

Anomaly detection at the European X-ray Free Electron Laser using a parity-space-based method

A. Eichler¹,* J. Branlard¹, and J. H. K. Timm¹

Deutsches Elektronen-Synchrotron DESY, Germany



(Received 28 February 2022; accepted 5 December 2022; published 4 January 2023)

A novel approach to detect anomalies in superconducting radio-frequency (rf) cavities is presented, based on the parity space method with the goal to detect quenches and distinguish them from other anomalies. The model-based parity space method relies on analytical redundancy and generates a residual signal computed from measurable rf waveforms. The residual is a sensitive indicator of deviation from the model and provides different signatures for different types of anomalies. This new method not only helps with detecting faults but also provides a catalog of unique signatures, based on the detected fault. The method was experimentally verified at the European X-ray Free Electron Laser (EuXFEL). Various types of anomalies incorrectly detected as quenches by the current quench detection system are analyzed using this new approach.

DOI: 10.1103/PhysRevAccelBeams.26.012801

I. INTRODUCTION

Large-scale superconducting particle accelerators such as the European X-ray Free Electron Laser (EuXFEL) comprise several hundreds of radio-frequency (SRF) cavities. A user facility of this size expects high rf availability in order to accelerate the beam and provide users with reliable and predictable photon light. Automation algorithms [1] are necessary to monitor the rf and recover as fast as possible stations where a trip took place. One typical issue causing rf downtime is a cavity quench, where an area of the cavity walls becomes normal conducting (thermal breakdown), which translates into a drop of the cavity quality factor by several order of magnitudes leading to a collapse of its accelerating field [2]. The likelihood of a quench taking place increases when cavities are operated at a high gradient, close to their quench limits (20–30 MV/m for EuXFEL cavities). For pulsed accelerators, the usual approach to quench detection consists of measuring the loaded quality factor Q_L during the decay phase of the rf pulse [3]. While this approach has proven to be robust in detecting actual quenches, it also provides false positive cases, when anomalies in machine conditions have an impact on the Q_L computation. In its current implementation, the quench detection system cannot discriminate a real from a “fake”

quench and provides by default the same reaction: switching off the rf. Events leading to such a “fake” quench are, e.g., controlled detuning events, field emitters events, and digital glitches, which are described in detail in the course of the paper in Sec. III. However, more information than what is currently used by the quench detection server is available in the control system and can be used to help analyze each rf pulse, with the intent to discriminate real quenches from other anomalies. The current quench detection server uses the cavity probe amplitude information to compute Q_L . The probe phase, the rf forward and reflected signals are also available and can be used to provide a more accurate description of the type of fault taking place. The electrical and mechanical behavior of SRF cavities is well understood and can be modeled [4]. Thus, model-based fault detection methods can be used to detect faults as it has been demonstrated in [5–8]. These model-based approaches provide a promising insight into anomaly detection and categorization. As an alternative, a fully data-based approach is also possible as presented in [9], where machine learning is applied for the identification and classification of faults in SRF cavities for the Continuous Electron Beam Accelerator Facility (CEBAF) at Jefferson Laboratory. The basis for this approach was a dataset of several thousand events, manually labeled.

In addition to cavity failures, fault and anomaly detection has gained significant interest for accelerator operation in general due to the great potential to increase availability. Mainly data-driven methods have been used, exploiting tools from machine learning. Application examples range from magnet fault detection at CERN [10] and at the Advanced Photon Source at Argonne National Laboratory [11] over the detection of faulty beam position monitors at

*Corresponding author.
annika.eichler@desy.de

Published by the American Physical Society under the terms of the *Creative Commons Attribution 4.0 International license*. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

the Large Hadron Collider [12] to fault detection at the digital electronics level [13,14].

In this paper, a novel approach to detect cavity anomalies is presented based on the parity space method [15,16] based on [5]. This method relies on the well-known cavity model since no extensive labeled dataset is available at DESY for a purely data-driven approach. The fault detection parity method makes use of analytical redundancy to generate a residual signal from measurable rf waveforms. This residual can be further analyzed with statistical tools to evaluate significant deviations from the model, i.e., a fault. The resulting measures provide different signatures depending on the type of detected anomaly. This method can detect faults but also provides a catalog of signatures based on the detected fault, thus allowing for a more robust cavity quench detection, by minimizing the number of false positives. The method was experimentally verified at EuXFEL and applied to various types of anomalies incorrectly detected as quenches by the current quench detection server. Although the method presented here is used for fault detection of SRF cavities, it is not limited to this specific example and can be applied to any part of an accelerator, where a physical model is available. The presented approach constitutes, therefore, a general tool for fault detection to improve availability and can support further automation of accelerator operation.

The next Sec. II provides general information about SRF cavities and their mathematical model. Different types of faults observed, based on 4 years of operation at EuXFEL are summarized in Sec. III. The following Sec. IV provides insight into the parity space method for fault detection and evaluation exploiting the cavity model. Experimental results are given in Sec. V.

II. SRF CAVITIES

SRF cavities are electromagnetic resonators that can be modeled as second order systems in the I (in-phase) and Q (quadrature) domain as [4]

$$\begin{bmatrix} \dot{V}_{P,I}(t) \\ \dot{V}_{P,Q}(t) \end{bmatrix} = \begin{bmatrix} -\omega_{1/2} & -\Delta\omega(t) \\ \Delta\omega(t) & -\omega_{1/2} \end{bmatrix} \begin{bmatrix} V_{P,I}(t) \\ V_{P,Q}(t) \end{bmatrix} + 2\omega_{1/2} \begin{bmatrix} V_{F,I}(t) \\ V_{F,Q}(t) \end{bmatrix} - \omega_{1/2} \begin{bmatrix} V_{B,I}(t) \\ V_{B,Q}(t) \end{bmatrix}, \quad (1)$$

where $V_F(t) = V_{F,I}(t) + jV_{F,Q}(t) \in \mathbb{C}$ is the forward field coupled into the cavity. Here, $V_{F,I}(t)$ and $V_{F,Q}(t)$ are the I and Q components, respectively. The state $V_P(t) = V_{P,I}(t) + jV_{P,Q}(t) \in \mathbb{C}$ is the probe signal, i.e., the field which builds up inside the cavity, and $V_B(t) \in \mathbb{C}$ is the field induced by the beam. Note that all signals are considered in the baseband. This means that $V_F(t)$ represents the envelope of the electromagnetic forward wave which is driven by the main oscillator with a frequency of

$f = 1.3$ GHz. Currently, the European XFEL is operated in pulsed mode with an rf pulse repetition rate of 10 Hz. Each pulse lasts for approximately 1.8 ms and can be divided into filling, flat-top, and decay. During filling, the electromagnetic forward wave is coupled into the cavity so that the probe's electromagnetic standing wave increases up to the desired field gradient. It is held constant via a feedback controller during the flat-top to accelerate the arriving electron beam by the desired energy level. This feedback controller is operating as a vector sum controller [4], where 32 cavities are controlled as one entity by considering their vector sum. Nevertheless, we can measure the individual forward signals, denoted here as $V_F(t)$. In the decay, the forward electromagnetic wave coupled into the cavity is switched off and the probe wave decays. Figure 1 shows the amplitude of the cavity signals in the different operating regions. Note that in addition to the probe and forward field, also the reflected field is shown. The reflected field $V_R(t) \in \mathbb{C}$ is the reflection of the forward field inside the cavity, so the following relationship holds:

$$V_P(t) = V_F(t) + V_R(t), \quad \forall t. \quad (2)$$

The parameters influencing the dynamic behavior of the SRF cavities are the half bandwidth $\omega_{1/2}$ and the detuning $\Delta\omega(t)$. The half bandwidth $\omega_{1/2} = \pi f / Q_L$ is defined by the loaded quality factor Q_L , which expresses the ratio of the energy stored inside the cavity to the dissipated energy. The external input power coupling to the cavity, defined by the penetration depth of the coupler antenna, is denoted as Q_{ext} . The unloaded, external, and loaded quality factors are linked by the following relationship [2]:

$$\frac{1}{Q_L} = \frac{1}{Q_0} + \frac{1}{Q_{\text{ext}}}. \quad (3)$$




FIG. 1. Amplitude of the fields for the probe $|V_P(t)|$, forward $|V_F(t)|$ and reflected $|V_R(t)|$ waveforms, and the field imposed by the beam $|V_B(t)|$ for the three regions of an rf pulse: (i) filling, (ii) flat-top, and (iii) decay.

Due to the superconductivity of the cavities at the European XFEL, the unloaded quality factor is very high ($\approx 10^{10}$). The external quality factor can be altered to tune the loaded quality factor. In nominal operation, the European XFEL is operated at a quality factor around $Q_L = 4.6 \times 10^6$, resulting in a small half bandwidth (≈ 140 Hz). This can be easily estimated during the decay region of the rf pulse as one over the time constant of the exponential decay of the probe signal.

Due to its small half bandwidth, the cavity is sensitive to the detuning $\Delta\omega(t)$, the second parameter that influences the cavity dynamics. The detuning is the difference between the driving frequency f and the cavity's resonance frequency $f_0(t)$. Ideally, this should be zero, but the resonance frequency is determined by the cavity's shape, which changes over time as the cavity is exposed to forces, e.g., mechanical forces and Lorentz forces. In principle, the detuning can be modeled as a second order mechanical resonator. However, due to the short duty cycle, an approximation as a sum of first order systems is given in [4] as

$$\begin{aligned}\Delta\dot{\omega}_n(t) &= -\frac{1}{\tau_n}\Delta\omega_n(t) + K_n[V_{P,I}^2(t) + V_{P,Q}^2(t)], \\ n &= 1, \dots, N, \\ \Delta\omega(t) &= \sum_{n=1}^N \Delta\omega_n(t).\end{aligned}\quad (4)$$

Here N defines the number of relevant mechanical modes (index n), τ_n is the time constant of the response of the respective modes, and K_n is the Lorentz force constant.

III. ANOMALIES THAT TRIGGERED THE CURRENT QUENCH DETECTION

The quench detection and reaction system deployed at EuXFEL [3] computes the loaded quality factor Q_L for all SRF cavities based on the probe waveform during the decay section of every rf pulse (iii in Fig. 1). This instantaneous Q_L is compared to a mean $\bar{Q}_L$, typically computed as a running average over the last 100 pulses. The quench detection system triggers if the Q_L computed for a single pulse drops below the running average by more than a user-defined threshold (typically 10% of the nominal Q_L value). The idea of this approach is based on the fact that a cavity quench is accompanied by a drop of the unloaded quality factor Q_0 by several orders of magnitude, which can be detected as a drop in Q_L , see Eq. (3). However, other mechanisms, physical (i.e., cavity detuning) or nonphysical (i.e., digital glitch) can also affect Q_L and fool the quench detection. In this contribution, four particular conditions or anomalies yielding false positives on the quench detection system are investigated: a controlled cavity detuning, a controlled change of cavity

bandwidth, an electron burst linked to spontaneous field emission, and finally a digital glitch corrupting the data transfer at the firmware or software level. These four cases triggered the quench detection system, although no real quench occurred.

A. Controlled change of bandwidth

Coupler heating (due to operation) can change the cavity external quality factor Q_{ext} by up to 25% [17] for TESLA cavities. When increasing or decreasing the XFEL linac energy to accommodate different user run requirements, the change in cavity forward power is accompanied by a net change in Q_{ext} due to heating effects. At EuXFEL, the input power couplers are motorized, allowing for remote cavity bandwidth adjustments. As a routine maintenance, the Q_{ext} are monitored and readjusted. While the thermal effects are very slow (several hours), a motorized coupling readjustment can produce a rapid change in Q_{ext} . This change is detected as a change in Q_L as it can be seen in Eq. (3) by the quench detection system. In some conditions, if the change is fast enough, the quench detection will falsely interpret this change as a quench. Several countermeasures are possible: one can mask the quench detection system when adjusting Q_{ext} or adjust the learning rate (by increasing the number of rf pulses over which the running average $\bar{Q}_L$ is computed) to low-pass filter these rapid controlled Q_L variations. Finally, one could slow down the motorized coupler drive or increase the complexity of the quench detection system by having it check for motor movements. Having an independent measurement that can provide additional information to help in the algorithm decision process is the approach chosen in this contribution.

B. Controlled detuning

As reported in [3], changing the cavity tuning has an influence on the measured cavity coupling. A 1-kHz change in cavity resonance using the stepper motor can infer a deviation of up to 15% in cavity bandwidth at EuXFEL. Several mechanisms can play a role in this observed coupling. Measurement errors can be attributed to the limited isolation between the forward and reflected ports of the waveguide bidirectional couplers and through waveguide coupling especially between adjacent cavities. Furthermore, the mechanical coupling between the deformation of the cavity due to the tuner stack and the penetration of the antenna at the fundamental input power coupler could also take place. Finally, digitizer nonlinearities expose limitations of the sampling scheme resulting in an asymmetrical signal detection if the detuning is above or below resonance. Further investigation of these effects is required to quantify their respective contributions but this goes beyond the scope of this work. Suffice here to say that coupling between measured detuning and Q_L is commonly observed at EuXFEL, and in some cases, a fast detuning (requested by operation for example) can fool the quench

detection system and result in a station trip. Some safety mechanisms are built in to avoid such false positives: for example, ignoring detected quenches if the cavity gradient is below a user-defined minimal threshold (2.5 MV/m typically). Heavily detuned cavities are hence ruled out, but this exception handling does not cover all cases.

C. Field emitter and other electronic processes

Spontaneous field emission and multipacting [18,19] can be observed sporadically on some cavities. During the cryomodule test phase [20], prior to XFEL tunnel installation, cavities prone to field emission were identified and flagged. In some cases, their operating gradient was limited to keep the emitted radiation below the acceptance threshold [21]. However, new field emitters can appear during operations and can be triggered through secondary emission due to the onset of another field emitter. Such events produce erratic beam loading, often observed as flickering and noisy cavity rf traces taking place at random times during the rf pulse. In some extreme cases, the produced dark current is large enough to discharge the cavity field within tens of nanoseconds. This effect also referred to as plasma discharge can affect neighboring cavities, which will see this beam loading, albeit in a reduced magnitude. If this disturbance takes place during the decay phase of the rf pulse, it will influence the computation of the Q_L value, falsely triggering the quench detection system. These spontaneous electronic processes are typically accompanied by radiation bursts, observed by gamma and neutron detectors in the tunnel. In most cases, field emitters are conditioned away after a few pulses, but not always. A simple conditioning attempt is not always successful or could result in generating new field emitters via secondary emission. Techniques such as plasma processing have been developed to cure field emissions [22] but are currently not implemented at DESY.

D. Digital glitches

Several failures in the digital domain have been observed so far, resulting in a corruption of the data transferred between electronic boards. Due to the fact that all electronics responsible for cavity field control are located inside the tunnel, they are subject to radiation showers that can produce single event upsets, flipping a bit in the digital data stream. There are countermeasures in place to track and fix such events (check-sum or cyclic redundancy checks) but multiple bit flips taking place simultaneously cannot be fixed with this approach. Another failure mode has been observed when the CPUs in the tunnel become temporarily overloaded, delaying data memory access to the time when the field-programmable gate array (FPGA) is writing to memory. These read/write collisions result in data loss, where not all data points are recorded. These two failure cases can result in discontinuities in the cavity waveforms, corrupting the Q_L computation. Such events typically

trigger the quench detection system, although no real quench occurred.

IV. FAULT DETECTION

In the following section, we present a method for fault detection, the parity space method. A model-based approach is chosen because it does not require large datasets for training. Furthermore, the standard electromagnetic model for SRF cavities [4] introduced in Sec. II, which is used in this approach, stays valid for all operational working points considered in this application, hence preserving the validity of the presented fault detection approach. With the parity space method, a residual is generated, as described in Sec. IV B, that can be further evaluated by statistical tests, see Sec. IV C.

A. Parity space method

The parity method is a method for fault detection that is based on analytical redundancy. Analytical redundancy relations are derived from an analytical model, as the cavity model in (2) and (4), and only involve measured variables [23]. These have to hold in absence of a fault. Thus, the analytical model has to represent the nominal behavior of the system. The deviation from the analytical relation is called residual, i.e., if the residual is zero, the system behavior is as described by the model, thus behaving nominally, otherwise, it is behaving faultily.

Consider the nonlinear system

$$\dot{x}(t) = f(x(t), u(t), d(t), \nu(t)), \quad (5)$$

$$y(t) = h(x(t), u(t), d(t), \nu(t)). \quad (6)$$

Here (5) is called the state equation and $x(t) \in \mathbb{R}^n$ is the state vector at time t , $u(t) \in \mathbb{R}^m$ is the control input, $d(t) \in \mathbb{R}^{m_d}$ is the disturbance vector, representing sensor noise, unmodeled dynamics, etc. The fault signals are described by $\nu(t) \in \mathbb{R}^{m_\nu}$. Equation (6) is called the output equation and $y(t) \in \mathbb{R}^n$ is the output vector. For fault detection, a residual $r(t)$ has to be described, depending on known signals only, that is zero in the absence of faults and different from zero otherwise [24], i.e.,

$$\begin{cases} \nu(t) = 0 & \Rightarrow & r(t) = 0 \\ \nu(t) \neq 0 & \Rightarrow & r(t) \neq 0 \end{cases} \quad \forall u(t), y(t), d(t). \quad (7)$$

It is clear that in a real system, the residual will never be exactly zero due to the disturbances $d(t)$, which cannot be measured. Therefore, instead of (7), the residuals need to be robust against disturbances $d(t)$ but at the same time sensitive to faults $\nu(t)$. We will tackle this in the following with a stochastic interpretation of the residuals and introduce with the generalized likelihood a measure when the residual significantly deviates from zero, thus the behavior

is not consistent with the nominal behavior, i.e., there is a fault. This will be presented in detail in Sec. IV C.

B. Residual generation

It is assumed in the following that the output $y(t)$, the input $u(t)$ as well as multiple derivations of these signals are known, but not the states $x(t)$ [and neither $d(t)$ nor $\nu(t)$]. To calculate the residuals, analytical redundant expressions of the model description are exploited. These can be generated by deriving the single equations of (5) and (6) multiple times and eliminating the unknown states to obtain residuals that are only dependent on the known signals

$$r_i = P_i(u, y, u^{(1)}, u^{(2)}, \dots, u^{(\alpha_i)}, u^{(\alpha_i)}), \quad i = 1, 2, \dots,$$

where i denotes the i th residual. With $u^{(k)}$, the k th derivative of $u(t)$ is defined. Depending on the residual i , the highest degree α_i of deviation might be different. While the elimination of the unknown states is obvious for linear systems [23], it can be very involved for nonlinear ones [24].

With the model equations of the SRF cavities (1) and (4), three (redundant) relations for the nonmeasurable detuning $\Delta\omega(t)$ can be derived, resulting in three different residuals as given in detail in Appendix A 1. As reasoned in the appendix, we will focus in this work on the following residual:

$$r(t) = \frac{-\dot{V}_{P,I}(t) + \omega_{1/2}[-V_{P,I}(t) + 2V_{F,I}(t) - V_{B,I}(t)]}{V_{P,Q}(t)} - \frac{\dot{V}_{P,Q}(t) + \omega_{1/2}[V_{P,Q}(t) - 2V_{F,Q}(t) + V_{B,Q}(t)]}{V_{P,I}(t)}, \quad (8)$$

which has been shown in [5] to be most informative. Here, the half bandwidth $\omega_{1/2}$ is either known or can be easily determined, as described in Sec. II.

The continuous physical cavity system, $y(t)$ with $t \in \mathbb{R}$, is sampled in discrete time $y(t_0 + kT)$ for $k \in \mathbb{Z}$, where T is the sampling period. This will be abbreviated as $y(k)$ in the following. In order to account for this, the residual (8) needs to be discretized. For the sampling rates in question (1 MHz or 9 MHz depending on which layer the algorithm is implemented), Euler forward discretization can be used with $\dot{y}(t) \approx \frac{y(k+1) - y(k)}{T}$. Replacing $\dot{V}_{P,I}(t)$ and $\dot{V}_{P,Q}(t)$ accordingly in (8) leads to

$$r(k) = \frac{1}{V_{P,Q}(k)T} \{-V_{P,I}(k+1) + V_{P,I}(k) + \omega_{1/2}T[-V_{P,I}(k) + 2V_{F,I}(k) - V_{B,I}(k)]\} - \frac{1}{V_{P,I}(k)T} \{V_{P,Q}(k+1) - V_{P,Q}(k) + \omega_{1/2}T[V_{P,Q}(k) - 2V_{F,Q}(k) + V_{B,Q}(k)]\}. \quad (9)$$

To improve the numerical properties of this residual, for implementation, (9) is multiplied with $V_{P,I}(k)V_{P,Q}(k)T$.

C. Residual evaluation

There are several approaches to evaluate residuals; in this work, we apply statistical tests. Here, likelihood ratios, i.e., the log-likelihood ratio, as given in (A1), are indicators of the goodness of fit of a null hypothesis H_0 versus an alternative one H_1 by the ratio of their likelihoods, assessing if the observed residuals support or significantly disagrees with the null hypothesis.

For the application at hand, we assume that the residuals follow a zero-mean Gaussian distribution with variance Σ in nominal operation, while in case of a fault or anomaly in the system, a jump in the mean value appears. With this assumption, the generalized likelihood ratio (GLR) can be derived as

$$\lambda_{\text{GLR}}(k) = \frac{K}{2} \left(\frac{1}{K} \sum_{i=k-K+1}^k r(i)^\top \right) \Sigma^{-1} \left(\frac{1}{K} \sum_{i=k-K+1}^k r(i) \right) \quad (10)$$

from the log-likelihood ratio. Details are given in Appendix A 2. Here, $r(k), \dots, r(k+K-1)$ are the observed residuals and Σ is the variance calculated from the given nominal data.

Under the assumption (A2), $2\lambda_{\text{GLR}}$ follows a $\chi^2(n)$ distribution. Here, n is the dimension of the residuals, which in our case is 1, since only the one-dimensional residual (8) is considered. With this, an error threshold $\bar{\lambda}_{\text{GLR}}$ can be chosen from the $\chi^2(1)$ distribution so that $\lambda_{\text{GLR}}(k) > \bar{\lambda}_{\text{GLR}}$ is considered to be erroneous (or anomalous), according to a chosen acceptable probability of false positives alarms given as $P(Q > 2\bar{\lambda}_{\text{GLR}})$ with $Q \sim \chi^2(1)$.

In order to fulfill the assumption (A2) for the nominal case, the residuals coming from the parity space (8) need to be zero mean. Thus, we correct them for the mean value as

$$r'(k) = r(k) - \bar{r}(k) \quad \text{with} \quad \bar{r}(k) = \frac{1}{P} \sum_{p=1}^P r_p(k), \quad (11)$$

where $\bar{r}(k)$ is the mean value of sample k over P pulses. It is clear that this would not be necessary if the model Eqs. (2) and (4) would perfectly fit and (A2) would only contain the measurement noise. But as discussed above, this is never the case in reality. Thus, repetitive disturbances or model mismatches can be corrected using Eq. (11). This approach also allows taking the beam contribution into account. As obvious in model (2), the beam is an input here. Although it can be measured, the beam information was not available in the given data sets. Thus, the beam loading is considered a repetitive disturbance, whose effect is canceled by the mean value correction in (11).

V. RESULTS

In the following section, the parity space method for fault detection is applied to real operational data, and the different cases described in Sec. III, which falsely trigger the quench detection system, are presented and analyzed.

A. Data

The data that will be analyzed in this section are provided by a snapshot recorder, triggered by any interlock occurring in the rf system. For each snapshot, the cavity signals of 250 pulses, i.e., 25 s, are saved, with 200 pulses before the event and 50 afterward. Unfortunately, the beam signal is not saved, but a good work-around to deal with the missing signal is presented in Sec. V B. Each pulse in the dataset is labeled with a unique pulse identifier, the ID, and the data also include a time vector. The data are saved as hdf5 files. One dataset contains the forward, reflected and probe signals in amplitude and phase, which can be easily transferred to the I and Q domain. Time domain signals are 1.82 ms long, sampled at 1 MHz yielding 1820 samples per waveform. Before calculating residuals, the forward and reflected signals are systematically cross calibrated in order for Eq. (2) to hold, as described in [25]. Here, the first pulses of a snapshot (which are assumed to be nominal) are used to determine the calibration coefficients.

B. Implementation

The fault detection infrastructure is implemented in C++. The analysis is automated with an easy-to-use console command. Although only results of a *posteriori* analysis are presented here, the algorithm is optimized for real-time operation in terms of calculation speed to cope with a 10-Hz rf pulse repetition rate for pulse-to-pulse detection. This run-time operation capability was demonstrated for a small number of cavities. See [26] for further implementation details.

C. Statistics

From July 8, 2020, till November 18, 2020, 34 snapshots were saved, triggered by the quench detection system. After a postmortem analysis of these trips, 18 of them have been identified as real quench. This yields a false positive rate of 47%. Thus, almost half of the events identified as quenches are no quenches during this period of operation. Within this period of 15 weeks, the European XFEL was operated at a high gradient for 5 weeks (i.e., 16.5-GeV beam energy) and 10 weeks at a low gradient (14-GeV beam energy). Two third of the real quenches occurred within the weeks of a high gradient.

Figure 2(a) shows the distribution of the events, i.e., real quenches and false positives, over the different stations (denoted as A2–A25). Station A11 is the highest performing station and was operating very close to its quench gradient during the high-gradient study period.





FIG. 2. Statistics of archived events triggered by the quench detection and reaction system.


This explains the higher quench rate. Otherwise, the quenches are relatively spread over the stations. Although the statistics is still insufficient to draw definitive conclusions, there seems to be a bias of occurrences toward the front of the linac. Possible explanations range from higher CPU loads on the front-end stations running additional algorithms, exposure to higher radiation levels due to beam losses in the proximity of the beam chicanes, and more frequent energy tuning on rf stations located at the beginning of the tunnel. There is ongoing research to further quantify and understand the distribution of trips along the linac. The duration of downtime caused by the respective event is shown in Fig. 2(b). Here one real quench with a downtime of more than an hour is missing. Due to this outlier, the mean downtime for the real quenches is almost 7 min, while it is 2 min for the false positives. The median, however, is very similar in both cases around 100 s. Whether real or false, a detected quench stops the rf, the beam, and photon delivery for the entire machine, interrupting user experiments. Minimizing the occurrence of such events is paramount. Thanks to a high level of automation, a quench trip is typically recovered in less than 2 min. In such cases, the resulting downtime is short enough not to cause any detrimental drifts in the thermal equilibrium of the test setup at the experimental site. Once the station is recovered, photon delivery is typically recovered with a similar intensity as before the trip allowing for the experiment to carry on.

D. Residual generation and evaluation

The residuals are calculated as given in (9), scaled with $V_{P,I}(k)V_{P,Q}(k)T$ for numerical issues as discussed in Sec. IV B. Figure 3 shows the residuals of nominal pulses. The averaged residuals without beam show that the model fits very well and the residuals are mean value free. Since the presence of a beam is not yet implemented as an input to the model and not available as measurements in the data,



(a) Without beam



(b) With beam

FIG. 3. Nominal residuals without and with the beam. In blue, the residuals of one pulse are shown, in red, the residuals are averaged over 50 pulses. The presence of the beam is clearly visible as a step up corresponding to the duration of the bunch train, between 0.8 and 0.9 ms. The spikes around 0.8 and 1.4 ms mostly visible on the residual trace without a beam are actually due to imperfections in the time alignment between the different rf waveforms and are not beam dependant.

the presence of an accelerated beam is clearly visible as a model mismatch as shown in Fig. 3. This can easily be corrected using (11), as long as the beam signal stays constant over the short snapshot dataset, which was the case for all datasets analyzed so far. If the beam changes over the dataset, the correction does not hold, and this could result in a false positive. The final solution consists of adding the beam signal as an additional input to the model, which is the subject of current work. Furthermore, it is obvious that with the residual calculation, we are in the resolution range of the considered signals as one clearly sees quantization effects. This is confirmed in the histogram in Fig. 4(a). As the Q-Q plot in Fig. 4(b) appears to be relatively linear, there might be an underlying Gaussian distribution, however, this is clearly distorted by quantization. With this, the assumption in (A2) does not hold. Nevertheless, the GLR can be used as a metric, as it can be interpreted as a weighted error norm and the experimental distribution of nominal data can be considered to choose a reasonable threshold. The upper bound, $\bar{\lambda}_{\text{GLR}} = 10.8$, is chosen




FIG. 4. Probability plots of residuals without beam. In red, a fitted normal distribution is shown.

empirically, such that nominal data do not trigger false-positive alarms and at the same time, the quench datasets generate no false-negative alarms. Since much more nominal data are available (compared to quench datasets), the bound is pushed toward the lower side to reduce the false-negative rate in general. This bound would have led to a false-positive rate of 0.0003% if the λ_{GLR} followed the χ^2 distribution.

E. Analysis of the generalized likelihood ratio

For all cases visualized and analyzed in the following section, the corresponding station, module, and cavity as well as the unique ID are given in Table I in the appendix for documentation and transparency.

1. Quenches

Figure 5 shows the cavity signals of a quench event together with the GLR for the first pulse, where the threshold of $\bar{\lambda}_{\text{GLR}} = 10.8$ is exceeded. The GLR for this pulse and the three consecutive ones is shown in Fig. 6. The rf was switched off on the following pulse by the quench detection




FIG. 5. Quench signals for pulse 1 from Fig. 6.




FIG. 6. GLR values of quenches for the four consecutive pulses after exceeding the threshold.

system. As shown here, there can be a delay of multiple pulses between quench detection and reaction, due to the current software implementation of the quench detection. Thus, it is unclear at what pulse the quench detection system has actually been triggered. All signals of the GLR follow a characteristic bell curve, initiating during the decay of the first pulse and occurring at earlier times for each successive pulse. One quench event is selected here for demonstration but all of them show very similar GLR traces.

2. Field emitter and other electronic processes

The rf signals observed during a plasma discharge event and the corresponding GLR for cavity 4 of cryomodule 3 (C4.M3) are shown in Fig. 7, where the accelerating field is discharged within tens of nanoseconds. This strong anomaly is confirmed by an extremely high GLR value. For this event, neighboring cavities have also seen this




FIG. 7. Field emitter signals for cavity C4.M3.




FIG. 8. Field emitter signals for cavity C1.M3, for the same pulse as in Fig. 7.

beam loading effect, albeit in a reduced magnitude. Figure 8 shows the discharge of a fraction of the probe field (see Box 1) for a neighboring cavity within the same cryomodule (C1.M3). The propagating effect downstream, but also upstream is further illustrated in Fig. 9 where the maximal GLR value over the whole pulse is plotted for all cavities within the station. The cavity where the event took place (C4.M3) is clearly identified, its GLR being 1 order of magnitude higher than all others. One can further see the slow decay of the effect on downstream cavities (C5–C8 in M3 and C1–C4 in M4) and upstream (C3–C1 in M3 and C8 in M2), where the threshold of $\bar{\lambda}_{\text{GLR}} = 10.8$, marked in red, is exceeded. Most of the energy is scattered at the quadrupole located at the end of each cryomodule, due to the mismatch between the quadrupole setting and the energy of the accelerated dark current.

The GLR signature is clearly distinct from that observed in the case of quenches. It can be further noticed, for example, in Fig. 8 that very small effects in the cavity signals get strongly emphasized by the GLR.




FIG. 9. Maximum GLR signal over the pulse as in Figs. 7 and 8 of all 32 cavities for the plasma discharge event observed on C4.M3. The threshold of $\bar{\lambda}_{\text{GLR}} = 10.8$ is marked in red, dashed-dotted.






FIG. 10. Glitch signals.

3. Glitch


The cavity signals for a digital glitch together with the GLR are shown in Fig. 10. Here the signals of one cavity are shown as an example, but all cavities within the module show the same distortions. It is obvious that this cannot be a physical fault since forward and reflected signals are shifted in time during the flattop region, one compared to the other (see Box 1), while the probe and reflected waveforms differ during the decay (see Box 2). This is not physically possible, as the probe equals the sum of the forward and reflected signal, and the forward signal is zero during the decay.

4. Q_L Change

To study the effect of a controlled change in Q_L , the motorized antenna of an input power coupler was deliberately moved over several pulses. This change in Q_{ext} directly impacts the measured Q_L , as shown in Fig. 11. To better isolate the effect of the Q_{ext} change and remove any possible detuning contribution resulting from adjusting Q_{ext} , the automatic frequency tuning (piezo-based Lorentz force detuning compensation) is kept on during this experiment. Figure 12 shows the GLRs corresponding to the eight successive pulses marked by red crosses in Fig. 11. While the first one looks normal (numerical noise), a clear

FIG. 11. Loaded quality factor Q_L .FIG. 12. GLR values for changed Q_L for pulses 10–17 corresponding to Fig. 11. The threshold of $\bar{\lambda}_{\text{GLR}} = 10.8$ is marked in red, dashed-dotted.

signature is visible starting from the second pulse, remaining relatively flat in the filling, followed by a first step up when the flattop begins and a second one at the beginning of the decay followed by a linear decay. The amplitude of this specific signature increases as Q_{ext} is further reduced. Two specific pulses (the second and the last one of those depicted in Fig. 12) are shown in detail in Figs. 13 and 14. While the threshold is not hit for the second pulse in Fig. 13, one clearly sees the distinct signature. The threshold is hit two pulses later. After the eighth pulse, shown in detail in Fig. 14, the quench detection system has switched off the cavity. Note that the signals have been calibrated for the nominal pulses (i.e., before the fault). As obvious in both figures, more prominently in Fig. 14, the calibration no longer satisfies Eq. (2) when Q_{ext} is changed; this is clearly visible during the decay where the probe and reflected traces

FIG. 13. Cavity rf signals and corresponding GLR for pulse 11 from Fig. 12 when Q_L is changed. The threshold of $\bar{\lambda}_{\text{GLR}} = 10.8$ is not hit.




FIG. 14. All signals for changed Q_L for pulse 17 from Fig. 12. The threshold of $\bar{\lambda}_{\text{GLR}} = 10.8$ is marked in red, dashed-dotted.

no longer lay on top of each other, while the forward signal is zero. This calibration error leads to a mismatch of the model (2) and thus to an increase in the GLR.

5. Detuning

Figures 15 and 16 show an event, where the quench detection system is fooled by a fast detuning change. The first change in cavity detuning is observed in Fig. 15 (positive tilt of the reflected signal during the flattop region). Larger detuning for the same cavity is evidenced in Fig. 16, characterized by the round drop of the cavity probe gradient and steeper increase of the reflected signal during the flattop. Also here the GLR shows a special signature, distinct from the ones observed before: a linear increase during the filling and the flattop while it is almost zero during the decay. It is interesting to note that




FIG. 15. Cavity signals and corresponding GLR when the detuning is changed.




FIG. 16. All signals for changed detuning, seven pulses after the one shown in Fig. 15.

the GLR stays below the threshold despite the heavy detuning. This is to be expected, as the residual in (8) is the difference between detuning values resulting from two differential equations of the same electromagnetic model (1). If we assume that the model holds, a common mode detuning change will affect both equations identically, hence yielding a zero residual. The reason that we see distinct traces can be explained by imperfections of the model and signal imperfections due to coupling but they do not significantly affect fault detection. These signal imperfections are given by the calibration, which does not take the coupling described in Sec. III B between detuning and Q_L into account. It is clearly visible in Figs. 15 and 16 that the calibration is affected by this coupling as the forward field is not equal to zero in the decay and the probe and reflected traces do not perfectly coincide. Note that the GLR is robust against this mismatch condition since it does not lead to an increase in the GLR value.

VI. CONCLUSION

This work presents the parity space method for fault detection in SRF cavities followed by statistical evaluation using the generalized likelihood ratio. It validates the good performance of the method in detecting faults and also shows that different fault types result in clearly distinct signatures of the generalized likelihood ratio. The underlying physical root cause of different faults can follow different time constants, which translates into variations of the GLR in the time domain. The distinction of different fault types could also be achieved by a thorough analysis of all six cavity signals (I and Q values of the probe, forward and beam signals). However, with the generalized likelihood ratio, all information is embedded into one single signal. This is of practical use for online operation, as the operator only needs to observe one signal but also facilitates the *a posteriori* classification of faults. Furthermore,

this opens the door toward an automatic classification of faults using classification tools from machine learning, exploiting the very distinct GLR signatures for the different failure types demonstrated here. This topic is subject to future work. Such an online automatic classification would be of great support for linac's operation.

It could be demonstrated that, due to the choice of residuals, the method presented here is robust against changes in cavity detuning. Thus, in contrast to the current quench detection system, the proposed approach would not result in false positives of this kind. Detecting such abrupt detuning changes could be realized using additional residuals. One possibility would be to solve the electromagnetic model for the half bandwidth. It is expected that this should be independent of a change of Q_L but detect a change in the detuning.

Experimental results of the implementation at the European XFEL are currently gathered for a statistical analysis over a longer time period. In particular, this could help quantify how well the new approach performs compared to the previous one. The outcome will be presented in a future contribution.

In this work, the cavity faults have been analyzed postmortem. However, the code for fault detection and evaluation is able to cope with real-time requirements and can be switched to support offline and online analysis. Running the analysis on *live* data would increase statistics, possibly catching events otherwise left unnoticed and in the long run, support operations with more accurate diagnostics and point to necessary preventive maintenance or counteractions. The online analysis has already been demonstrated for a small number of cavities [26]. It will be subject to future work to bring this to normal operation for all cavities.

ACKNOWLEDGMENTS

The authors wish to thank Sven Pfeiffer, Lutz Lilje, and Denis Kostin for fruitful discussions about the data analysis during the writing of this paper. This work was supported by the European X-ray Free Electron Laser research and development program.

APPENDIX

1. Residual generation for SRF cavities

Given the model of SRF cavities (1) and (4), three (redundant) relations for the nonmeasurable detuning $\Delta\omega(t)$ can be derived as

$$\begin{aligned}\Delta\omega_I(t) &= \frac{-\dot{V}_{P,I}(t) + \omega_{1/2}(-V_{P,I}(t) + 2V_{F,I}(t) - V_{B,I}(t))}{V_{P,Q}(t)}, \\ \Delta\omega_{II}(t) &= \frac{\dot{V}_{P,Q}(t) + \omega_{1/2}(V_{P,Q}(t) - 2V_{F,Q}(t) + V_{B,Q}(t))}{V_{P,I}(t)}, \\ \Delta\dot{\omega}_{III}(t) &= \sum_{n=1}^N \Delta\dot{\omega}_n(t).\end{aligned}$$

While $\Delta\omega_I(t)$ and $\Delta\omega_{II}(t)$ are already functions of known signals only, $\Delta\omega_{III}(t)$ is a function of the nonmeasurable states $\Delta\omega_n(t)$ for $n = 1, \dots, N$. As given in [23], standard procedures exist to eliminate the nonmeasurable states due to the linearity of the relationship. Then, three different residuals can be derived, i.e., $\Delta\omega_I(t) - \Delta\dot{\omega}_{III}(t)$, $\Delta\dot{\omega}_I(t) - \Delta\omega_{III}(t)$, and $\Delta\dot{\omega}_{II}(t) - \Delta\dot{\omega}_{III}(t)$. As the first has been shown in [5] to be most informative, we will concentrate on

$$\begin{aligned}r(t) &= \Delta\omega_I(t) - \Delta\omega_{II}(t) \\ &= \frac{-\dot{V}_{P,I}(t) + \omega_{1/2}(-V_{P,I}(t) + 2V_{F,I}(t) - V_{B,I}(t))}{V_{P,Q}(t)} \\ &\quad - \frac{\dot{V}_{P,Q}(t) + \omega_{1/2}(V_{P,Q}(t) - 2V_{F,Q}(t) + V_{B,Q}(t))}{V_{P,I}(t)}.\end{aligned}$$

Further reasons for this choice are that (i) for this residual, no additional derivation of the data is required, which amplifies noise, and (ii) $\Delta\omega_I(t)$ and $\Delta\omega_{II}(t)$ only depend on $\omega_{1/2}$, which is either known or can be easily determined, as described in Sec. II, while $\Delta\omega_{III}(t)$ depends on τ_n and K_n for $n = 1, \dots, N$, for which an additional system identification step is necessary.

2. Generalized likelihood ratio

Assuming that the observed residuals $r(k), \dots, r(k + K - 1)$ follow a statistical model with parameter μ , with $\mu = \mu_0$ being the null hypothesis and $\mu = \mu_1$ being the alternative one, the log-likelihood ratio given the observed residuals is defined as [27]

$$\lambda_{\text{LR}}(k) = \sum_{i=k-K+1}^k \ln \frac{p(\mu_1|r(i))}{p(\mu_0|r(i))}. \quad (\text{A1})$$

Here $p(\mu_1|r(k))$ is the probability of the alternative hypothesis H_1 given the observed residual $r(k)$; the complementary probability $p(\mu_0|r(k))$ is defined accordingly. For the application at hand, we assume in case of a fault or anomaly, i.e., in the case of the alternative hypothesis, a jump in the mean value of a Gaussian distribution as

$$\begin{aligned}r(k) &= \mu + \mathcal{N}(0, \Sigma), \quad \text{with} \\ \theta &= \begin{cases} \mu_0 = 0, & H_0(\text{no change}), \\ \mu_1 \neq 0, & H_1(\text{change}). \end{cases} \quad (\text{A2})\end{aligned}$$

With this in nominal operation, we expect that the residuals follow a zero-mean Gaussian distribution with variance Σ , while in case of a fault or anomaly in the system, a jump in the mean value appears. So that the data are still Gaussian distributed with the same variance but the mean value is different. While the variance Σ can be calculated

TABLE I. Documentation of location and time of all cases analyzed in the paper.

Figure	Cavity.Module.Station	ID
Figure 5	C3.M1.A11	794495576
Figure 6	C3.M1.A11	794495576–794495579
Figure 7	C4.M3.A23	1209385508
Figure 8	C1.M3.A23	1209385508
Figure 9	C1-8.M1-4.A23	1209385508
Figure 10	C1.M4.A10	985812832
Figure 11	C1.M1.A13	982283119–982283135
Figure 12	C1.M1.A13	982283128–982283135
Figure 13	C1.M1.A13	982283129
Figure 14	C1.M1.A13	982283135
Figure 15	C2.M1.A2	764126695
Figure 16	C2.M1.A2	764126702

from the given nominal data, the mean μ_1 is unknown. For estimating μ_1 , the maximum log-likelihood ratio, i.e., generalized likelihood ratio (GLR), is considered. This is derived in (10) from (A1) by replacing μ_1 with its maximum likelihood estimate, given as [27]

$$\lambda_{\text{GLR}}(k) = \max_{\mu_1} \lambda_{\text{LR}}(k) \\ = \frac{K}{2} \left(\frac{1}{K} \sum_{i=k-K+1}^k r(i)^{\top} \right) \Sigma^{-1} \left(\frac{1}{K} \sum_{i=k-K+1}^k r(i) \right).$$

[1] J. Branlard, S. Choroba, M. Grecki, S. Köpke, D. Kostin, D. Nölle, V. Vogel, N. Walker, and S. Wiesenberger, Four years of successful operation of the European XFEL, in *Proceedings of 2021 International Conference on Superconductivity* (2021), <https://accelconf.web.cern.ch/srf2021/papers/moofav06.pdf>.

[2] H. Padamsee, J. Knobloch, and T. Hays, *RF Superconductivity for Accelerators*, Wiley Series in Beam Physics and Accelerator Technology (Wiley, New York, 1998).

[3] J. Branlard, V. Ayvazyan, O. Hensler, H. Schlarb, C. Schmidt, and W. Cichalewski, Superconducting cavity quench detection and prevention for the European XFEL, in *Proceedings of the 14th International Conference on Accelerator and Large Experimental Physics Control Systems, ICALEPCS'13, San Francisco, CA* (JACoW, Geneva, Switzerland, 2013), pp. 1239–1241.

[4] T. Schilcher, Vector sum control of pulsed accelerating fields in Lorentz force detuned superconducting cavities, Ph.D. thesis, Universität Hamburg, 1998.

[5] A. Nawaz, S. Pfeiffer, G. Lichtenberg, and P. Rostalski, Anomaly detection for the European XFEL using a non-linear parity space method, in *Proceedings of 10th IFAC Symposium on Fault Detection, Supervision and Safety for Technical Processes, SafeProcess 2018, Warsaw, Poland* (2018), Vol. 51, p. 1379, [10.1016/j.ifacol.2018.09.554](https://doi.org/10.1016/j.ifacol.2018.09.554).

[6] A. S. Nawaz, C. Herzog, J. Grasshoff, S. Pfeiffer, G. Lichtenberg, and P. Rostalski, Probabilistic model-based fault diagnosis for the cavities of the European XFEL, *at - Automatisierungstechnik* **69**, 538 (2021).

[7] W. H. Syed, A. Eichler, A. Nawaz, B. Sharan, and H. Werner, Koopman-based Kalman filter for fault detection for the superconducting radio frequency cavities of the European XFEL, in *Proceedings of 2021 60th IEEE Conference on Decision and Control (CDC), Austin, TX* (IEEE, New York, 2021), pp. 6855–6860.

[8] A. Bellandi, Ł. Butkowski, B. Dursun, A. Eichler, Ç. Gümüüş, M. Kuntzsch, A. Nawaz, S. Pfeiffer, H. Schlarb, C. Schmidt, K. Zenker, and J. Branlard, Online detuning computation and quench detection for superconducting resonators, *IEEE Trans. Nucl. Sci.* **68**, 385 (2021).

[9] C. Tennant, A. Carpenter, T. Powers, A. Shabalina Solopova, L. Vidyaratne, and K. Iftekharuddin, Superconducting radio-frequency cavity fault classification using machine learning at Jefferson Laboratory, *Phys. Rev. Accel. Beams* **23**, 114601 (2020).

[10] M. Wielgosz, A. Skoczni, and M. Mertik, Using LSTM recurrent neural networks for monitoring the LHC superconducting magnets, *Nucl. Instrum. Methods Phys. Res., Sect. A* **867**, 40 (2017).

[11] J. P. Edelen and N. M. Cook, Anomaly detection in particle accelerators using Autoencoders, in *Proceedings of 2021 Improving Scientific Software Conference* (2021), pp. 5–11, [10.26024/p6mv-en77](https://doi.org/10.26024/p6mv-en77).

[12] E. Fol, J. M. Coello de Portugal, and R. Tomás, Un-supervised machine learning for detection of faulty beam position monitors, in *Proceedings of 10th International Particle Accelerator Conference, IPAC-2019, Melbourne, Australia* (JACoW, Geneva, Switzerland, 2019), pp. 2668–2671.

[13] A. Grünhagen, J. Branlard, A. Eichler, G. Martino, G. Fey, and M. Tropmann-Frick, Fault analysis of the beam acceleration control system at the European XFEL using data mining, in *Proceedings of 2021 IEEE 30th Asian Test Symposium (ATS), Matsuyama, Ehime, Japan* (2021), pp. 61–66, [10.1109/ATS52891.2021.00023](https://doi.org/10.1109/ATS52891.2021.00023).

[14] G. Martino, A. Grünhagen, J. Branlard, A. Eichler, G. Fey, and H. Schlarb, Comparative evaluation of semi-supervised anomaly detection algorithms on high-integrity digital systems, in *Proceedings of 2021 24th Euromicro Conference on Digital System Design (DSD), Palermo, Italy* (2021), pp. 123–130, [10.1109/DSD53832.2021.00028](https://doi.org/10.1109/DSD53832.2021.00028).

[15] R. J. Patton and J. Chen, A review of parity space approaches to fault diagnosis, *IFAC Proceedings Series* **24**, 65 (1991).

[16] J. J. Gertler, Parity equation implementation of residual generators, in *Fault Detection and Diagnosis in Engineering Systems* (CRC Press, Boca Raton, FL, 1998), 1st ed., Chap. 6.

[17] W. Cichalewski, J. Sekutowicz, A. Napieralski, R. Rybaniec, J. Branlard, and V. Ayvazyan, Continuous wave operation of superconducting accelerating cavities with high loaded quality factor, *IEEE Trans. Nucl. Sci.* **67**, 2119 (2020).

[18] P. T. Farnsworth, Television by electron image scanning, *J. Franklin Inst.* **218**, 411 (1934).

- [19] J. Knobloch, W. Hartung, and H. Padamsee, Multipacting in 1.5-GHz superconducting niobium cavities of the CEBAF shape, in *Proceedings of 8th International Conference on RF Superconductivity (SRF'97), Abano Terme, Padua, Italy* (JACoW, Geneva, 1997), pp. 1017–1027.
- [20] J. Branlard *et al.*, LLRF Tests of XFEL Cryomodules at AMTF: First Experimental Results, in *Proceedings of 16th International Conference on RF Superconductivity (SRF'13), Paris, France* (JACoW Publishing, 2013), pp. 1132–1134.
- [21] D. Reschke, V. Gubarev, J. Schaffran, L. Steder, N. Walker, M. Wenskat, and L. Monaco, Performance in the vertical test of the 832 nine-cell 1.3 GHz cavities for the European X-ray Free Electron Laser, *Phys. Rev. Accel. Beams* **20**, 042004 (2017).
- [22] J. Knobloch and H. Padamsee, Explosive field emitter processing in superconducting RF cavities, *Part. Accel.* **61**, 169 (1998), <https://accelconf.web.cern.ch/SRF97/papers/srf97d22.pdf>.
- [23] M. Kinnaert, Fault diagnosis based on analytical models for linear and nonlinear systems—a tutorial, *IFAC Proceedings Series* **36**, 37 (2003).
- [24] J. Bokor and Z. Szabó, Fault detection and isolation in nonlinear systems, *Annu. Rev. Control* **33**, 113 (2009).
- [25] S. Pfeiffer, V. Ayvazyan, J. Branlard, L. Butkowski, R. Rybaniec, H. Schlarb, C. Schmidt, and R. Rybaniec, Virtual cavity probe generation using calibrated forward and reflected signals, in *Proceedings of 6th International Particle Accelerator Conference, IPAC-2015, Richmond, VA* (JACoW, Geneva, Switzerland, 2015).
- [26] J. H. K. Timm, J. Branlard, A. Eichler, and H. Schlarb, The Trip Event Logger for Online Fault Diagnosis at the European XFEL, in *Proceedings of the 12th International Particle Accelerator Conference, IPAC-2021, Campinas, SP, Brazil* (JACoW, Geneva, Switzerland, 2021), pp. 3344–3346.
- [27] S. X. Ding, *Model-based Fault Diagnosis Techniques* (Springer, Berlin, Heidelberg, 2008).