startup or software development days). Over this period, 124 trips snapshots were recorded. The QDS tripped 65 times, 55 trips were actual quenches ($TP = 55$), 10 were false positives. Another 59 trips occurred but were triggered by other technical interlock mechanisms such as coupler interlock, machine protection system or high-power inhibit (i.e. klystron gun arc for example). Among those faulty pulses, three also presented a quench, not caught by the QDS because another interlock triggered first ($FN = 3$).

All these trips were also analyzed using the GLR method. All 55 quenches correctly detected by the QDS were also correctly identified by the GLR. Furthermore, among the 10 trips that “fooled” the QDS, the GLR discredited nine of these trips as “non-quench”. Amongst the remaining cases, four were identified as possible quenches: post analysis proves that one wasn’t ($FP = 1$) but three were ($FN = 3$). These cases are counted as errors, since the GLR failed to recognize with enough accuracy real and false quenches. Table 1 summarizes the accuracy and their different components for the QDS and GLR approaches.

Table 1: Accuracy of QDS and GLR

	TP	TN	FP	FN	a
QDS	55	56	10	3	89.5%
GLR	55	65	1	3	96.8%

The accuracy comparison is not entirely fair, since the GLR only runs on post-mortem data, and doesn’t carry the responsibility to interlock the RF, as is the case for the QDS. It is reasonable to assume that the accuracy of the GLR would change if this algorithm were to look at all pulses. However, one should stress that all real quenches detected

by the QDS were also correctly identified as quenches by the GLR and the GLR accurately discarded the fake quenches (nine out of ten instances). Figure 4 gives an example of a “fake” quench that triggered the QDS, but was correctly identified as a non-quench by the GLR. The cavity gradient amplitude is shown for three consecutive pulses. The first one being nominal and the subsequent two flagged as faulty. In this example, the QDS stopped the RF after two faulty pulses. The corresponding Q_L displayed in the legend show a drop from nominal value in excess of the 5×10^5 threshold, hence triggering the QDS to trip the station.

The GLR traces are also displayed (right axis) and are very distinct from the GLR traces corresponding to a real quench (Figure 2). The GLR algorithm accurately tagged these pulses as faulty, but not as quenches.

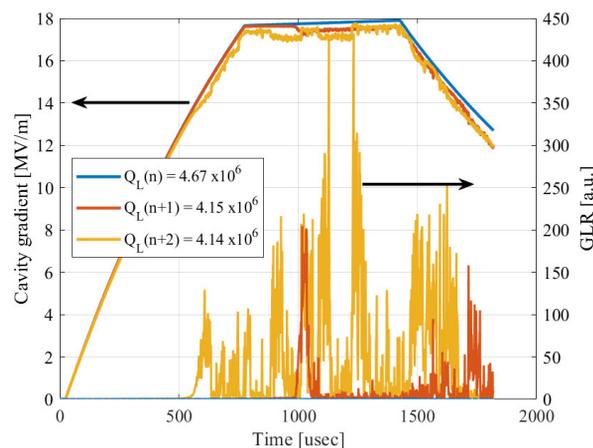


Figure 4: Gradient amplitude for three consecutive pulses: the last nominal and the next two faulty ones (left axis) and the corresponding GLR (right axis). The GLR for the nominal pulse is in the noise (not visible on this scale).

CONCLUSION AND FUTURE WORK

An overview of an approach relying on machine-learning methods to categorize trips as quench or not-quench recently implemented at the European XFEL was presented. This approach relies on existing trip snapshots to compute additional metrics (referred to as residuals) to evaluate if the quench is real. The next step consists of running this analysis on live data (as opposed to post-mortem). There are 2 options: a software- and a firmware-based approach. Running the analysis is computationally expensive so that the software approach cannot be implemented on the front-end CPUs. A solution would be to use external CPUs sharing a direct PCIe bus connection to the front-end CPU. The firmware solution is attractive because it doesn’t require additional hardware, but might be quite expensive in terms of FPGA resources. Both options are currently under evaluation. Another future work consists of looking into adapting the GLR algorithm for normal conducting cavities (such as the RF gun). The GLR approach could then help categorize different gun trips.

REFERENCES

- [1] N. Walker “XFEL linac automated trip analysis”, *Intelligent Process Control Seminar*, Oct. 6th 2020
<https://indico.desy.de/event/26983/>
- [2] J. Branlard, V. Ayvazyan, O. Hensler, H. Schlarb, Ch. Schmidt, and W. Cichalewski, “Superconducting Cavity Quench Detection and Prevention for the European XFEL”, in *Proc. ICALEPCS'13*, San Francisco, CA, USA, Oct. 2013, paper THPPC072, pp. 1239-1241.
- [3] A. Eichler, *et al.* “Anomaly Detection at the European XFEL using a Parity Space based Method”, arXiv 2202.02051
- [4] C. Tennant *et al.* “Superconducting Radio-Frequency Cavity Fault Classification using Machine Learning at Jefferson Laboratory” *Phys. Rev. Accel. Beams* 23, 114601, 2020
- [5] T. Schilcher, “Vector Sum Control of Pulsed Accelerating Fields in Lorentz force Detuned Superconducting Cavities”, Ph.D. thesis (1998)
- [6] A. Nawaz, *et al.* “Anomaly Detection for the European XFEL using a Nonlinear Parity Space Method”, in *Proc. SafeProcess 51*, 1379 (2018).
- [7] J. H. K. Timm, J. Branlard, A. Eichler, and H. Schlarb, “The Trip Event Logger for Online Fault Diagnosis at the European XFEL”, in *Proc. IPAC'21*, Campinas, Brazil, May 2021, pp. 3344-3346.
doi:10.18429/JACoW-IPAC2021-WEPAB293