

Noise Analysis and Variance Reduction of Continuous Impedance Measurements of Energized Superconducting Magnets

Magnus Bøgh Borregaard Christensen ^{1b}, Peter Koch ^{1b}, Tomasz Podzorny ^{1b},
and Jan Østergaard ^{1b}, *Senior Member, IEEE*

Abstract—Superconducting magnet impedance measurements are vital for assessing magnet health and electrical integrity. To further enhance monitoring capabilities beyond contemporary methods that largely rely on manual intervention, recent efforts have focused on enabling in situ and continuous measurements during magnet operation. This evolution is becoming increasingly relevant due to the growing complexity and aging of modern particle accelerator facilities. However, implementing such measurements presents challenges, particularly due to operational constraints and interference from the power converter. This article focuses on noise reduction techniques aimed at reducing the variance of impedance estimates derived from samples collected by a differential probing measurement system. A key contribution is the analysis of a unique, high-resolution dataset comprising uninterrupted impedance measurements across both steady-state magnet current plateaus and current ramping stages. This dataset enables inspection of the interaction between injected stimuli and power converter noise throughout key stages of a magnet’s powering cycle, an aspect not previously explored and reported in the literature. Using a differential measurement configuration, we extract a reference of the power converter noise and apply Wiener filtering to reduce the variance of impedance estimates. We evaluate two denoising strategies, a static approach with fixed filter coefficients and an adaptive method with periodically updated coefficients. For long estimation windows (1 s), neither approach yields significant improvements. However, for short windows (10 ms), both methods achieve substantial variance reduction of up to two orders of magnitude. Under certain operating conditions, the adaptive method provides a further improvement of approximately one order of magnitude over the static approach, highlighting the potential advantage of adaptivity for real-time impedance monitoring.

Index Terms—Accelerator magnets, impedance, measurement, testing.

I. INTRODUCTION

IN LARGE-SCALE particle accelerators, such as the Large Hadron Collider (LHC) at CERN, the continuous monitoring

Received 20 July 2025; revised 13 November 2025; accepted 19 December 2025. Date of publication 24 December 2025; date of current version 15 January 2026. This work was supported by High Field Magnet R&D (HFM). (Corresponding author: Magnus Bøgh Borregaard Christensen.)

Magnus Bøgh Borregaard Christensen is with CERN, 1211 Geneva, Switzerland, and also with Aalborg University, 9220 Aalborg, Denmark (e-mail: magnus.christensen@cern.ch).

Peter Koch and Jan Østergaard are with Aalborg University, 9220 Aalborg, Denmark.

Tomasz Podzorny is with CERN, 1211 Geneva, Switzerland.

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TASC.2025.3648278>.

Digital Object Identifier 10.1109/TASC.2025.3648278

of superconducting magnet impedance is becoming an increasingly critical area of study as existing installations age. The LHC was first commissioned in 2008, and with much of its original infrastructure being retained in the high-luminosity LHC [1] (HL-LHC) upgrade, currently projected to operate into the 2040s, maintenance and diagnostics will become increasingly crucial in the next two decades. In this context, recently developed tools have enabled the measurement of superconducting magnet impedance during operation [2], [3], [4], and evidence suggests that such measurements may offer elevated diagnostic insight. For example, it has been shown that impedance variations can be used for detecting interturn short circuits, where, in particular, the resonance peak undergoes a pronounced transition [5], [6], [7].

A previously designed impedance measurement system [4] operates at low voltage to avoid interference with peripheral equipment. This limitation renders the impedance estimates susceptible to noise, particularly from the magnets’ connected power converter. This article investigates a Wiener filter-based signal processing approach to reduce estimation variance by incorporating knowledge of the power converter noise. To the best of the authors’ knowledge, the overall approach of measuring the impedance of an actively powered superconducting magnet circuit and applying denoising represents a novel contribution to the instrumentation of superconductors.

Ideally, power converter noise should be modeled for each specific measurement scenario and incorporated into the impedance estimation method, including variations in power converter models, connected loads, and operating scenarios, such as current ramp rates. However, this approach is impractical, due to the wide variety of power converter technologies [8], each exhibiting distinct noise spectra, rendering a comprehensive general model unobtainable. A key aspect influencing the noise behavior is the control scheme: most power converters rely on control loops that regulate a voltage source to achieve and maintain the target current [8]. This control affects the noise characteristics, which vary depending on the magnet’s powering profile—typically involving a linear ramp-up or ramp-down of current followed by a constant current, onward referred to as a plateau [9], [10]. As the control loop responds to different demands during these stages, the noise properties change. This variability complicates open-loop modeling. To address this




Fig. 1. Voltage measurements (blue) recorded across the terminals of a prototype of an HL-LHC quadrupole superconducting magnet (MQXF57j) during a powering cycle at CERN’s SM18 test facility. The magnet is powered by a 20 kA thyristor-based power converter driving a current ramp (orange) with three constant current (plateau) stages.

challenge, the noise must be measured directly and filtered in accordance with the system’s dynamic behavior. Fig. 1 illustrates an example of a current ramp profile and the associated voltage measured across the power converter, showing relatively stable noise levels during plateau stages and more visibly fluctuating behavior during current ramps.

In this study, we apply a finite impulse response (FIR) filter designed using the Wiener method [11] to reduce noise in the impedance measurements. The process involves directly measuring the power converter noise, filtering it with designed FIR filters, shaping the spectrum of the noise to match that present in the voltage and current measurements of the injected stimulus, and then subtracting this filtered noise from the corresponding voltage and current measurements—aiming to improve the impedance estimates by reducing the estimation variance. We implement both a time invariant version of the Wiener filter, using fixed coefficients over the entire data, and a blockwise adaptive version that updates the filter as the system evolves. This allows us to evaluate whether adaptivity provides a measurable benefit under operating conditions, or whether a simpler static implementation suffices.

To assess the performance of the method across different temporal resolutions, we examine two window lengths representing different functional properties: 1 s and 10 ms. The longer 1-s window is expected to provide superior signal-to-noise ratio (SNR) and more stable impedance estimates, making it suitable for identifying long-term trends and gradual drift in the impedance. In contrast, the shorter 10-ms window is used to evaluate to what extent the method can resolve rapid impedance changes that may serve as indicators of transient events, such as quenches.

In addition to the methodological focus on variance reduction of impedance estimates, the central contribution of this work lies in the experimental dataset itself. Unlike previous studies that either presented data only at fixed current levels [4] or lacked accurate measurements of the circuit current [3], the data analyzed here are acquired continuously throughout both

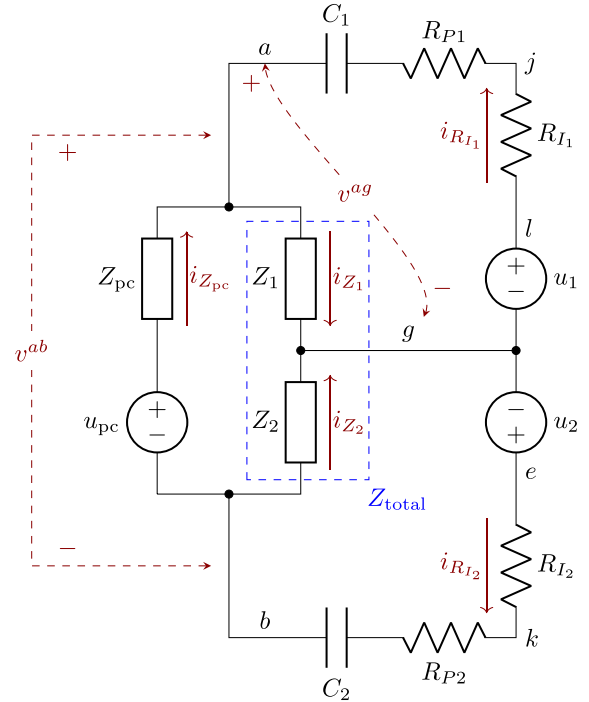


Fig. 2. Adapted from [4]. Schematic of the differential probing circuit, featuring two stimulus sources applied to a magnet divided into two partitions. The impedances of the magnet partitions are Z_1 and Z_2 , the power converter voltage is u_{pc} , and its source impedance is Z_{pc} . The stimulus voltages u_1 and u_2 are each in series with reference current resistors R_{I1} and R_{I2} . Two capacitors, C_1 and C_2 , along with two resistors, R_{P1} and R_{P2} , protect the stimulus injection electronics from the magnet circuit in case of a quench. Key currents ($i_{R_{I1}}$, $i_{R_{I2}}$, i_{Z_1} , i_{Z_2} , and $i_{Z_{pc}}$) and voltages (v^{ag} and v^{ab}) are highlighted.

plateau and ramping stages. Crucially, the instantaneous circuit current is measured simultaneously with the voltage and current responses to the injected stimuli, thus enabling comparison between impedance signals and the magnet’s operating conditions.

The rest of this article is organized as follows. Section II provides a self-contained review of the necessary background, including an overview of the signal model, the notations used throughout the study, and the processing steps employed to recover the magnet impedance from the measured signals. In Section III, we present the experimental data collected specifically for this investigation, highlighting the challenges posed by noise and motivating the need for effective noise reduction. Section IV introduces the Wiener approach for denoising, followed by Section V, where we present the results obtained from applying this technique to the experimental data. Section VI discusses directions for future research. Finally, Section VII concludes this article.

II. BACKGROUND

A. Circuit Model, Assumptions, and Notations

A powered magnet can be probed using the circuit topology illustrated in Fig. 2 [3], [4].

The circuit consists of a magnet with total impedance Z_{total} , partitioned into two sections, Z_1 and Z_2 , such that $Z_{total} = Z_1 + Z_2$. The magnet is powered by a voltage source u_{pc}

with an associated source impedance Z_{pc} . A stimulus signal is introduced via voltage sources u_1 and u_2 , each connected in series with known reference current resistors R_{I_1} and R_{I_2} . The stimulus signal determines the impedance spectrum being excited. In addition, the circuit also shows two capacitors and additional resistors, between the stimulus output stages and the magnet connections. These are to protect the magnet and injection electronic in case of high voltage being generated by a quench.

To indicate variables specific to a particular pair of nodes, superscripts are used. For example, the voltage between nodes a and g is written as v^{ag} . If no superscript is used, it implies the expression applies for any given node pair.

This article employs a discrete-time modeling framework. Although the underlying physical system is inherently continuous, our analysis assumes sampled data. Such discrete-time modeling is valid, provided all stimulus and response signals are sufficiently band-limited and sampled at an adequately high rate, preventing aliasing and preserving system dynamics [12]. Specifically, to remain in compliance with the Nyquist sampling theorem, we define a sampling frequency f_s , to exceed at least twice the highest frequency component of the stimulus.

Discrete-time domain variables are denoted by lowercase letters with index n , indicating sample number. Their corresponding frequency-domain representations are denoted by uppercase letters

$$u[n], v[n], i[n] \in \mathbb{R} \xleftrightarrow{\mathcal{F}} U(f), V(f), I(f) \in \mathbb{C} \quad (1)$$

where frequency-domain functions are defined only on the domain $f \in [-\frac{f_s}{2}, \frac{f_s}{2}]$.

In the following sections, time-domain and frequency-domain descriptions are used interchangeably depending on context. For improved readability, the discrete-time index $[n]$ and continuous functional argument (f) are often omitted.

To enable impedance estimation and measurement of power converter noise, specific conditions must be satisfied. It is assumed that reference resistors R_{I_1} and R_{I_2} are known and that the bandwidth of the stimulus signals, u_1 and u_2 , is predefined and controlled. Under these conditions, when the voltages of the stimulus between nodes a and g , v^{ag} , and between nodes b and g , v^{bg} , are equal, the circuit configuration allows for the estimation of the magnet impedances Z_1 and Z_2 within the bandwidth of the injected stimulus signal, although the resulting estimates will still be influenced by power converter noise [4]. The approach is broadly applicable and not restricted to a particular magnet: any load that can be divided into two partitions can, in principle, be characterized using this powered impedance measurement method.

In this article, the voltage balancing condition is achieved, as shown in previous work [4], by imposing three key design and measurement constraints: 1) the magnets are assumed equal, $Z_1 = Z_2$; 2) the reference resistors and other injection-stage components are identical, $R_{I_1} = R_{I_2}$, $R_{P_1} = R_{P_2}$, and $C_1 = C_2$; and 3) the stimulus signals are identical, $u_1 = u_2$. These constraints were explicitly enforced in the system design to enable impedance estimation, although, in practice, they are only

approximately satisfied. For this analysis, they are assumed as exact.

A key requirement for impedance estimation is that the system remains linear time-invariant (LTI) within the measurement window. However, in practice, a magnet's impedance varies as a function of the current level, leading to deviations from ideal LTI behavior. This is driven by effects, such as magnetization of the iron yoke [13], magnetization of the conductors [14], and induced eddy currents [15]. For accelerator magnets, inductance can in certain cases change by more than 50% across the full powering range [13], [16], [17]. To minimize errors introduced by this effect, measurement windows must be sufficiently short relative to the rate of change in impedance, ensuring that the system remains invariant during estimation.

Under these conditions, this differential probing topology not only enables impedance estimation but also results in the cancellation of the stimulus signals when the voltage is measured across the power converter. This cancellation isolates the power converter noise, which can then be measured at its output and leveraged in the noise reduction process discussed in Section IV. In this context, the "power converter output" refers to the measurable voltage v^{ab} . If the power converter was ideal with zero source impedance ($Z_{pc} = 0$), then v^{ab} would exactly reflect the converter's output voltage. Although this condition is not strictly met in practice, v^{ab} remains the best available proxy for observing the converter's contribution to the circuit.

B. Discrete-Time Modeling of Stimulus and Power Converter Sources

With the circuit configuration and underlying assumptions established, we now present a model of the signals in the circuit, including the stimulus voltage, the power converter output, and how these signals may be used to recover the magnet impedance. The stimulus voltage is designed as identical multisine signals [4], [18] with K frequency components

$$u_1[n] = u_2[n] = \sum_{k=1}^K |\alpha_k| \cos(\psi_k n + \angle \alpha_k) \quad (2)$$

where the following holds.

$u_1[n], u_2[n] \in \mathbb{R}$ represent the multisine signals.

$\alpha = [c_1 e^{j\phi_1}, c_2 e^{j\phi_2}, \dots, c_K e^{j\phi_K}]^T \in \mathbb{C}^K$ contains the complex amplitudes (real amplitude and phase) of the K sinusoids.

$c \in \mathbb{R}^K$, $\phi \in (-\pi, \pi]^K$, and $\psi_k = 2\pi \frac{f_{stim,k}}{f_s}$ are the amplitude, phase, and angular frequency vectors of the sinusoids, respectively, and $\mathbf{f}_{stim} = [f_{stim,1}, f_{stim,2}, \dots, f_{stim,K}]^T \in [0, \frac{f_s}{2}]^K$ denotes the vector of stimuli frequencies in Hertz.

The designed stimulus signal provides a controlled input for system identification. The power converter output, in contrast, cannot be easily modeled in general terms, as its behavior depends on the specific hardware, control loops, and operating conditions [8]. In the present circuit, the converter introduces both deterministic effects, such as harmonics of the grid frequency and controller-induced oscillations, and random noise




Fig. 3. Illustration of subtransfer function for u_{pc} .

sources, such as switching noise. These contributions vary with the circuit current and the converter's ramping rate, making the power converter output a nonstationary noise source that interferes with impedance estimation.

The measurement approach adopted in this article considers a topology in which voltages are measured differentially between pairs of nodes without a common reference. In such cases, the classical interpretation of a transfer function or impulse responses, often defined as a voltage-to-voltage ratio relative to ground, is inadequate. We therefore adopt the multiport methodology [19], which is suited to differential voltage measurements, i.e., where input and output signals are not referenced to a common node.

Under the assumption of linearity, the measurable voltage between any two node pairs can be expressed as a superposition of contributions from the system's three voltage sources ($u_{pc}[n]$, $u_1[n]$, and $u_2[n]$). Specifically, the voltage is equivalent to the sum of the convolution of each source signal with its corresponding impulse response between the node pairs. The individual impulse responses correspond to the voltage that would be measured between the output node pair if the associated source in Fig. 2 was replaced by an ideal impulse generator, while the other two sources were short-circuited. This concept is illustrated in Fig. 3. Specifically, this figure illustrates how the impulse response, $h_{pc}[n]$, corresponding to sources u_{pc} is derived from the circuit at the node pairs (a, g) and (a, b) .

As with previous notation, the absence of node-specific superscripts implies a general expression. Similarly, if no specific source is indicated (e.g., $h[n]$ instead of $h_{pc}[n]$), the expression refers to a general case.

Signals in this study are represented either in discrete-time (e.g., $v[n]$, $i[n]$) or frequency-domain forms (e.g., $V(f)$ and $I(f)$). For discrete-time impulse responses, it is natural to define their dual representation in the z -domain, resulting in a transfer function denoted $H(f)$. However, as this article is concerned only with sinusoidal steady-state excitation, all transfer functions are evaluated on the unit circle. We therefore simplify and define the frequency-domain transfer function as follows:

$$H(f) := H(z) \Big|_{z=e^{j2\pi f/f_s}} \quad \text{for } f \in \left[-\frac{f_s}{2}, \frac{f_s}{2}\right]. \quad (3)$$

Using these impulse responses along with the superposition principle, the measurable voltage between any given node pair can be expressed as follows:

$$v[n] = h_{pc}[n] * u_{pc}[n] + h_1[n] * u_1[n] + h_2[n] * u_2[n] \quad (4)$$

$$\Downarrow \mathcal{F}$$

$$V(f) = H_{pc}(f) U_{pc}(f) + H_1(f) U_1(f) + H_2(f) U_2(f). \quad (5)$$

Given that the system is assumed to be LTI and the two stimulus sources generate identical multisine signals, their contribution to the measured voltage between any node pair will be a similar multisine with the same frequency components, but with modified amplitudes and phase. We denote this $y[n]$

$$y[n] = h_1[n] * u_1[n] + h_2[n] * u_2[n] \quad (6)$$

$$= \sum_{k=1}^K |\beta_k| \cos(\psi_k n + \angle \beta_k) \quad (7)$$

where

$\beta \in \mathbb{C}^K$ is a complex-valued vector where each entry encodes the amplitude and phase of a sinusoidal component at one of the K stimuli frequencies, as observed between the chosen nodes.

Similarly, we define the stimulus-induced current through an arbitrary circuit element as $l[n]$

$$l[n] = \sum_{k=1}^K |\gamma| \cos(\psi_k n + \angle \gamma_k) \quad (8)$$

where the following holds.

$\gamma \in \mathbb{C}^K$ is a complex-valued vector containing the amplitudes and phases of the current components at each stimulus frequency.

For a given circuit element with impedance $Z \in \mathbb{C}$, the relation between the current and voltage amplitudes is given by $\gamma_k = \frac{\beta_k}{Z}$.

As with previous sequences, the corresponding frequency-domain representations of $y[n]$ and $l[n]$ are denoted by $Y(f)$ and $L(f)$, respectively.

In contrast to $y[n]$ and $l[n]$, the contribution from the power converter is not constrained to the frequencies in \mathbf{f}_{stim} . Hence, the total voltage measured between two nodes is the sum of the stimulus-induced multisine and the noise contribution from the power converter, leading to the following model:

$$v[n] = \sum_{k=1}^K |\beta_k| \cos(\psi_k n + \angle \beta_k) + \epsilon[n] \quad (9)$$

$$= y[n] + \epsilon[n] \quad (10)$$

where the following holds.

$v[n]$ represents the measured voltage between any two unspecified nodes.

$\epsilon[n] = h_{pc}[n] * u_{pc}[n]$ represents the contribution of power converter noise to the measured voltage, where $h_{pc}[n]$ is the impulse response between the power converter and the specific node pair.

It has been shown [4] that the impedance of the two individual magnet segments at the K stimuli frequencies can be estimated as follows:

$$\hat{\gamma}_k^{R_{I_1}} = \frac{\hat{\beta}_k^{lj}}{R_{I_1}}, \quad \hat{\gamma}_k^{R_{I_2}} = \frac{\hat{\beta}_k^{ek}}{R_{I_2}} \quad (11)$$

$$\hat{Z}_1(f_{\text{stim},k}) = \frac{\hat{\beta}_k^{ag}}{\hat{\gamma}_k^{R_{I_1}}}, \quad \hat{Z}_2(f_{\text{stim},k}) = \frac{\hat{\beta}_k^{bg}}{\hat{\gamma}_k^{R_{I_2}}} \quad (12)$$

$$\hat{Z}_{\text{total}} = \hat{Z}_1 + \hat{Z}_2 \quad (13)$$

where the following holds.

$\hat{\cdot}$ denotes an estimation of the target variable.

$\hat{\gamma}_k^{R_{I_1}}$ and $\hat{\gamma}_k^{R_{I_2}}$ correspond to the complex amplitude estimates of the currents $I_{R_{I_1}}$ and $I_{R_{I_2}}$ displayed in Fig. 2, respectively, for the k stimuli frequencies.

The impedance estimation process requires four key parameters: $\hat{\beta}^{ag}$, $\hat{\beta}^{lj}$, $\hat{\beta}^{bg}$, and $\hat{\beta}^{ek}$, corresponding to the complex amplitude across the magnet partitions ($\hat{\beta}^{ag}$ and $\hat{\beta}^{bg}$), and the reference current resistor ($\hat{\beta}^{lj}$ and $\hat{\beta}^{ek}$). Thus, to proceed, these complex amplitudes must be estimated.

C. Complex Amplitudes Estimation

Numerous methods exist to estimate the amplitude and phase angle at specific and known frequencies. These include the Fourier transform, least squares, weighted least squares, and matched filter-banks, all of which are discussed in [18]. When dealing with a large number of samples, N , and frequencies that are not closely spaced, the mean-squared error (mse) performance of each of these four methods converges. These estimates are dependent on the local SNR around each of the stimuli frequencies [20]

$$\text{SNR}_k = \frac{|\beta_k|^2}{\int_{f_{\text{stim},k-1}/(2N)}^{f_{\text{stim},k+1}/(2N)} |U_{\text{pc}}(f)|^2 df} \quad (14)$$

where the following holds.

SNR_k is the local SNR at stimuli frequency k .

$|U_{\text{pc}}(f)|^2$ denotes the power spectral density of the power converter noise.

N is the window size used to form the estimate.

While any of the mentioned method would be equally valid for this study, we retain the least squares method for consistency with previous studies [2], [3], [4]. This process takes a vector \mathbf{v} of size N to produce an estimate of the complex amplitudes, $\hat{\beta}$, at the stimuli frequencies. For our system where the digital-to-analog converters and analog-to-digital converters are driven from the same sample clock, this method is mathematically identical to using a lock-in amplifier.

To aid interpretation, Fig. 4 provides a generic, end-to-end visualization of the key signals involved in measuring the voltage between any two nodes and subsequently estimating the complex amplitudes at the stimuli frequencies. The amplitude- and phase-shifted stimulus signal $y[n]$ is added to the noise $u_{\text{pc}}[n]$ filtered by the system's frequency response $H(f)$. This combined signal

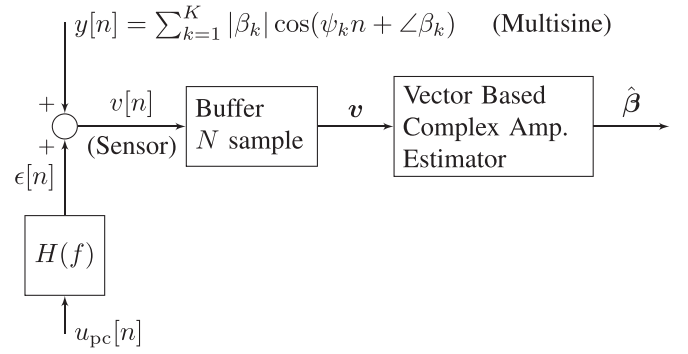


Fig. 4. Generic block diagram illustrating the signal flow from stimulus injection to complex amplitude estimation.

is measured by the voltage sensor as $v[n]$, then digitally buffered into a vector \mathbf{v} of length N , and subsequently processed by a complex amplitude estimation algorithm to recover the amplitudes and phases of the stimulus signal's frequency components. The specific forms of $v[n]$, $H(f)$, and $y[n]$ depend on the particular pair of nodes selected for the voltage measurement.

III. REFERENCE DATA

A. Data Collection

Having introduced the signal model and a method for impedance estimation, we now present the dataset used in this study. These data demonstrate the impact of system noise on the impedance estimates and serve as a benchmark for the remainder of the study.

The data were collected from the MQXFS7j magnet [21] at CERN's SM18 magnet test facility [22]. The magnet remained at a temperature of 1.9 K, and the power converter used was a 20 kA thyristor-based system. The differential coupling system from [4] was employed for probing the magnet, with the key difference that the previously used ac-coupled analog channels were replaced by dc-coupled ones, as used in [23]. Data were acquired continuously throughout the measurement using the setup described in [3]. In that study, a sampling rate of 200 kS/s exceeded the system's data handling capacity, resulting in dropped samples. In this study, we reduced the sampling rate to 40 kS/s, ensuring uninterrupted and uniformly spaced samples.

To measure the current, the injection system reference resistors, R_{I_1} and R_{I_2} , were set to 1 Ω . As part of the circuit protection, capacitors of 5 μmF were used for C_1 and C_2 , and 10 Ω resistors for R_{P1} and R_{P2} (nominal values; component tolerances are not considered in this work). The ability to probe impedance over a wide frequency range is generally valuable. At low frequencies (below 1 Hz), measurements are useful for studying magnet inductance, while frequencies up to about 100 kHz allow access to the resonance peak, which recent studies indicate is particularly relevant for magnet health monitoring, for instance in the detection of interturn shorts [5], [6], [7]. Beyond this range, capacitive coupling to ground tends to dominate the response. For this study, the injection system was configured to inject three frequency components (f_{stim}): 10, 110, and 1000 Hz. The frequencies were chosen to span a broad spectral range to

assess denoising performance at different regions, while complementing previous measurements that focused on higher frequencies (above 1 kHz) [3], [4]. Ideally lower frequencies would be tested; however, the current iteration of the injection electronics is not capable of injecting signals below this frequency. To compensate for the greater attenuation at 10 and 110 Hz caused by capacitive elements (C_1 and C_2) in the injection path, the digital-to-analog converter (DAC) output amplitudes for the stimuli frequencies of 10, 110, and 1000 Hz were scaled by factors of 100, 10, and 1, respectively. The multisine signal, $y[n]$ had a peak-to-peak amplitude of approximately 10 mV when measured between nodes a and g in Fig. 2, corresponding to one partition.

The powered measurement lasted 1800 s, during which the power converter was initially set to its standby current of 200 A, then ramped at 10 A/s to three consecutive current plateaus, each held for 300 s: 2, 5.5, and finally, 8.23 kA.

The current levels were selected to cover a broad operational range while remaining well below the magnet's nominal operating current of 16.23 kA [24]. This precaution minimized the risk of quenching at high energy. Beyond this constraint, the selected values are not of significance, but were chosen to ensure a varied dataset across different current levels.

To reduce inductance variations over time, the ramp rate should be low. Furthermore, a reduced current ramp rate would provide additional samples per current level, enabling impedance estimates with longer windows. However, this must be balanced against practical time constraints for conducting the measurements. In this study, the focus is not on precisely quantifying impedance but rather on analyzing noise characteristics, with particular emphasis on preserving the previously defined LTI assumptions. Based on this, a ramp rate of 10 A/s is chosen. For a magnet [13] similar to the one considered in our study, a maximum inductance change of approximately 5% per 1 kA has been observed. Although this figure only reflects inductance variation, it indicates that at the chosen ramp rate of 10 A/s, the system parameters vary slowly enough for the circuit to be reasonably approximated as time invariant within short estimation windows (e.g., ≤ 1 s).

During the measurement, voltage was recorded at five distinct sources: v^{ag} , v^{bg} , v^{lj} , v^{ek} , and v^{ab} . These correspond to the voltages across partitions Z_1 and Z_2 (v^{ag} and v^{bg}), across the reference resistors (v^{lj} and v^{ek}), and across the entire magnet (v^{ab}).

B. Overview of Collected Data

To provide an overview of the system dynamics during powering, a downsampled view of the power converter noise, v^{ab} , is shown in Fig. 1. This coarse overview reveals that the power converter noise is nonstationary. As current increases, the variance increases. Furthermore, the instantaneous change in current, di/dt , combined with the magnet's inductance results in a dc voltage offset during ramping.

To see how the impedance estimates evolve during the ramp, the data were segmented into 1 s and 10 ms intervals (equivalent to 40 000 and 400 samples, respectively), and the impedances

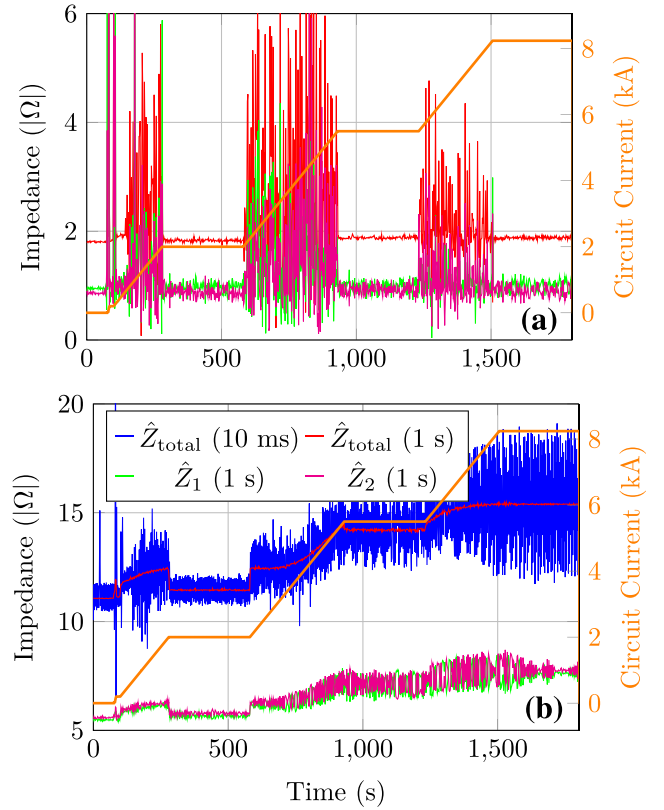


Fig. 5. Estimated impedance over time at (a) 10 Hz and (b) 1 kHz. 1 s windows (40 000 samples) and 10 ms windows (400 samples) were used to estimate each point. For (a), 10 ms estimates (blue) are omitted as they are too noisy to display. Shown for Z_{total} , Z_1 , and Z_2 .

of Z_1 , Z_2 , and Z_{total} were estimated for each of the three stimuli frequencies at each interval. These estimates are shown in Fig. 5. Fig. 5(a) and (b) shows the estimate for the 10 and 1000 Hz stimuli frequencies, respectively. The 110 Hz frequency is omitted for conciseness. Further, only the magnitude of the impedance is presented; phase angle information is also omitted.

First, considering the 1-s estimates in Fig. 5, we observe that the behavior of the impedance estimates differs significantly for each frequency. From independent di/dt -based inductance measurements of a similar magnet type (MQXFS5) [13], one would generally expect the inductance, and thus the impedance, to decrease with increasing current due to the progressive magnetic saturation of the iron yoke, particularly above approximately 2 kA. However, in the presented measurement, the impedance magnitude instead exhibits a gradual increase with current. The underlying physical mechanisms responsible for this unexpected trend are outside the scope of this work but are the subject of a separate study [25].

In all cases, Fig. 5 shows the Z_{total} estimate to be less noisy than the estimates of the individual components, Z_1 and Z_2 . At the 10-Hz estimates, the noise level is significantly higher during the ramping parts of the current cycle. On the other hand, in the 1-kHz estimates shown, the noise increases with current level.

To better understand the mechanism causing the estimation discrepancy between the individual segments, Z_1 and Z_2 , and

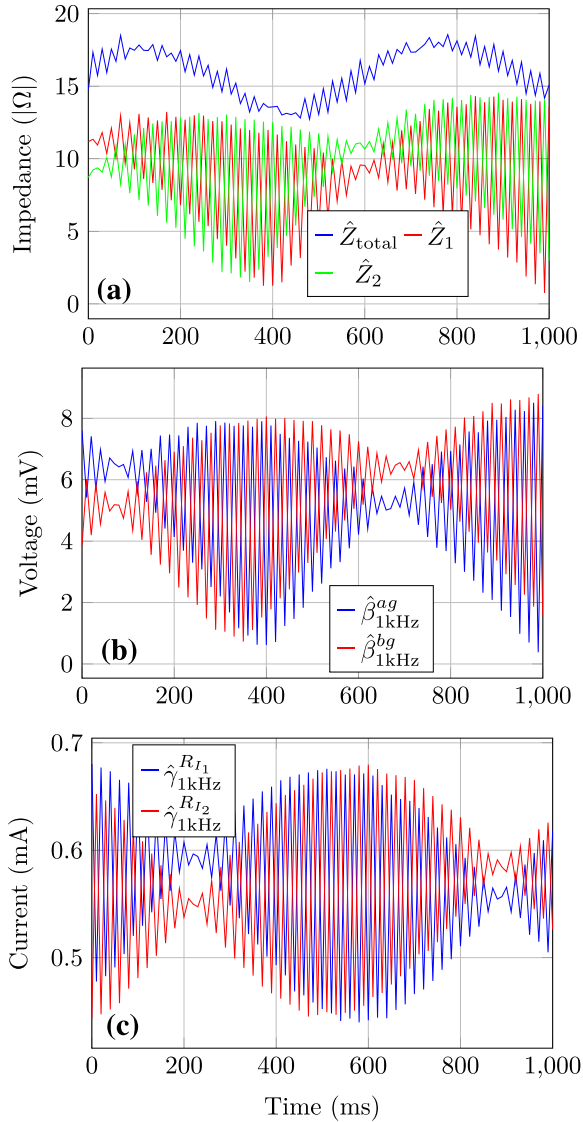


Fig. 6. Estimated (a) impedance, (b) voltage, and (c) current over time for the 1 kHz stimulus frequency. Estimates are computed using 10 ms windows (400 samples) over a 1-s interval during the 8.23 kA plateau stage.

the total impedance, Z_{total} , it is appropriate to zoom in and observe the short-term evolution of the estimates. An example of this is shown in Fig. 6, where the 10-ms segmented estimated impedances, as well as the underlying voltages and current estimates are shown for 1 s during the 8.23 kA stage.

The plots reveal significant amplitude modulation, with the two voltage estimates 180° out of phase across the partitions, and similarly for the two current estimates, leading to partial cancellation in the Z_{total} estimate. This destructive interference between v^{ag} and v^{bg} explains why Z_{total} appears markedly smoother than the individual partition estimates, even under noisy conditions.

To better understand the origins of these observations we will take a closer look at the time- and frequency-domain behavior of the collected data.




Fig. 7. 40 ms snippet of the full capture during the 8.23 kA plateau. Shows the voltage across both magnet halves ($v^{ag}[n]$ and $v^{bg}[n]$), across the whole magnet ($v^{ab}[n]$) and an estimate of the injected stimulus for reference ($\hat{y}[n]$).

C. Time- and Frequency-Domain View of Data

As the total data consist of 72 million data points for each of the five channels, it is impractical to present the raw data in its entirety within the scope of this article. Thus, we select representative time slices and data transformations to illustrate the overall dynamics, explain patterns observed in the impedance-over-time plots (see Figs. 5 and 6), and lay the groundwork for the denoising analysis presented in Section IV. Furthermore, we will limit specific discussion regarding the stimuli frequencies to the 10-Hz and 1-kHz cases, choosing to omit 110 Hz in much of the subsequent presentation, as the discussion on the other two cases largely applies to it.

Fig. 7 shows a 40-ms segment of voltage data recorded during the 8.23 kA plateau, approximately 1550 s into the measurement cycle. It shows the voltages across each magnet half, $v^{ag}[n]$ and $v^{bg}[n]$, and across the full magnet, $v^{ab}[n]$. An estimate of the injected signal, \hat{y} , is also plotted for reference. Since the injected signal is not directly measurable, it is reconstructed using the least-squares method described in Section II.

The power converter noise at 50 Hz dominates the voltage measurement and contributes significantly to the variance in the impedance estimates. The differential injection causes the 50-Hz noise to appear approximately 180° out of phase in $v^{ag}[n]$ and $v^{bg}[n]$, resulting in partial cancellation when forming \hat{Z}_{total} from \hat{Z}_1 and \hat{Z}_2 , further explaining the smoother \hat{Z}_{total} impedance estimates observed earlier in Fig. 6.

The dominance of the power converter noise renders it difficult to discern how interference affects the impedance estimates at each of the three stimuli frequencies. This makes it challenging to interpret the previously shown unfiltered time-domain plots. To address this, we apply band-limiting FIR filters to isolate the time-domain behavior near the 10 Hz and 1 kHz stimuli frequencies. The filter targeting 10 Hz was designed as a low-pass filter with a -3 dB cutoff frequency at 23 Hz, using 14 001 coefficients. Similarly, to highlight behavior near the 1 kHz frequency, a bandpass filter using 6001 coefficients and with -3 dB bandwidth ranging from 930 to 1070 Hz was applied on the signal.




Fig. 8. Estimated voltage during the 2–5.5 kA ramping stage, shown after (a) applying a low-pass FIR filter with a 23 Hz to -3 dB cutoff frequency, and (b) applying a band-pass FIR filter with a passband from 930 to 1070 Hz.

The resultant time series following filtering are shown in Fig. 8 during the ramp from 2 to 5.5 kA. Fig. 5(a) and (b) shows for the 23-Hz low-pass filter and 1-kHz bandpass filter, respectively.

At 10 Hz, the signal exhibits recurring voltage spikes with varying amplitude. These transients are presumed to result from flux jumps, sudden redistribution of magnetic flux that occur during current ramping in superconducting magnets [26]. In terms of a circuit analogy, they manifest as abrupt, pulse-like changes in the magnet’s inductance, which during ramping give rise to momentary voltage spikes. From a modeling perspective, they can be treated as randomly occurring pulses with variable amplitude, shaped by a bandpass filter [27]. As the magnet current increases, both the intensity and rate of these events initially rise, and then gradually decline as the current approaches the magnet’s nominal operational current value [27]. This trend is visible in the impedance estimates at 10 Hz [see Fig. 5(a)], where the noise level increases during ramping up to around 4 kA, before decreasing again at higher currents. These observations strongly suggest that the poor 10 Hz impedance estimation performance during ramping is driven by flux jump activity.

At 1000 Hz, as shown in Fig. 8(b), the filtered voltage traces exhibit highly regular sinusoidal waveforms. There is no visible distortion attributable to flux jumps at this frequency. However, a clear amplitude modulation at 50 Hz is present, consistent with the base grid frequency.

To complement the earlier time-domain analysis and better characterize the spectral content of the voltage signals, we estimate their power spectra. All power spectra in this article are estimated using Welch’s method [28] with 50% overlap between segments. Specifically, we examine a zoomed-in view of the power spectrum surrounding each frequency, as shown in Fig. 9. Stimuli frequencies are marked with dots at their respective locations in the plots.

Fig. 9(a) presents the power spectrum from 0 to 55 Hz. During the ramping stage, the noise floor is significantly elevated compared to the plateau stages, which both show similar




Fig. 9. Zoomed power spectra of v^{ag} : (a) from 0 to 55 Hz, and (b) from 945 to 1055 Hz. Spectra were estimated using 19 segments of 400 000 samples each.

spectral profiles. The spectral shape during ramping is consistent with flux jump-related noise described in [27], reinforcing the interpretation that these transients are responsible for the disturbances observed during ramping.

Fig. 9(b) shows the power spectrum from 945 to 1055 Hz, centered around the 1 kHz stimulus frequency. Unlike the lower frequency case, the spectral content during ramping remains fixed. However, the noise level increases with current, particularly during the 8.23 kA plateau. Distinct spectral peaks are visible at 950, 1000, and 1050 Hz, corresponding to harmonics of the 50-Hz grid interference.

In contrast to the voltage measurements, which exhibit considerable variability and noise, especially during ramping, the current measurements exhibit less variation across all stages of powering. Fig. 10 shows the current signal’s frequency representation. Across all powering stages, compared to the voltage measurements, the current exhibits a much lower noise floor at the stimuli frequencies, with no temporal or spectral evidence of flux jumps during ramping. The low noise level is attributed to the capacitor in the injection path, which suppresses low-frequency disturbances before they reach the reference resistor.

IV. METHOD FOR NOISE REDUCTION

A. Principles of Noise Reduction

Reducing the variance in the impedance estimates requires denoising of the voltage and current signals. Although flux jumps were identified in Section III as a dominant source of low-frequency disturbances during ramping, they are not addressed




Fig. 10. Zoomed spectrum of the estimated current through reference resistor R_{I_1} . Shown from 0 to 55 Hz. Estimated using 19 segments of 400 000 samples each.

in the denoising strategy presented here. Instead, we focus exclusively on reducing the effect of power converter noise. While both power converter and flux jump noise is contained in v^{ab} , flux jumps manifest as localized, asymmetric voltage drops across the magnet [29]. Thus, this interference is distributed differently across the circuit elements than noise originating from the power converter. Therefore, to treat both disturbances, they would have to be decomposed and isolated. However, we leave this for future work, and for the rest of this article, we focus on plateau stages where flux jumps do not occur.

For clarity and to distinguish it from the other signal channels, we define the voltage measurement across the magnet as $\eta[n] = v^{ab}[n]$ and introduce $d[n] = h_{pc}^{ab}[n]$ as the impulse response shaping the power converter noise across the magnet. In this context, by “power converter noise,” we refer exclusively to the voltage source u_{pc} , excluding its associated source impedance Z_{pc} .

Ideally, $\eta[n]$ should contain only the power converter noise as convolved with the impulse response between nodes a and b of the circuit in Fig. 2

$$\eta[n] = \sum_{k=1}^K |\beta_k^{ab}| \cos(\psi_k n + \angle \beta_k^{ab}) + d[n] * u_{pc}[n]. \quad (15)$$

In theory, the stimulus-related terms, β^{ab} , should go to zero, as the differential injection scheme is designed to cancel them entirely. This would leave $\eta[n]$ as a pure observation of the power converter noise convolved with the impulse response.

Assuming that we have perfect cancellation, the core objective of the denoising process is to exploit this reference to reconstruct and subtract the noise contributions present in the voltage and current measurement channels. In other words, the goal is to design a filter with frequency response $B(f)$ that maps the noise shaping response across the magnet, $D(f)$, to the noise shaping response in a given measurement channel, $H(f)$. This ideal mapping is given exactly by

$$B(f) = \frac{H(f)}{D(f)}. \quad (16)$$

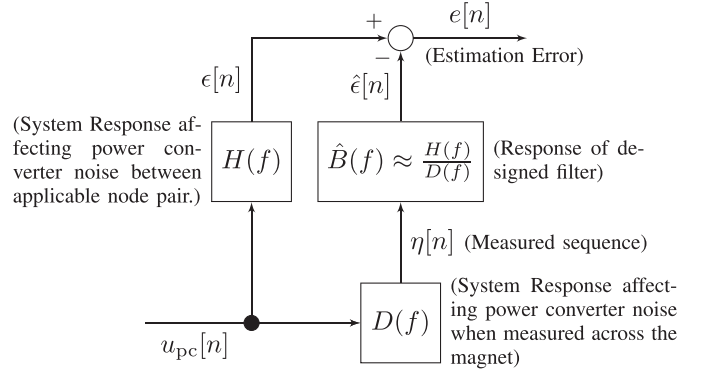


Fig. 11. Block diagram of denoising scheme key signals and goals.

In practice, however, we are faced with a system identification problem: the true $B(f)$ is unknown, and the best we can achieve is to estimate it, resulting in an approximation $\hat{B}(f)$.

In the time domain, the designed filter is realized by a set of coefficients $\hat{b}[n]$. Filtering the noise reference $\eta[n]$ through $\hat{b}[n]$ yields an estimate, $\hat{\epsilon}[n]$, of the power converter noise contribution in the measurement channel

$$\hat{\epsilon}[n] = \hat{b}[n] * \eta[n] \quad (17)$$

where

$\hat{b}[n]$ is the estimated impulse response corresponding to $\hat{B}(f)$.

The estimation error, $e[n]$, is defined by the difference between the true noise contribution $\epsilon[n]$ and its estimate $\hat{\epsilon}[n]$

$$e[n] = \epsilon[n] - \hat{\epsilon}[n]. \quad (18)$$

The goal of estimating $B(f)$ is to minimize the estimation error according to a suitable optimality criterion.

Fig. 11 provides a conceptual overview of the denoising strategy, showing the relationship between the noise reference $\eta[n]$, the estimated system $\hat{B}(f)$, the resulting noise estimate $\hat{\epsilon}[n]$, and the estimation error $e[n]$.

B. Characterization of Stimulus Leakage Into Noise Measurement

The noise reduction approach relies on the assumption that the reference signal $\eta[n]$ contains only power converter noise, with negligible residual stimulus contributions. To evaluate the validity of this assumption, we compare the spectral content of the voltage measured across a single partition, $v^{ag}[n]$, with that of the full magnet voltage, $\eta[n]$ ($v^{ab}[n]$). Fig. 12 presents the power spectrum estimates for both signals, in the vicinity of the 10 Hz stimuli frequency as an example.

The results confirm that the differential injection effectively suppresses the stimulus contributions in the noise measurement channel, as intended. However, the suppression is not perfect, as some leakage of the stimulus signal into $\eta[n]$ remains. This leakage is likely due to imperfections in the experimental setup, such as imbalances in the circuit that deviate from the idealized assumptions.




Fig. 12. Zoomed power spectrum of v^{ag} and v^{ab} (i.e., η), during the 8.23 kA plateau. Shown from 0 to 55 Hz. Estimated using 19 segments of 400 000 samples each.

C. Noise Reduction Using a Wiener Filter

Having confirmed in the previous section that the noise reference $\eta[n]$ predominantly contains power converter noise with minimal stimulus leakage, we proceed to the realization of the denoising scheme. To implement the denoising scheme in practice, we must estimate the filter coefficients $\hat{b}[n]$ from data. Among the available techniques, Wiener filtering represents a classical and well-established method for noise reduction [30]. Under the assumptions that the system is linear and the signal is stationary within the estimation window, Wiener filtering yields the linear filter that minimizes the mse between the true and estimated noise contribution in the measurement channel, $e[n]$ [30].

However, Wiener filtering is inherently nonadaptive and thus less suited for situations where the system is time-varying, such as in our case where the magnet impedance varies throughout the powering cycle. Nonetheless, we adopt Wiener filtering as a baseline method due to its well-established theoretical foundation, straightforward implementation, and interpretability. In the following, we first introduce and develop a static Wiener filtering approach based on long segments of data. This is subsequently extended to a blockwise adaptive version, which allows the filter to update over time in response to changes in system dynamics. The two methods are later compared in terms of denoising performance, providing insight into whether the added complexity of adaptivity is justified and to what extent the system dynamic varies during operation.

We restrict our attention to FIR filters, as infinite impulse response (IIR) realizations are practically infeasible with finite data. For a filter of length q , the optimal coefficients are given by the Wiener–Hopf equations [30]

$$\underbrace{\begin{bmatrix} R_\eta[0] & R_\eta[1] & \dots & R_\eta[q-1] \\ R_\eta[1] & R_\eta[0] & \dots & R_\eta[q-2] \\ \vdots & \vdots & \ddots & \vdots \\ R_\eta[q-1] & R_\eta[q-2] & \dots & R_\eta[0] \end{bmatrix}}_{\mathbf{R}_\eta} \underbrace{\begin{bmatrix} \hat{b}[0] \\ \hat{b}[1] \\ \vdots \\ \hat{b}[q-1] \end{bmatrix}}_{\hat{\mathbf{b}}}$$

$$= \begin{bmatrix} r_{\eta,\epsilon}[0] \\ r_{\eta,\epsilon}[1] \\ \vdots \\ r_{\eta,\epsilon}[q-1] \end{bmatrix} \quad (19)$$

$\mathbf{r}_{\eta,\epsilon}$

where the following holds.

$R_\eta[k] = E[\eta[n]\eta[n+k]]$ is the autocorrelation of $\eta[n]$.
 $r_{\eta,\epsilon}[k] = E[\eta[n]\epsilon[n+k]]$ is the cross-correlation between $\eta[n]$ and $\epsilon[n]$.

Solving this system yields the filter coefficients

$$\hat{\mathbf{b}} = \mathbf{R}_\eta^{-1} \mathbf{r}_{\eta,\epsilon}. \quad (20)$$

Since the statistical properties of $\eta[n]$ and $\epsilon[n]$ are unknown, we estimate the statistics empirically from data. The autocorrelation of $\eta[n]$ for an M length sequence is estimated as follows:

$$\hat{R}_\eta[k] = \sum_{i=0}^{M-k-1} (\eta[i] - \bar{\eta})(\eta[i+k] - \bar{\eta}) \quad (21)$$

where

$$\bar{\eta} = \frac{1}{M} \sum_{k=0}^{M-1} \eta[k]$$

is the sample mean of $\eta[n]$.

Similarly, the cross-correlation between $\eta[n]$ and $\epsilon[n]$ is estimated as follows:

$$\hat{r}_{\eta,\epsilon}[k] = \sum_{i=0}^{M-k-1} (\eta[i] - \bar{\eta})(\epsilon[i+k] - \bar{\epsilon}) \quad (22)$$

where

$$\bar{\epsilon} = \frac{1}{M} \sum_{k=0}^{M-1} \epsilon[k]$$

is the sample mean of $\epsilon[n]$.

However, we do not have direct access to $\epsilon[n]$, as it is superimposed with the injected stimulus signal in the measured voltage. We therefore substitute $\epsilon[n]$ with $v[n]$, the full measured signal. This approximation relies on the assumption that the multisine stimulus is uncorrelated with the noise reference $\eta[n]$. This condition holds when the stimulus is fully suppressed in the reference channel, i.e., when $\beta^{ab} = 0$. With this substitution, the final Wiener filter equation becomes

$$\hat{\mathbf{b}} = \hat{\mathbf{R}}_\eta^{-1} \hat{\mathbf{r}}_{\eta,v} \in \mathbb{R}^q \quad (23)$$

where the filter length, q , is an adjustable parameter.

As discussed previously, while the differential injection scheme is designed to suppress the stimuli frequency components in the noise reference channel, $\eta[n]$, this suppression is not perfect. Consequently, the stimulus signal will be correlated with $\eta[n]$. Thus, in practice, the noise suppression may potentially alter the signal.

The complete denoising and amplitude estimation process is summarized in Fig. 13. In this scheme, the multisine stimulus signal $y[n]$ is combined with power converter noise, modeled as $u_{pc}[n]$ filtered by $H(f)$, and measured by the voltage sensor as $v[n]$. Simultaneously, the power converter noise filtered by $D(f)$ is measured as $\eta[n]$, providing a noise reference signal. In

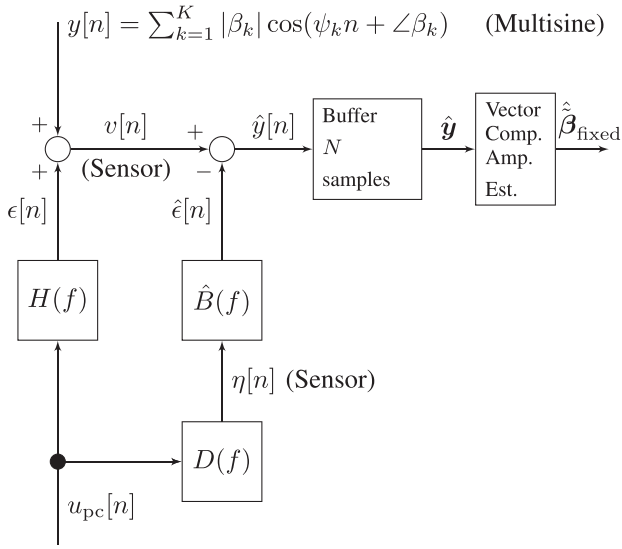


Fig. 13. Block diagram of denoising and amplitude estimation scheme.

postprocessing, $\eta[n]$ is filtered by the estimated system $\hat{B}(f)$ and subtracted from $v[n]$ to form the denoised signal. This denoised sequence is then buffered into vectors of length N , upon which least squares estimation is applied to recover the amplitudes and phases at each stimulus frequency. The resulting estimates, obtained after denoising using the fixed set of filter coefficients, are denoted $\hat{\beta}_{\text{fixed}}$.

D. Figure of Merit for Noise Reduction

Evaluating the effectiveness of noise reduction requires a suitable figure of merit. However, this presents a challenge, as to the best of the authors' knowledge, no traceable benchmark measurements exist for powered impedance monitoring since this is a new area of research. Consequently, the bias of the impedance estimates cannot be directly quantified. While a previous study [4] has validated the methodology against a trusted reference in nonpowered conditions, confirming its reliability in simpler scenarios, extending this validation to powered measurements would require detailed theoretical modeling and comprehensive error and sensitivity analyses, which is beyond the scope of this article.

Given these constraints, we adopt an empirical approach by assessing variance reduction as a measure of noise suppression. Specifically, we define the relative variance of the complex amplitude estimates, comparing results from the denoised signals to those from the raw, unfiltered data

$$\text{Relative variance at the } i\text{th frequency} := \frac{\text{var}\left(|\hat{\beta}_i|\right)}{\text{var}\left(|\beta_i|\right)} \quad (24)$$

where a relative variance less than 1 indicates that denoising has reduced fluctuations in the impedance estimates.

The variance of the impedance estimates depends primarily on two factors: noise characteristics and the length of the estimation window. Reliable variance estimates require the assumption of

signal stationarity, which, in the context of superconducting magnet operation, is only valid under specific conditions. While short segments during ramping stages may vary slowly, the most stable operating conditions are found during plateau stages, where the circuit current is held constant. In the measured data, these plateaus last for 300 s, during which we assume the system to be stationary for the purpose of variance estimation.

E. FIR Filter Length

The length of the FIR filters, q , is a pivotal, configurable parameter that affects the effectiveness of the denoising strategy. As we must estimate the complex amplitudes of four parameters, i.e., the stimulus voltages across the two partitions, and their corresponding currents, it is possible to have four unique values of q , one for each signal. However, as previously, we assume that the voltages and currents in the system are symmetrical, and thus for simplicity, we opt for using only two distinct filter lengths: one for the set of voltage measurements and another for the currents.

To evaluate the impact of varying q , we use the relative variance metric defined previously in (24). This requires a set of β estimates computed under comparable noise conditions. We use the 300 s of data from the 8.23 kA plateau, where the noise is expected to be approximately stationary, and the estimate will not be influenced by flux jumps (as they only occur during current ramping stages).

We applied FIR filters with lengths ranging from 1 to 1000 to the 8.23 kA plateau data, using 10 million samples per channel. Each denoised stream was divided into 1-s and 10-ms segments. For each segment, we estimated the signal amplitude and computed the relative variance as defined in (24). For some values of q , the Wiener–Hopf system of equation was ill-conditioned. In these cases, we applied ridge regression [31] to obtain well-behaved solutions. The resulting variance curves are shown in Fig. 14, illustrating the tradeoff between filter length and the reduced variance.

We observe that denoising leads to improvements in the amplitude estimates for both voltage and current across all frequencies. However, the improvements are less pronounced in the current estimates—especially for the 1-s segments. For example, at 10 Hz and 110 Hz, the relative variances improved to 0.5 and 0.2, respectively. Conversely, at 1 kHz, the relative variance decreased by two orders of magnitude. For both the 1-s and 10-ms segments, the trend for the current data displays a “U-shaped” curve; increasing the filter length initially leads to substantial improvements, which stabilizes around ten coefficients. Beyond this, the performance remains steady up to 200 coefficients, at which point the relative variance rises. One feasible explanation of this decline is leakage of the stimulus into the noise reference channel $\eta[n]$, which longer filters are more likely to pick up. In such cases, the denoising begins to attenuate the stimulus itself, rather than isolating just the power converter noise.

The behavior of the voltage estimates differs markedly from that of the currents; they are largely insensitive to changes in filter length, exhibiting a flat response up to several hundred

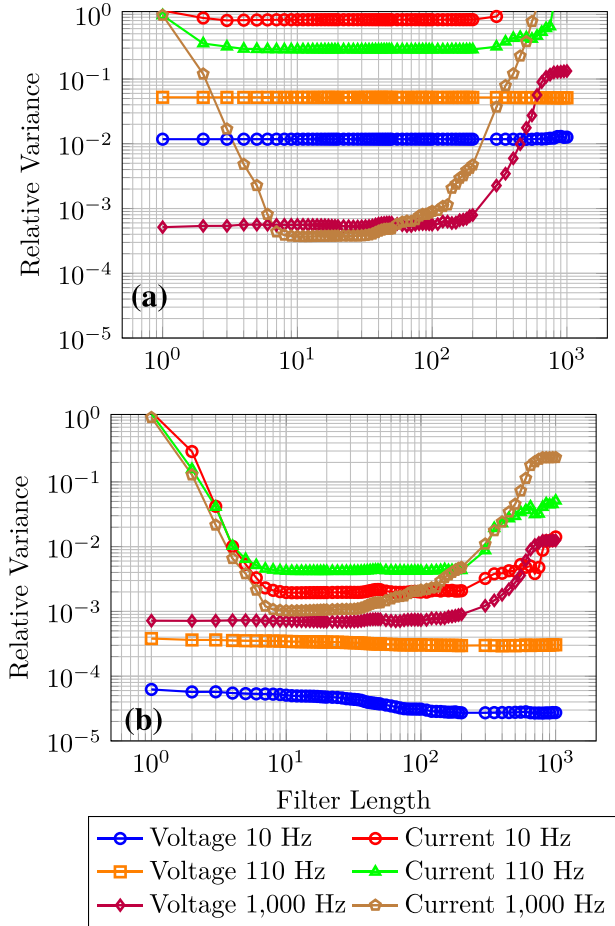


Fig. 14. Relative variance versus filter length for both voltage and current across three stimulus frequencies: (a) using 1-s segments and (b) using 10-ms segments. Data correspond to the 8.23 kA plateau and are shown for half the circuit, i.e., estimates of β^{ag} and $\gamma^{R_{I1}}$.

coefficients. A minor exception occurs at 10 Hz for the 10-ms segments, where the variance drops by roughly half an order of magnitude with increased filter length. Aside from this case, however, the near-flat response across filter lengths suggests that the noise contribution is predominantly scaled, and thus only one coefficient (i.e., a constant scaling of the signal with no shift in phase angle) is generally sufficient. This insensitivity to filter length is consistent with the underlying circuit model (see Fig. 2): assuming a balanced injection and negligible source impedance, Z_{pc} , the voltage noise across each partition is simply a scaled version of the power converter noise, with no phase distortion. Hence, only a scalar is needed to align the noise reference. In contrast, the current path includes capacitors, which introduce phase and amplitude distortion that requires longer filters to model effectively. In addition, the voltage estimates generally exhibited significantly greater improvement from denoising compared to the current estimates. This is likely due to their initially higher noise levels (both relative and absolute), as is evident in the voltage power spectrum plots (see Fig. 9) compared to the current power spectra (see Fig. 10), which provided a greater margin for noise reduction.

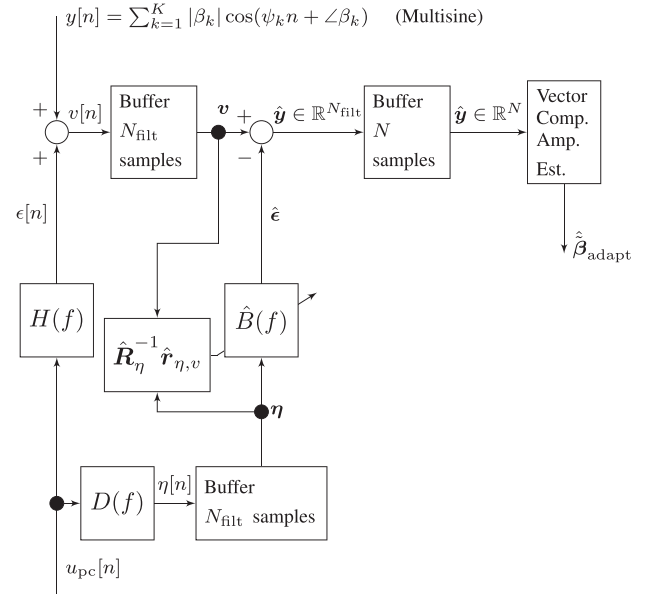


Fig. 15. Simplified block diagram of adaptive denoising and amplitude estimation scheme.

Based on this analysis, no single filter length provides optimal relative variance reduction across all stimulus frequencies. Instead, a tradeoff must be made between maximizing noise suppression and maintaining practical filter complexity. For voltage denoising, a filter length of $q = 1$ (i.e., a scalar) is selected, as higher orders provide negligible additional benefit. For current denoising, $q = 10$ is chosen since improvements stabilize across all frequencies beyond this point.

F. Adaptive Noise Reduction Approach

In addition to the static filtering approach, we also explore a simple adaptive variant. In this method, the Wiener filter coefficients are periodically re-estimated, using the same filter length q determined previously. The scheme is illustrated in the block diagram, as shown in Fig. 15, and is structurally similar to the static filtering approach of Fig. 13, with the key difference that the filter coefficients are periodically updated. The denoised estimate of this approach is denoted as $\hat{\beta}_{\text{adapt}}$. For the adaptive approach, two window lengths must be selected: one for the amplitude estimation, already set to 1 s or 10 ms, and one for the filter re-estimation procedure, denoted N_{filt} . We set $N_{\text{filt}} = 400000$ samples (equivalent to 10 s) to ensure that the filter estimation has sufficient samples to converge reliably, while assuming that the system dynamics do not vary significantly over this window.

To avoid issues related to startup transients when switching to newly estimated coefficients, overlapping data segments are used, and the initial transient portions are discarded.

V. RESULTS

A. Designed Responses

It is instructive to examine the frequency responses designed by the Wiener–Hopf equations, (19), based on the measured

data, and to compare these responses with the expected circuit behavior. This comparison allows us to assess how well the underlying circuit assumptions match reality.

We derive analytical estimates of the expected transfer functions based on the circuit assumptions outlined in Section II. These derivations further assume that the power converter source impedance is negligible, i.e., $Z_{pc} = 0$. The full derivations are provided in Appendix A. Since the derivations are performed in the continuous-time s -domain, care must be taken when comparing them to the transfer functions used throughout this article, which are expressed in the discrete-time z -domain. To obtain the frequency response from an s -domain transfer function, we evaluate it on the imaginary axis by substituting $s = j2\pi f$, and define

$$H(f) := H(s) \Big|_{s=j2\pi f} \quad \text{for } f \in \left[-\frac{f_s}{2}, \frac{f_s}{2}\right]. \quad (25)$$

Assuming ideal brick-wall antialiasing filtering and restriction to the frequency range $f \in [-\frac{f_s}{2}, \frac{f_s}{2}]$, the frequency responses in the s - and z -domains are equivalent [12].

To summarize, the transfer functions from Appendix A are given symbolically in terms of f , along with their evaluated forms using the component values listed in Section III

$$B_{\text{anly}}^{ag}(f) = 0.5 \quad B_{\text{anly}}^{bg}(f) = -0.5 \quad (26)$$

$$B_{\text{anly}}^{lj}(f) = \frac{-0.5R_{I_1}C_12\pi jf}{1 + (R_{P_1} + R_{I_1})C_12\pi jf} = \frac{-5 \cdot 10^{-6} \pi jf}{1 + 11 \cdot 10^{-5} \pi jf} \quad (27)$$

$$B_{\text{anly}}^{ek}(f) = \frac{0.5R_{I_2}C_22\pi jf}{1 + (R_{P_2} + R_{I_2})C_22\pi jf} = \frac{5 \cdot 10^{-6} \pi jf}{1 + 11 \cdot 10^{-5} \pi jf} \quad (28)$$

where the subscript ‘‘anly’’ denotes the analytically derived transfer functions.

Focusing first on the analytical estimates of the partition voltages, $B_{\text{anly}}^{ag}(s)$ and $B_{\text{anly}}^{bg}(s)$, it is expected that they take values of 0.5 and -0.5 , respectively, consistent with the balance constraint discussed earlier in relation to power converter noise cancellation. This also supports the earlier analysis indicating that a single filter coefficient is a suitable choice.

The fixed filter estimates obtained from the measured data, using ten million samples from the 8.23 kA plateau, are close to the analytical values

$$B_{\text{fixed}}^{ag}(f) = 0.498 \quad B_{\text{fixed}}^{bg}(f) = -0.502. \quad (29)$$

This close agreement further indicates that the two halves of the circuit behave nearly identically, consistent with the low levels of stimulus leakage previously demonstrated in Section IV-B. To compare this to the adaptive approach, where the filter coefficients are re-estimated every 10 s, we plot the evolution of the filter gains as a function of circuit current in Fig. 16.

The y -axis spans only from 0.48 to 0.52, indicating that the coefficient variations are relatively small. The estimated gains at the 2 and 8.23 kA plateaus are very similar. However, at the 5.5 kA plateau, a sudden shift in the coefficients is observed. The reason for this abrupt change is unclear. Notably, the relative

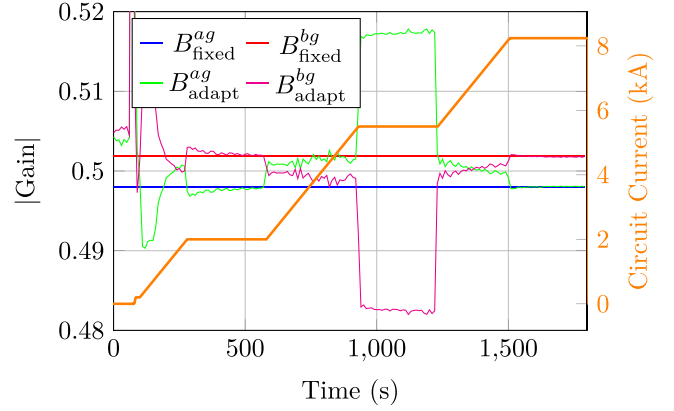


Fig. 16. Estimated gain versus time of B_{adapt}^{ag} and B_{adapt}^{bg} . Static 8.23 kA gains are also shown for reference.

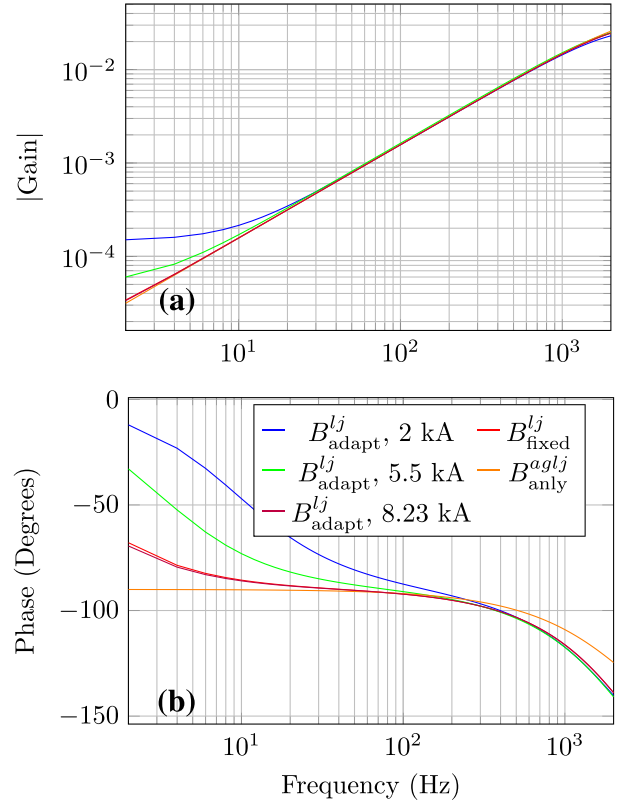


Fig. 17. Comparison between analytically derived, fixed, and adaptively designed current filter responses B_{adapt}^{lj} during the three different plateau stages: (a) magnitude response and (b) phase response.

weighting at 5.5 kA shifts toward B_{adapt}^{ag} , which receives greater weight than during the other two plateaus.

Next, to assess the differences between the analytical and data-derived estimates of the current denoising response, we plot the magnitude and phase responses in Fig. 17. As this filter involves multiple coefficients, we cannot simply plot a time series. Instead, we plot their magnitude and phase responses, selecting a representative filter realization from each plateau.




Fig. 18. 40 ms snippets of denoised $v^{ag}[n]$ during (a) the 2–5.5 kA ramp and (b) the 5.5 kA plateau stage.

The magnitude responses show good agreement between the analytical and estimated models. Although the phase responses do not match perfectly, the discrepancy remains within 20° across the plotted frequency range. The observed deviations are likely due to the simplifying assumptions made in the analytical model, and the presence of additional circuit elements not explicitly accounted for. Nonetheless, the results demonstrate that the key circuit dynamics are well captured by the simplified model.

B. Time- and Frequency- Domain Improvement From Denoising

To evaluate how effectively the denoising attenuates power converter noise, and to compare the performance of the static and adaptive approaches, we examine the time- and frequency-domain results with and without denoising. For brevity, we focus on the voltage across one partition ($v^{ag}[n]$) and the voltage across the reference resistor R_{I_1} ($v^{lj}[n]$), omitting the second pair of voltage and current signals, as they exhibit similar behavior.

Voltage signals are shown in Fig. 18. Each plot presents 40 ms excerpts of the raw measurements, the denoised signals using fixed filter coefficients, and the results from the adaptive approach. Fig. 18(a) and (b) corresponds to extracts from the ramp from 2 to 5.5 kA, and 5.5 kA, respectively.

In both cases, the denoising procedure significantly attenuates the power converter noise. Where the raw signal is almost entirely dominated by converter noise, after denoising, the signal

appears clean and periodic, consistent with the expected superposition of the three sinusoidal stimulus frequencies. During the ramping stage, while the 50-Hz noise and its harmonics are effectively removed, sharp transients remain. These transients are consistent with the flux jump phenomena discussed earlier and are not suppressed by the denoising process, which is designed specifically to target power converter noise. Finally, we note that the denoised traces from the static and adaptive denoising approaches are nearly indistinguishable across the plots. The only clear exception is the voltage across the magnet during the 5.5 kA plateau [see Fig. 18(b)], where slight differences between the two methods can be observed.

Turning to the zoomed power spectra shown in Fig. 19 for voltage, we again focus on the regions surrounding the stimuli frequencies at 10 Hz and 1 kHz. As in the time-domain analysis, we present the estimated spectra for the raw measurements, the denoised signals using constant filter coefficients, and the adaptively denoised signals. We show results for the different powering stages: the 2 kA plateau, the ramp from 2 to 5.5 kA, and the 5.5 kA plateau. The 8.23 kA plateau is omitted, as it does not contribute additional insight compared to the 2 kA plateau.

The frequency-domain plots for the voltage measurements, as shown in Fig. 19, confirm the overall trends seen in the time domain. In all powering stages, the 50-Hz harmonics are strongly attenuated by the denoising procedure. At 10 Hz, a significant reduction of the background noise is also observed, except during the ramping stage. While the 50-Hz harmonics near the stimulus frequency are still effectively removed, the overall noise floor remains elevated, likely due to flux jump activity that the denoising process does not suppress.

As in the case of the voltage signals, denoising also aids in the current measurements. However, the effect is not as pronounced. This is illustrated in Fig. 20, plotting the current spectrum before and after denoising from 0 to 55 Hz. While denoising also reduces the 50-Hz harmonics in the current spectra, its performance is less consistent than for the voltage signals. In some regions, it can even introduce additional artifacts. For example, as seen in Fig. 20(b), around the 1-kHz stimulus frequency, added noise appears in the surrounding ± 10 Hz range.

Finally, we note that the static and adaptive denoising approaches produce nearly identical results in this measurement scenario. As in the time-domain analysis, the traces from both methods almost fully overlap. The one notable exception is the voltage spectra during the 5.5 kA plateau, as seen in Fig. 19(c) and 19(f), where the adaptive approach achieves noticeably better suppression of the 50-Hz harmonics.

C. Impedance Estimates With and Without Denoising

Until now, we have examined the effect of denoising on the raw voltage and current measurements. In this section, we evaluate its impact on the impedance estimates, our primary quantities of interest.

After filtering the full dataset, we replicate the procedure from Section III: the signal is split into 1-s and 10-ms segments, and the impedance is estimated for each segment. Both the static and adaptive denoising approaches were applied to the

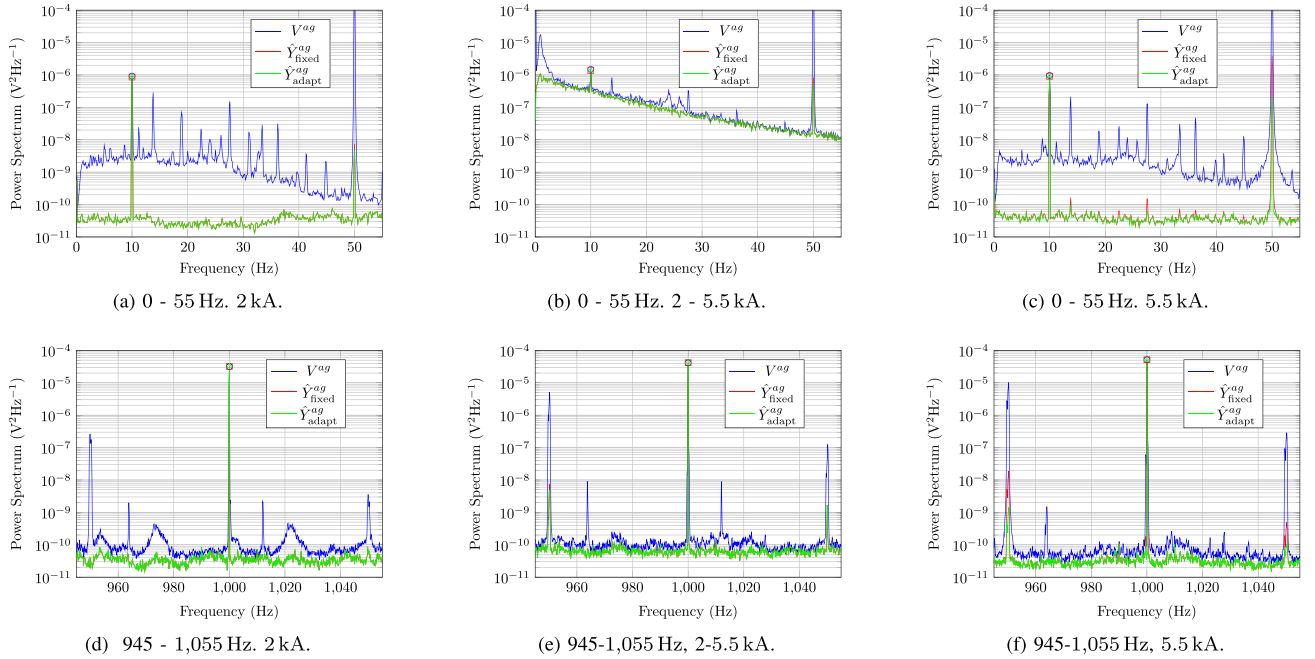


Fig. 19. Zoomed power spectrum of voltage across Z_1 during the 2 kA plateau (first column), ramp from 2 to 5.5 kA (middle column) and 5.5 kA plateau (last column). Shown around 10 Hz (first row) and 1 kHz (second row). Each plot shows the estimated power spectrum with no denoising (V^{ag}), with denoising using the static coefficients (\hat{Y}_{fixed}^{ag}), and the adaptive approach (\hat{Y}_{adapt}^{ag}). Each spectrum is estimated using 19 segments of 400 000 samples. (a) 0–55 Hz, 2 kA. (b) 0–55 Hz, 2–5.5 kA. (c) 0–55 Hz, 5.5 kA. (d) 945–1,055 Hz, 2 kA. (e) 945–1,055 Hz, 2–5.5 kA. (f) 945–1,055 Hz, 5.5 kA.

full dataset. However, to limit the number of plotted traces and avoid overcrowding the figures, we only show results obtained using the adaptive filter. This is shown in Fig. 21. The plots compare the estimated impedance of Z_{total} and Z_1 at 10 and 1000 Hz, with and without denoising for the 1-s and 10-ms window impedance estimates. The Z_2 curve is omitted as it closely follows the behavior of Z_1 . In addition, note that Z_{total} is omitted for the 10-ms window case prior to applying denoising, as the estimate is incoherent and varies wildly, reaching values of several hundred ohms.

Fig. 21 shows broadly similar results for both frequencies. For the 10-ms windows, denoising improves the results at both 10 Hz and 1 kHz. At 10 Hz, the improvement is particularly striking: the denoised estimate is now within the same order of magnitude as the 1-s reference, although it remains quite noisy. At 1 kHz, the effect of denoising is more moderate. Up to approximately 4 kA, little difference is observed between the denoised and nondenoised cases; however, as the circuit current increases toward 8 kA, the nondenoised estimate becomes increasingly noisy while the variance of the denoised estimate appears constant. This behavior is likely a result of the power converter noise increasing substantially within a 100 Hz bandwidth surrounding the 1 kHz stimulus frequency at higher currents (as shown previously in Fig. 9), with the denoising effectively suppressing these contributions.

For the 1-s estimates and the Z_{total} results, essentially no difference is seen between denoised and nondenoised data. However, this is not the case for the partition impedances, illustrated by Z_1 , where substantial improvement after denoising is observed for both the 10-Hz and 1-kHz cases.

This discrepancy is again likely due to the fact that when calculating Z_{total} , much of the common noise cancels since the majority of the measured power converter noise across the two magnet partitions is approximately 180° out of phase and thus cancels when summed.

Fig. 21 shows the impedance evolution over more than a thousand seconds, making it difficult to appreciate how the denoising impacts individual 10 ms estimates. To better visualize this effect, we zoom in on a 1-s window during the 8.23 kA plateau and examine the 1 kHz impedance, voltage, and current estimates. In Fig. 22, we plot both the denoised signals and their nondenoised counterparts, previously shown in Fig. 6.

Across all estimates, current, voltage, impedance of each partition, and impedance of the full magnet, we see a complete elimination of the sinusoidal out-of-phase amplitude modulation, present prior to applying denoising.

To summarize the impact of the denoising procedure across different operating conditions, Table I reports the relative variance, as defined in (24), of the impedance amplitude estimates at each stimulus frequency and plateau current level. The analysis is restricted to the plateau stages, as these represent quasi-stationary conditions where sufficient data are available to compute statistics with low uncertainty. For clarity, the table focuses on \hat{Z}_1 and \hat{Z}_{total} ; the corresponding results for \hat{Z}_2 and the individual voltage and current components are provided in Appendix B.

The table supports the trends observed in the time-domain impedance plots. For \hat{Z}_{total} , minimal variance reduction is achieved at 10 Hz and 110 Hz when using 1 s windows, regardless of whether static or adaptive filtering is applied. In fact, at the




Fig. 20. Zoomed power spectra of the current through Z_1 during the 2 kA plateau: (a) from 0 to 55 Hz, and (b) from 945 to 1055 Hz. Each plot compares the estimated spectra with no denoising (\hat{I}_{R1}), with static-coefficient denoising ($\hat{L}_{R1, \text{fixed}}$), and with adaptive denoising ($\hat{L}_{R1, \text{adapt}}$). Spectra are estimated using 19 segments of 400 000 samples each.

2 kA plateau, denoising slightly increases the variance in some cases. Only the 1-kHz component of \hat{Z}_{total} shows a consistent improvement, with relative variances reaching values between 0.39 and 0.011 times the original, depending on the conditions. In contrast, \hat{Z}_1 exhibits consistently reduced variance after denoising across all frequencies and plateau levels.

For the 10-ms estimates, denoising has a more pronounced effect on \hat{Z}_{total} across all frequencies. The comparison between static and adaptive filters reveals that their performance is generally similar, with the exception of the 5.5 kA plateau. In this case, the adaptive filter significantly outperforms the static one, indicating that it effectively compensates for changes in the system dynamics, such as changes in magnet inductance, occurring during powering.

VI. DISCUSSION

Building on the results presented in the previous section, we provide general observations and recommendations for future research.

A. Window Length

During steady-state (plateau) conditions, there is little reason to restrict window length, as longer windows improve the SNR and yield less noisy impedance estimates. This makes them well suited for slow health-monitoring applications and

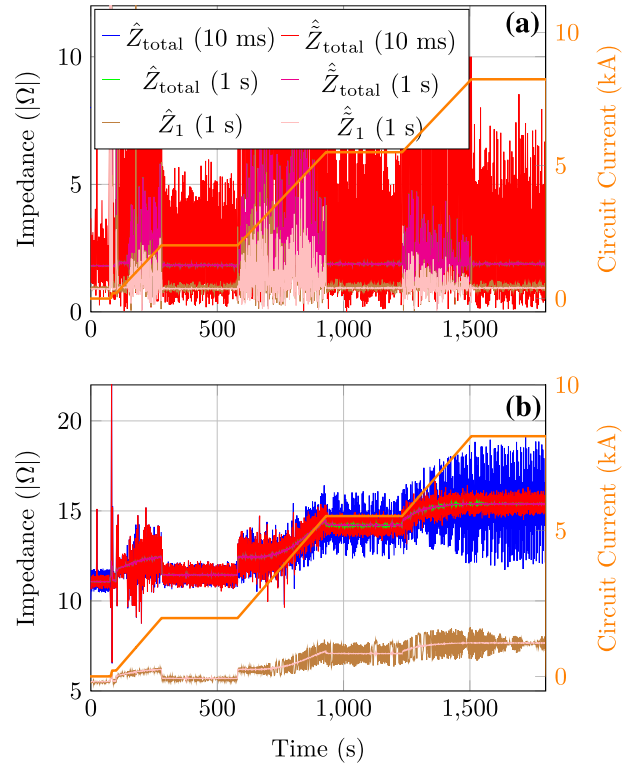


Fig. 21. Estimated impedance over time for (a) 10 Hz and (b) 1 kHz. Estimates are computed using 1 s windows (40 000 samples) and shown both with and without denoising for Z_{total} and Z_1 . 1 s estimates are also shown as a reference. For (a), Z_{total} without denoising at 10 ms (blue) is omitted to improve visualization.

trend analysis. In contrast, during ramping stages, the optimal window length depends strongly on the ramp rate and on how much temporal “smearing” of the impedance evolution can be tolerated. Shorter windows enable finer temporal resolution but result in increased variance, which can obscure subtle impedance changes. Applications targeting transient phenomena, such as the onset of quenches, are one of the use cases where the very short windows studied here (i.e., 10 ms) could be justified. However, achieving a tolerable SNR at these time scales remains a challenge, and further work is needed to improve the sensitivity of the measurement system and associated signal processing to make such studies practical.

B. Frequency Choice

The results clearly indicate that the power converter noise spectrum strongly influences the variance of the estimates. Future measurement campaign employing this approach should consider placing excitation frequencies in spectral regions with low converter noise. For example, using the power converter in the present study, this implies avoiding harmonics of the 50 Hz grid frequency. More generally, future implementations could benefit from automatic frequency-selection tools that characterize the noise spectrum in real time and identify suitable excitation frequencies. Such approaches could maximize SNR dynamically and ensure consistent measurement quality under varying operating conditions.




Fig. 22. Estimated (a) impedance, (b) voltage, and (c) current over time for the 1 kHz stimulus frequency. Estimates are computed using 10 ms windows (400 samples) over a 1-s interval during the 8.23 kA plateau stage. Shown with and without denoising applied.

C. Filter Implementations

This work applied Wiener filtering in postprocessing to establish a theoretical performance baseline. For real-time implementations, however, adaptive schemes, such as least-squares [30] or recursive least-squares [30] filtering, are likely more appropriate, as they can continuously track variations in the noise characteristics and system dynamics. Future research should explore these adaptive methods in combination with online impedance estimation to evaluate their performance, convergence stability, and computational feasibility in embedded control hardware.

TABLE I
RELATIVE VARIANCE AFTER DENOISING FOR THE THREE PLATEAUS

	Freq. (Hz)	Relative Variance After Denoising			
		1 s		10 ms	
		Static	Adapt	Static	Adapt
\hat{Z}_1	10	1.1e-02	1.1e-02	2.2e-04	2.1e-04
	110	1.4e-01	1.4e-01	2.3e-03	2.3e-03
	1 000	2.0e-02	2.1e-02	2.9e-02	2.8e-02
\hat{Z}_{total}	10	1.1e+00	1.1e+00	4.2e-01	4.1e-01
	110	1.0e+00	1.0e+00	7.3e-01	7.2e-01
	1 000	3.9e-01	3.9e-01	6.1e-01	6.1e-01

(a) Relative variance at 2 kA.

	Freq. (Hz)	Relative Variance After Denoising			
		1 s		10 ms	
		Static	Adapt	Static	Adapt
\hat{Z}_1	10	1.2e-02	1.2e-02	7.5e-04	8.9e-05
	110	1.6e-01	1.5e-01	6.4e-03	8.0e-04
	1 000	2.6e-03	1.1e-03	3.8e-03	2.1e-03
\hat{Z}_{total}	10	1.0e+00	1.0e+00	4.9e-01	5.8e-02
	110	1.0e+00	1.0e+00	1.2e+00	1.5e-01
	1 000	5.7e-02	2.2e-02	1.0e-01	5.3e-02

(b) Relative variance at 5.5 kA.

	Freq. (Hz)	Relative Variance After Denoising			
		1 s		10 ms	
		Static	Adapt	Static	Adapt
\hat{Z}_1	10	1.2e-02	1.2e-02	2.2e-05	2.2e-05
	110	5.4e-02	5.4e-02	3.9e-04	3.9e-04
	1 000	7.6e-04	7.7e-04	1.1e-03	1.1e-03
\hat{Z}_{total}	10	1.0e+00	1.0e+00	1.1e-01	1.1e-01
	110	9.8e-01	9.8e-01	4.7e-01	4.7e-01
	1 000	1.1e-02	1.1e-02	1.6e-02	1.6e-02

(c) Relative variance at 8.23 kA.

D. Flux Jumps

A key remaining challenge for continuous impedance monitoring during ramping is the presence of flux jumps. Their random occurrence and variable amplitude make them difficult to model and filter using conventional linear techniques. Future work should focus on developing methods to identify and mitigate the influence of flux jumps. One possible direction is to incorporate statistical models of flux-jump behavior, such as those discussed in [27], into the signal processing framework. These models could form the basis for model-based filtering schemes capable of distinguishing flux-jump transients from genuine impedance variations, thereby reducing the variance of impedance estimation during ramping.

VII. CONCLUSION

This article describes the progress made on continuous impedance measurements of superconducting magnets by introducing a new dataset and investigating the impact of active

power converters on noise levels in impedance estimates. We presented a Wiener filtering strategy designed to suppress this noise based on a reference signal acquired from the differential probing configuration. This article investigated two schemes: one based on a set of constant filters applied to the whole dataset and an adaptive approach where the coefficients were periodically recalculated.

The results show that these denoising methods substantially decreased the variance of the impedance estimates for individual magnet partitions. While the approach yields minimal benefit for full-magnet impedance estimates over 1 s estimation windows, it performs better when the window length is reduced to 10 ms. In that regime, we observe strong suppression of periodic noise artifacts, leading to stable estimates that were previously unusable.

However, at 10 Hz, particularly during ramping, flux jumps introduce large transients that remain unmitigated by this method. These pose a significant challenge for impedance estimation, limiting its viability as a tool for transient analysis and, for example, quench detection. Future work should explore detection and exclusion strategies to address them.

When comparing static and adaptive denoising, overall performance was similar across most operational conditions. However, at the 5.5 kA magnet current plateau, the adaptive method outperformed the static approach, with the relative variance reduced by up to a factor of 10 in this specific scenario. This highlights that adaptive filtering offers tangible benefits when system conditions vary notably. As our current adaptive method relies on noncausal segment-based calculation of filter coefficients, it is only suitable for offline applications, thus further research should aim toward testing real-time capable adaptive filtering approaches.

Finally, we emphasize that this study focused primarily on variance reduction. Due to the absence of independent reference measurements, we did not evaluate potential estimation biases introduced by the denoising methods. Addressing this limitation by establishing traceable references or performing systematic error analysis remains an interesting direction for future research.

APPENDIX A ANALYTICAL TRANSFER FUNCTION

To determine reasonable estimates for the transfer functions affecting the power converter output across the magnet, $D(s)$, across the magnet halves, $H^{ag}(s)$ and $H^{bg}(s)$, as well as across the reference resistors, $H^{lj}(s)$ and $H^{ek}(s)$, we analyze the circuit topology shown in Fig. 2.

Although discrete-time modeling is employed throughout the main text due to practical considerations in sampling and analysis, analytical derivations of transfer functions for the circuit topology are naturally conducted in the continuous-time Laplace (s) domain. This simplifies the analysis of circuit elements, especially reactive components, and their interrelations. The resulting continuous-time transfer functions can then be straightforwardly evaluated at discrete frequencies on the unit




Fig. 23. Simplified model after applying balance assumptions and superposition.




Fig. 24. Simplified schematic for analyzing $H^{lj}(s)$ and $H^{ek}(s)$.

circle, assuming that signals are sufficiently bandlimited and sampled without aliasing, ensuring direct correspondence with the discrete-time models used in this article.

We adopt the same balance assumptions as in the main text

$$\begin{aligned} Z_1 &= Z_2 \\ R_{I_1} &= R_{I_2} \\ R_{P_1} &= R_{P_2} \\ C_1 &= C_2 \\ U_1 &= U_2 \end{aligned}$$

As we are only interested in the transfer functions with respect to the power converter input U_{pc} , we can simplify the circuit via the superposition principle, replacing the stimulus sources U_1 and U_2 with short circuits. Furthermore, we assume that the power converter source impedance is negligible, i.e., $Z_{pc} = 0$. This yields the simplified schematic shown in Fig. 23.

TABLE II
VARIANCE OF IMPEDANCE ESTIMATES. DURATION: 1 S

	F	Raw	DN		DN Adapt	
			Abs	Rel	Abs	Rel
$\hat{\beta}^{ag}$	10	3.3e-09	3.8e-11	1.1e-02	3.8e-11	1.1e-02
	110	2.0e-10	2.7e-11	1.4e-01	2.7e-11	1.4e-01
	1000	2.3e-09	3.8e-11	1.6e-02	3.8e-11	1.6e-02
$\hat{\beta}^{bg}$	10	3.5e-09	3.8e-11	1.1e-02	3.8e-11	1.1e-02
	110	2.0e-10	2.7e-11	1.4e-01	2.7e-11	1.4e-01
	1000	2.4e-09	3.9e-11	1.6e-02	3.9e-11	1.6e-02
$\hat{\beta}^{lj}$	10	9.1e-14	8.9e-14	9.7e-01	9.0e-14	9.8e-01
	110	4.7e-14	3.3e-14	7.1e-01	3.4e-14	7.2e-01
	1000	2.0e-12	1.0e-13	5.2e-02	1.0e-13	5.1e-02
$\hat{\beta}^{ek}$	10	9.1e-14	1.4e-13	1.5e+00	1.4e-13	1.5e+00
	110	4.7e-14	3.9e-14	8.3e-01	3.9e-14	8.2e-01
	1000	2.0e-12	1.2e-13	5.9e-02	1.2e-13	5.9e-02
\hat{Z}_1	10	1.1e-02	1.2e-04	1.1e-02	1.2e-04	1.1e-02
	110	5.3e-04	7.2e-05	1.4e-01	7.2e-05	1.4e-01
	1000	6.5e-03	1.3e-04	2.0e-02	1.3e-04	2.1e-02
\hat{Z}_2	10	1.1e-02	1.3e-04	1.2e-02	1.3e-04	1.2e-02
	110	5.2e-04	7.7e-05	1.5e-01	7.7e-05	1.5e-01
	1000	1.0e-02	1.6e-04	1.6e-02	1.6e-04	1.6e-02
\hat{Z}_{total}	10	4.8e-04	5.1e-04	1.1e+00	5.0e-04	1.1e+00
	110	2.9e-04	3.0e-04	1.0e+00	3.0e-04	1.0e+00
	1000	1.4e-03	5.5e-04	3.9e-01	5.5e-04	3.9e-01

(a) Variance at 2 kA (1s window).

$\hat{\beta}^{ag}$	10	2.9e-09	3.3e-11	1.1e-02	3.3e-11	1.1e-02
	110	3.5e-10	5.2e-11	1.5e-01	5.1e-11	1.5e-01
	1000	7.4e-08	1.7e-10	2.3e-03	4.7e-11	6.3e-04
$\hat{\beta}^{bg}$	10	2.7e-09	3.3e-11	1.2e-02	3.3e-11	1.2e-02
	110	3.0e-10	5.2e-11	1.7e-01	5.1e-11	1.7e-01
	1000	6.4e-08	1.7e-10	2.6e-03	4.7e-11	7.3e-04
$\hat{\beta}^{lj}$	10	1.7e-14	1.5e-14	9.1e-01	1.5e-14	9.2e-01
	110	4.6e-14	3.1e-14	6.8e-01	3.2e-14	6.8e-01
	1000	6.3e-11	1.9e-13	3.0e-03	1.1e-13	1.8e-03
$\hat{\beta}^{ek}$	10	1.7e-14	2.0e-14	1.2e+00	2.0e-14	1.2e+00
	110	4.6e-14	2.5e-14	5.3e-01	2.5e-14	5.3e-01
	1000	6.3e-11	2.0e-13	3.2e-03	7.6e-14	1.2e-03
\hat{Z}_1	10	9.1e-03	1.1e-04	1.2e-02	1.1e-04	1.2e-02
	110	9.3e-04	1.4e-04	1.6e-01	1.4e-04	1.5e-01
	1000	1.9e-01	4.8e-04	2.6e-03	2.1e-04	1.1e-03
\hat{Z}_2	10	8.6e-03	1.1e-04	1.3e-02	1.1e-04	1.3e-02
	110	8.0e-04	1.5e-04	1.8e-01	1.4e-04	1.8e-01
	1000	2.8e-01	4.2e-04	1.5e-03	1.5e-04	5.3e-04
\hat{Z}_{total}	10	4.2e-04	4.4e-04	1.0e+00	4.4e-04	1.0e+00
	110	5.7e-04	5.8e-04	1.0e+00	5.7e-04	1.0e+00
	1000	3.1e-02	1.8e-03	5.7e-02	6.8e-04	2.2e-02

(b) Variance at 5.5 kA (1s window).

$\hat{\beta}^{ag}$	10	3.0e-09	3.5e-11	1.2e-02	3.5e-11	1.2e-02
	110	7.7e-10	4.0e-11	5.2e-02	4.0e-11	5.2e-02
	1000	4.7e-08	2.4e-11	5.1e-04	2.4e-11	5.2e-04
$\hat{\beta}^{bg}$	10	3.2e-09	3.6e-11	1.1e-02	3.6e-11	1.1e-02
	110	8.1e-10	4.1e-11	5.0e-02	4.1e-11	5.0e-02
	1000	4.7e-08	2.5e-11	5.3e-04	2.5e-11	5.3e-04
$\hat{\beta}^{lj}$	10	2.6e-14	2.1e-14	8.2e-01	2.1e-14	8.3e-01
	110	6.8e-14	1.9e-14	2.9e-01	1.9e-14	2.9e-01
	1000	4.0e-11	1.5e-14	3.7e-04	1.6e-14	4.0e-04
$\hat{\beta}^{ek}$	10	2.6e-14	2.0e-14	7.8e-01	2.0e-14	7.8e-01
	110	6.8e-14	2.4e-14	3.5e-01	2.4e-14	3.5e-01
	1000	4.0e-11	2.2e-14	5.4e-04	2.3e-14	5.6e-04
\hat{Z}_1	10	9.7e-03	1.2e-04	1.2e-02	1.2e-04	1.2e-02
	110	2.0e-03	1.1e-04	5.4e-02	1.1e-04	5.4e-02
	1000	1.1e-01	8.6e-05	7.6e-04	8.7e-05	7.7e-04
\hat{Z}_2	10	1.0e-02	1.2e-04	1.2e-02	1.2e-04	1.2e-02
	110	2.2e-03	1.1e-04	5.3e-02	1.1e-04	5.3e-02
	1000	2.1e-01	9.0e-05	4.4e-04	9.1e-05	4.4e-04
\hat{Z}_{total}	10	4.7e-04	4.7e-04	1.0e+00	4.7e-04	1.0e+00
	110	4.6e-04	4.5e-04	9.8e-01	4.5e-04	9.8e-01
	1000	3.1e-02	3.4e-04	1.1e-02	3.5e-04	1.1e-02

(c) Variance at 8.23 kA (1s window).

TABLE III
VARIANCE OF IMPEDANCE ESTIMATES. DURATION: 10 MS

	F	Raw	DN		DN Adapt	
			Abs	Rel	Abs	Rel
$\hat{\beta}^{ag}$	10	1.6e-04	4.9e-08	3.0e-04	4.7e-08	2.9e-04
	110	2.5e-06	5.6e-09	2.3e-03	5.6e-09	2.3e-03
	1000	1.4e-07	3.5e-09	2.5e-02	3.5e-09	2.5e-02
$\hat{\beta}^{bg}$	10	1.6e-04	5.3e-08	3.2e-04	5.1e-08	3.1e-04
	110	2.5e-06	5.7e-09	2.3e-03	5.7e-09	2.3e-03
	1000	1.4e-07	3.6e-09	2.5e-02	3.6e-09	2.5e-02
$\hat{\beta}^{lj}$	10	1.7e-08	6.7e-11	3.9e-03	8.4e-11	4.9e-03
	110	7.4e-10	5.6e-12	7.6e-03	5.9e-12	8.0e-03
	1000	7.6e-11	5.4e-12	7.2e-02	5.3e-12	7.0e-02
$\hat{\beta}^{ek}$	10	1.7e-08	7.5e-10	4.4e-02	7.7e-10	4.5e-02
	110	7.4e-10	8.3e-12	1.1e-02	8.5e-12	1.1e-02
	1000	7.6e-11	5.9e-12	7.8e-02	5.9e-12	7.8e-02
\hat{Z}_1	10	7.4e+02	1.6e-01	2.2e-04	1.5e-01	2.1e-04
	110	6.7e+00	1.5e-02	2.3e-03	1.5e-02	2.3e-03
	1000	4.4e-01	1.2e-02	2.9e-02	1.2e-02	2.8e-02
\hat{Z}_2	10	7.5e+02	1.7e-01	2.3e-04	1.7e-01	2.2e-04
	110	6.7e+00	1.6e-02	2.4e-03	1.6e-02	2.4e-03
	1000	5.3e-01	1.4e-02	2.6e-02	1.4e-02	2.6e-02
\hat{Z}_{total}	10	1.6e+00	6.6e-01	4.2e-01	6.4e-01	4.1e-01
	110	8.6e-02	6.2e-02	7.3e-01	6.2e-02	7.2e-01
	1000	8.3e-02	5.1e-02	6.1e-01	5.1e-02	6.1e-01

(a) Variance at 2 kA (10ms window).

$\hat{\beta}^{ag}$	10	9.7e-04	1.4e-06	1.4e-03	1.7e-07	1.7e-04
	110	1.1e-05	6.9e-08	6.5e-03	8.4e-09	7.9e-04
	1000	2.4e-06	8.0e-09	3.3e-03	3.7e-09	1.5e-03
$\hat{\beta}^{bg}$	10	8.5e-04	1.4e-06	1.6e-03	1.7e-07	2.0e-04
	110	9.2e-06	6.9e-08	7.5e-03	8.5e-09	9.2e-04
	1000	2.0e-06	8.0e-09	4.0e-03	3.7e-09	1.8e-03
$\hat{\beta}^{lj}$	10	2.6e-08	1.1e-10	4.1e-03	5.9e-11	2.3e-03
	110	1.0e-09	6.0e-12	6.0e-03	4.1e-12	4.1e-03
	1000	1.9e-09	8.3e-12	4.3e-03	5.0e-12	2.6e-03
$\hat{\beta}^{ek}$	10	2.6e-08	7.7e-10	3.0e-02	7.3e-10	2.8e-02
	110	1.0e-09	7.7e-12	7.7e-03	5.9e-12	5.9e-03
	1000	1.9e-09	7.4e-12	3.9e-03	4.0e-12	2.1e-03
\hat{Z}_1	10	6.0e+03	4.5e+00	7.5e-04	5.4e-01	8.9e-05
	110	2.8e+01	1.8e-01	6.4e-03	2.3e-02	8.0e-04
	1000	6.5e+00	2.5e-02	3.8e-03	1.3e-02	2.1e-03
\hat{Z}_2	10	5.2e+03	4.7e+00	9.0e-04	5.5e-01	1.1e-04
	110	2.5e+01	1.9e-01	7.4e-03	2.3e-02	9.3e-04
	1000	8.7e+00	2.4e-02	2.7e-03	1.3e-02	1.5e-03
\hat{Z}_{total}	10	3.8e+01	1.8e+01	4.9e-01	2.2e+00	5.8e-02
	110	6.3e-01	7.4e-01	1.2e+00	9.2e-02	1.5e-01
	1000	9.5e-01	9.6e-02	1.0e-01	5.1e-02	5.3e-02

(b) Variance at 5.5 kA (10ms window).

$\hat{\beta}^{ag}$	10	1.7e-03	1.0e-07	6.1e-05	1.0e-07	6.1e-05
	110	1.8e-05	7.0e-09	3.8e-04	7.0e-09	3.8e-04
	1000	4.8e-06	3.8e-09	7.9e-04	3.8e-09	7.9e-04
$\hat{\beta}^{bg}$	10	1.7e-03	1.1e-07	6.2e-05	1.1e-07	6.2e-05
	110	1.9e-05	7.1e-09	3.8e-04	7.1e-09	3.8e-04
	1000	4.7e-06	3.9e-09	8.3e-04	3.9e-09	8.3e-04
$\hat{\beta}^{lj}$	10	3.0e-08	6.1e-11	2.0e-03	6.2e-11	2.1e-03
	110	1.2e-09	4.8e-12	4.1e-03	4.8e-12	4.1e-03
	1000	4.2e-09	4.1e-12	9.8e-04	4.2e-12	1.0e-03
$\hat{\beta}^{ek}$	10	3.0e-08	7.3e-10	2.4e-02	7.3e-10	2.4e-02
	110	1.2e-09	6.3e-12	5.4e-03	6.3e-12	5.5e-03
	1000	4.2e-09	4.4e-12	1.1e-03	4.5e-12	1.1e-03
\hat{Z}_1	10	1.5e+04	3.4e-01	2.2e-05	3.4e-01	2.2e-05
	110	5.0e+01	1.9e-02	3.9e-04	1.9e-02	3.9e-04
	1000	1.2e+01	1.3e-02	1.1e-03	1.4e-02	1.1e-03
\hat{Z}_2	10	1.5e+04	3.5e-01	2.3e-05	3.5e-01	2.3e-05
	110	5.1e+01	2.0e-02	3.9e-04	2.0e-02	3.9e-04
	1000	2.1e+01	1.4e-02	6.5e-04	1.4e-02	6.5e-04
\hat{Z}_{total}	10	1.3e+01	1.4e+00	1.1e-01	1.4e+00	1.1e-01
	110	1.7e-01	7.8e-02	4.7e-01	7.8e-02	4.7e-01
	1000	3.4e+00	5.3e-02	1.6e-02	5.3e-02	1.6e-02

(c) Variance at 8.23 kA (10ms window).

Under these conditions, it is trivial to determine by inspection that

$$D(s) = 1$$

$$H^{ag}(s) = 0.5$$

$$H^{bg}(s) = -0.5.$$

To determine the remaining transfer functions $H^{lj}(s)$ and $H^{ek}(s)$, we first observe that, by symmetry, the two functions have equal magnitude and opposite sign (180° phase difference).

It suffices to analyze one case. The corresponding simplification is shown in Fig. 24.

It is clear that the relevant circuit is a simple RC high-pass filter. The transfer function between nodes l and j is

$$H^{lj}(s) = -0.5 \frac{R_{I_1} C_1 s}{1 + (R_{P_1} + R_{I_1}) C_1 s} \quad (30)$$

and similarly, the transfer function between nodes e and k is

$$H^{ek}(s) = 0.5 \frac{R_{I_2} C_2 s}{1 + (R_{P_2} + R_{I_2}) C_2 s}. \quad (31)$$

Since the system signals are bandlimited and sampled at a sufficiently high rate, the discrete-time transfer functions used in the estimation process are obtained directly by evaluating the continuous-time expressions at the sampled frequencies. Thus, for the purposes of discrete-time modeling, the corresponding z -domain representations $D(f)$, $H^{ag}(f)$, $H^{bg}(f)$, $H^{lj}(f)$, and $H^{ek}(f)$ are equivalent to their s -domain forms evaluated on the unit circle $z = e^{j\omega T_s}$.

APPENDIX B RELATIVE VARIANCE TABLES

These tables list the estimated absolute and relative variances of the voltage estimates ($\hat{\beta}^{ag}$ and $\hat{\beta}^{bg}$), the currents ($\hat{\beta}^{lj}$ and $\hat{\beta}^{ek}$), the impedance estimates of the two partitions (\hat{Z}_1 and \hat{Z}_2), and the total magnet impedance (\hat{Z}_{total}).

The variances are listed for each of the three plateau stages. Table II details results based on the 1-s estimation windows, while Table III references from the 10-ms windows.

REFERENCES

- [1] O. Bruning and M. Zerlauth, "Overall status of the HL-LHC project," *JACoW IPAC*, vol. 2023, 2023, Art. no. TUYG1. doi: [10.18429/JACoW-IPAC2023-TUYG1](https://cds.cern.ch/record/2886319). [Online]. Available: <https://cds.cern.ch/record/2886319>
- [2] M. B. B. Christensen et al., "Enabling real-time impedance measurements of operational superconducting circuits of accelerator magnets," *IEEE Trans. Appl. Supercond.*, vol. 34, no. 5, Aug. 2024, Art. no. 4004705.
- [3] T. Podzorny et al., "Continuous diagnostics for powered superconducting circuits," *IEEE Trans. Appl. Supercond.*, vol. 35, no. 5, Aug. 2025, Art. no. 4700105.
- [4] M. Christensen, M. Bednarek, P. Koch, T. Podzorny, J. Steckert, and J. Østergaard, "Enabling impedance measurements of energised superconducting circuits through differential probing," *J. Instrum.*, vol. 20, no. 05, May 2025, Art. no. P05042.
- [5] M. Janitschke, E. Ravaoli, A. P. Verweij, G. Willering, and U. van Rienen, "Complex impedance of the LHC main dipole magnet in the presence of artificial short circuits," *IEEE Trans. Appl. Supercond.*, vol. 35, no. 5, 2025, Art. no. 4002905, [Online]. Available: <https://cds.cern.ch/record/2932588>
- [6] V. Reynaud et al., "Evaluating the electro-magnetic effects of electrical short-circuits in a Nb-Ti accelerator magnet," *IEEE Trans. Appl. Supercond.*, vol. 35, no. 5, Aug. 2025, Art. no. 4700405.
- [7] H. Ehmler et al., "Comparative analysis of impulse and impedance tests to detect short circuits within the W7-X magnets," *IEEE Trans. Appl. Supercond.*, vol. 16, no. 2, pp. 767–770, Jun. 2006.
- [8] R. Bailey Ed., *CAS - CERN Accelerator School: Power Converters*, *CERN Yellow Report*, Geneva, Switzerland, 2015, p. 460, doi: [10.5170/CERN-2015-003](https://cds.cern.ch/record/1641409). [Online]. Available: <http://cds.cern.ch/record/1641409>
- [9] O. S. Bruning, A. Butterworth, M. Lamont, and J. Wenninger, "The Nominal Operational Cycle of the LHC with Beam (version 1)," CERN, Geneva, Switzerland, Tech. Rep., 2003. [Online]. Available: <http://cds.cern.ch/record/691955>
- [10] A. Chmielińska, L. Fiscarelli, M. Hostettler, W. Kozanecki, S. Russenschuck, and E. Todesco, "Magnetization in superconducting corrector magnets and impact on luminosity–calibration scans in the Large Hadron Collider," *Eur. Phys. J. Plus*, vol. 138, no. 9, 2023, Art. no. 796.
- [11] N. Wiener, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series: With Engineering Applications*. Cambridge, MA, USA: MIT Press, Aug. 1949.
- [12] J. R. Raggazzini and G. F. Franklin, *Sampled-Data Control Systems, 1st ed.* Berlin, Germany: McGraw-Hill Book Co., 1958.
- [13] H. Bajas et al., "Test result of the short models MQXF53 and MQXF55 for the HL-LHC upgrade," *IEEE Trans. Appl. Supercond.*, vol. 28, no. 3, Apr. 2018, Art. no. 4007006.
- [14] M. Janitschke, M. Bednarek, J. Ludwin, E. Ravaoli, A. P. Verweij, and U. van Rienen, "Analyzing the complex impedances of all LHC main dipole magnets," *IEEE Trans. Appl. Supercond.*, vol. 35, no. 5, Aug. 2025, Art. no. 4603706.
- [15] R. Shafer, "Eddy currents, dispersion relations, and transient effects in superconducting magnets," *IEEE Trans. Magn.*, vol. M-17, no. 1, pp. 722–725, Jan. 1981.
- [16] G. Kirby et al., "Models and experimental results from the wide aperture Nb-Ti magnets for the LHC upgrade," in *contribution to WAMSDO 2013: Workshop on Accelerator Magnet, Superconductor, Design and Optimization; 15–16 Jan. 2013*, Geneva, Switzerland: CERN, 2013, pp. 57–64. [Online]. Available: <http://cds.cern.ch/record/1643440>
- [17] V. Marinuzzi, M. Sorbi, G. Manfreda, F. Bellina, H. Bajas, and G. Chlachidze, "Effect of coupling currents on the dynamic inductance during fast transient in superconducting magnets," *Phys. Rev. ST Accel. Beams*, vol. 18, Mar. 2015, Art. no. 032401. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevSTAB.18.032401>
- [18] P. Stoica, H. Li, and J. Li, "Amplitude estimation of sinusoidal signals: Survey, new results, and an application," *IEEE Trans. Signal Process.*, vol. 48, no. 2, pp. 338–352, Feb. 2000.
- [19] F. Kuo, *Network Analysis and Synthesis*, 2nd ed. New York, NY, USA: Wiley, Dec. 1966.
- [20] S. M. Kay, *Modern Spectral Estimation: Theory and Application*. Englewood Cliffs, NJ, USA: Prentice Hall, 1988.
- [21] P. Ferracin et al., "The HL-LHC Low- β quadrupole magnet MQXF: From short models to long prototypes," *IEEE Trans. Appl. Supercond.*, vol. 29, no. 5, Aug. 2019, Art. no. 032401.
- [22] A. Schaeffer, "The SM18 test facility in the HL-LHC era," Online, Dec. 2022. [Online]. Available: <https://home.cern/news/news/engineering/sm18-test-facility-hl-lhc-era>
- [23] J. Steckert, R. Denz, S. Mundra, T. Podzorny, J. Spasic, and D. G. Vancea, "Application of the universal quench detection system to the protection of the high-luminosity LHC magnets at CERN," *IEEE Trans. Appl. Supercond.*, vol. 32, no. 6, Sep. 2022, Art. no. 4006305.
- [24] F. J. Mangiarotti et al., "Performance of a HL-LHC Nb₃Sn Quadrupole Magnet in the 100–200MPa Range of Azimuthal Stress," *IEEE Trans. Appl. Supercond.*, vol. 35, no. 5, 2025, Art. no. 4000307.
- [25] M. B. B. Christensen et al., "Impedance characterization of a HL-LHC Nb₃Sn MQXF5 model magnet during powered operation at nominal current," *IEEE Trans. Appl. Supercond.*, p. 8, 2025, under review.
- [26] M. N. Wilson, *Superconducting Magnets, Ser. Monographs on Cryogenics*. Oxford, U.K.: Clarendon Press, Mar. 1987.
- [27] M. Martino, P. Arpaia, and S. Ierardi, "Impact of flux jumps on high-precision powering of Nb₃Sn superconducting magnets," *J. Phys.: Conf. Ser.*, vol. 1350, 2019, Art. no. 012180. [Online]. Available: <https://cds.cern.ch/record/2694057>
- [28] P. Welch, "The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms," *IEEE Trans. Audio Electroacoust.*, vol. AC-15, no. 2, pp. 70–73, Jun. 1967.
- [29] G. Willering et al., "Flux jumps in Nb₃Sn magnets—Types and characterization: Flux Jump characteristics in MBH coils and a possible link to 'spikes' in MBHA 001," Presentation at the Special Meeting on the Results of the Cold Tests of MBHA-001 – Update 10, CERN, p. 28, Apr. 2020, Accessed: Jun. 26, 2025. [Online]. Available: <https://indico.cern.ch/event/906182/>
- [30] S. Haykin, *Adaptive Filter Theory: International Edition*, 5th ed. London, U.K.: Pearson Education, Jul. 2013, p. 912.
- [31] M. H. J. Gruber, *Improving Efficiency by Shrinkage, Ser. Statistics: A Series of Textbooks and Monographs*. Boca Raton, FL, USA: CRC Press, Feb. 1998, p. 648.