Understanding the Inception of 14.7 Hz Oscillations Emerging from a Data Center

Chetan Mishra[a,*], Luigi Vanfretti[b], Jaime Delaree Jr.[a], T.J. Purcell[a], Kevin D. Jones[a]

[a]*Dominion Energy, Glen Allen, VA 23060, USA*
[b]*Dept. of ECSE, Rensselaer Polytechnic Institute, Troy, NY 12180, USA*

Abstract

As Dominion Energy's power grid has grown to include an increasing number of inverter-based resources and power electronic components, we have begun to observe an increase in previously unforeseen and difficult-to-explain dynamic behaviors. In this paper, we present a measurement-based analysis of 14.7-14.8 Hz oscillations emerging from a data center in Dominion Energy's power grid. Synchrophasor data collected over several months were analyzed using non-parametric spectral analysis approaches to understand its inception and, more importantly, to characterize its behavior. What makes this case unique is the presence of an independent and unplanned exogenous excitation in the system. This excitation plays a critical role in understanding the inception and emergence of an otherwise unobservable dynamic process that, as it destabilizes, results in oscillations at a higher and narrow frequency range than that of conventional power system dynamics.

© 2017 Elsevier Inc. All rights reserved.
Keywords: Forced oscillations, load dynamics, spectral analysis

1. Introduction

The growth of high-concentration large-scale commercial and industrial loads, including data centers, in Dominion Energy's system has increased on a considerable scale over the last decade. Furthermore, the concentration of such loads in populous regions, which is generally limited in local synchronous machine-based generation, poses challenges that should be addressed not only by the utility but also by large-scale commercial and industrial customers. Such challenges show the need for a coordinated and concerted effort to ensure that reliability is prioritized as more critical loads rely on it. Industrial loads such as data centers are not the only drivers, other inverter-based resources (IBRs, e.g. solar sites) and companies with investments in sensitive electronic manufacturing are just a few of those who could be impacted. The analysis of an oscillation emerging from a data center presented in this paper is a cautionary tale about how utilities and customers must work together to remain vigilant.

Challenges with data centers are especially of concern in areas that lack the tight interconnection of a 500KV system and instead operate in areas ripe for diluted reliability conditions at lower voltages, which can be exacerbated by weakened conditions such as shoulder-month maintenance [1]. The interconnection of large-scale loads in relatively weak parts of the transmission system can cause local controllers to respond to dynamic performance irregularities. When controllers tuned for nominal operating conditions fail to account for significant changes on the grid, they can and have responded erratically. This has been observed in Dominion Energy's grid manifest as unwanted grid-to-data center controller interactions, as discussed in this paper, resonance across IBR controllers [2], and other behaviors [3].

* Corresponding author.

A real-world example that motivates the work in this paper is presented in Fig. 2, which shows high-frequency undamped oscillations that appeared in a region rich in data centers in Dominion Energy's power system (a key use case of the work analyzed in this paper). Due to the nascent nature of observing such a phenomenon, operators and utilities alike are still learning what to analyze, how to monitor and how to optimally respond to such instances. In the case of this 'event', operators remained cautious and stoic but were unable to respond proactively. Fortunately, the oscillation in question lasted only two hours before organically disappearing, as it naturally damped from system changes.

Analyzing and modeling this behavior is difficult due to lackluster tools and visibility available to key decision makers. The inability to access accurate transparent models involving customer power electronics means we, the utility, are forced to make educated guesses about what is behind the meter. For data centers, this manifests itself when Uninterruptible Power Supplies (UPS) provides frequency response services [4] that include complex power electronic controls [5] designed to protect the sensitive components housed within their server farms. However, we can only guess how such controllers are configured, leaving measurement data as the utility's only option for analysis and monitoring. In addition, the composition of a service region, i.e., the types of dynamic devices connected at a specific site and their operational schedule, may be completely unknown in worst-case scenarios; this is representative of the situation discussed in this paper.

It is worth noting that in the context of data centers, most of the literature is concerned with exploring how to meet the growing power demand for data centers from the point of view of planning and their operation. In the case of planning, data center power supply is being revisited due to regulatory and policy pressures, sustainability targets and ancillary service opportunities [6]. These catalysts are bringing operators of data centers to explore means to improve their energy efficiency [7], [8] and alternative back-up power sources to diesel-power systems, including renewable energy sources in the form of nearby distributed energy sources (DERs) [9] or from longer distances [10]. Moreover, a great effort has been made in finding ways to operate data centers in order to reduce their carbon footprint[11] and [12]. However, understanding the dynamic performance of data centers from real-world field measurements remains largely unaddressed in the literature, and this paper aims to bridge this gap.

To gain insight into the dynamic performance and operation of data centers, it is possible to leverage existing approaches used to understand power system dynamics, such as oscillation monitoring, that have proven successful in real-world applications. Detecting oscillations with low-damping, especially once they become unstable and cause sustained oscillations, is a well-studied area [13]. The techniques used to pinpoint 'events' seek large increases in measured spectral power across specific frequency bands – this allows for the detection of *known* phenomena. However, this approach poses a significant shortcoming, as by the time an oscillation is identified and detected, it may not be possible to mitigate it, resulting in damage or a decrease in quality of service. To address oscillations as the ones shown in Figure 2, the system operators are left with only a few options and may feel compelled to take drastic measures, such as taking the offending asset offline or applying actions that may or may not resolve the issue. Meanwhile, addressing these issues post-facto, can result in necessary but high-cost solutions in the form of STATCOMS and SVC's [14]. Although they alleviate the symptoms and are essential to meet expected reliability standards, they do not address the underlying cause. Fortunately, in most cases of oscillations caused by unstable controllers, there are hints during the equipment's normal operation (ambient conditions), indicating that the controller is close to or on the verge of instability [15]. Therefore, identifying and promptly correcting emergent dynamic issues requires a shift in conventional engineering approaches, using for example the approach presented in this paper. The basis of our methodology is to screen through what has often been referred to as "noise," and use ambient data to intently identify unfavorable dynamics long before they occur [7].

The fundamental difficulty in working with ambient data is the stochastic nature of the system, which results from various uncertain and persistent inputs to the system [18]. In this regard, spectral analysis techniques [19] provide effective means to characterize the underlying system dynamics from ambient measurements (owing much to the nearly linear response of the system) and therefore have been widely used in mode estimation [18], [20], [21]. It's worth interjecting our concern here behind the broad assumption behind mode metering applications: while some oscillations (typically wide-area modes) operate under known frequency bands, this may not be the case for real-world control modes of converter-interfaced apparatuses [16], which makes the use of time-frequency analysis techniques necessary [22]. In our experience, time-frequency analysis of the time window of the inception of converter-interfaced control modes yields unique insights into the oscillation mechanism, as will be demonstrated in this work.

The remainder of this paper is organized as follows; Section 2, gives necessary background on signal processing

approaches utilized across this corpus of work and explains how they are used to gain insight from synchrophasor data. Section 3 provides an overview of a real-world oscillation event, as well as the region of interest. Section 4 performs analyses using the measurements. The goal of this work is not to share specific examples of observed phenomena, but rather to provide an evolving set of guidelines for performing data-driven analysis of forced oscillations in situations where model scarcity makes such events difficult to diagnose. Finally, Section 5 summarizes the main conclusions of this work.

2. Time Frequency Analysis Background

In this section we review key concepts used in the time-frequency analysis techniques that will be used to analyze measurement data.

2.1. Spectrum of Power System Ambient Dynamics

For a stochastic process $x(t)$, power spectral density denoted by $S(f)$ is defined as the distribution of average signal power over frequency,

$$\lim_{T\to\infty} E\left(\int_0^T \frac{x^2(t)dt}{T}\right) = \int S(f)df \tag{1}$$

where, $E(.)$ is the expectation operator. Processes with time-invariant mean and correlation functions are called wide-sense stationary and are characterized by a time invariant $S(f)$.

Although for most power system ambient dynamics (especially electromechanical modes) satisfy this assumption, it may not hold true in other dynamic processes contained in real-world data. The power system's time-varying nature results from the changing operating conditions that are caused by changes in load, generation dispatch, etc., thereby making it a nonstationary process. Therefore, a more general characterization using a time-varying linear system driven by a stationary white noise input $u(t)\sim N(0,1)$ is defined by

$$x(t) = \int h(t,\tau)u(t-\tau)d\tau = \int e^{j2\pi ft}H(t,f)dz(f) \tag{2}$$

where, $z(f)$ is a zero mean process with orthogonal increments, i.e., $E\big(dz(f)\big) = 0\ \forall f$ and $E\big(dz(f)dz^*(f_1)\big) = S_u(f)\delta(f-f_1)dfdf_1 \forall (f,f_1)$ [23]. The PSD of $u(t)$ is denoted by $S_u(f)$ and is equal to a constant (equal to 1 here) for white noise. $h(t,\tau)$ represents the time varying impulse response and is defined as response at time t for impulse input at $t-\tau$, $H(t,f) = \int h(t,\tau)e^{-j2\pi f\tau}d\tau$.

For nonstationary processes, a time varying notion for the spectrum, called the evolution spectrum (denoted by $S(t,f)$), was introduced in [23], for which the above process equals $|H(t,f)|^2$. Here, the most prominent dynamics of the underlying system appear as peaks (local maxima) in $S(t,f)$, which we are interested in capturing. Note that the discrete system changes, e.g. a change in network topology, are considered by assuming $S(t,f)$ to be piecewise continuous in t.

2.2. Time-Frequency Analysis using the Wavelet Transform

Due to the time-varying nature of real-world power systems [16], it is natural to utilize time-frequency analysis [10]. These techniques operate by decomposing the measurement signal into several constituent components, which (in general) correspond to distinct underlying phenomena in an actual system. In the linear class of techniques, the signal is characterized by inner products with a pre-assigned family of templates, which are normally skewed in the time and frequency plane.

In the wavelet transform [24], a pre-defined component called the 'mother wavelet' $\psi(t)$ is translated and scaled to create family of templates called 'wavelets'. It is chosen in accordance with the dominant behavior in the measurements to be analyzed, thus providing great flexibility. Under the assumption that slow processes generally vary less frequently than fast ones, frequency resolution is increasingly favored over time resolution at lower frequencies. The Continuous Wavelet Transform (CWT) coefficient $W(a,b)$ for a signal $x(t)$ is obtained as,

$$W_x(a,b) = \int x(t) a^{-\frac{1}{2}} \bar{\psi}\left(\frac{t-b}{a}\right) dt \tag{3}$$

where, a is the scale factor, b is the shift factor and $\bar{}$ is the complex conjugate operator. Increasing a has the effect of stretching $\psi(t)$ (shift towards lower frequency), while b introduces a phase shift in time. In practice, the CWT is realized for sampled data by discretizing it as $a = 2^{\frac{j}{\nu}}$ for $j = 1,2,3 \ldots.$, ν is an integer > 1 called number of voices per octave and b takes integer values. Note that the CWT is a highly redundant transformation, that is, a signal x of N samples is converted to $M \times N$ wavelet coefficients with M being the number of scales. Note that this is needed for better characterization and localization, as illustrated later.

The mother wavelet $\psi(t)$ is designed to have compact time and frequency support. This means that translating and scaling is akin to moving around a band pass filter paired with a windowing function. In the present work, we use the Morlet wavelet, which consists of a complex sinusoid localized with a Gaussian window,

$$\psi(t) = e^{-\frac{2\pi^2 t^2}{k_0^2}} \times \left(e^{j2\pi t} - e^{-\frac{k_0^2}{2}}\right).$$

A major challenge in time-frequency analysis is to make a suitable trade-off between time and frequency resolutions, otherwise known as the Heisenberg's uncertainty principle. Note that the CWT coefficients $W_x(a,b)$ are in a time-scale plane and not the time-frequency plane, where a single scale a is associated with a range of frequencies. To convert it to time-frequency representation, the synchrosqueezing wavelet transform was proposed in [25], which is a post-processing approach that allocates the wavelet coefficient $W_x(a,b)$ to points in time-frequency plane. This is achieved by estimating the dominant frequency $\omega_x(a,b)$ from $W_x(a,b)$ to which $W_x(a,b)$ value is allocated, resulting in

$$\omega_x(a,b) = -i\big(W_x(a,b)\big)^{-1} \partial_b W_x(a,b). \tag{4}$$

3. Case Study Description

3.1. Study System

The region of Dominion Energy's power grid under study is shown in Fig. 1 and is predominantly a 115 kV network rich in converter interfaced resources. A cluster of data centers at substations D1, D2, and D3 lies at the center of this region. Substations H and C have solar photovoltaic (PV) generation with 85 MW and 50 MW rating, respectively. A 267 MW hydro power plant is located at substation K, which serves as an important source of fault current. The nearest connection to higher kV levels is through a 115/230 kV transformer at substation C. This connects directly to a nearby 500/230 kV substation, which contains a ±125 MVAR STATCOM.

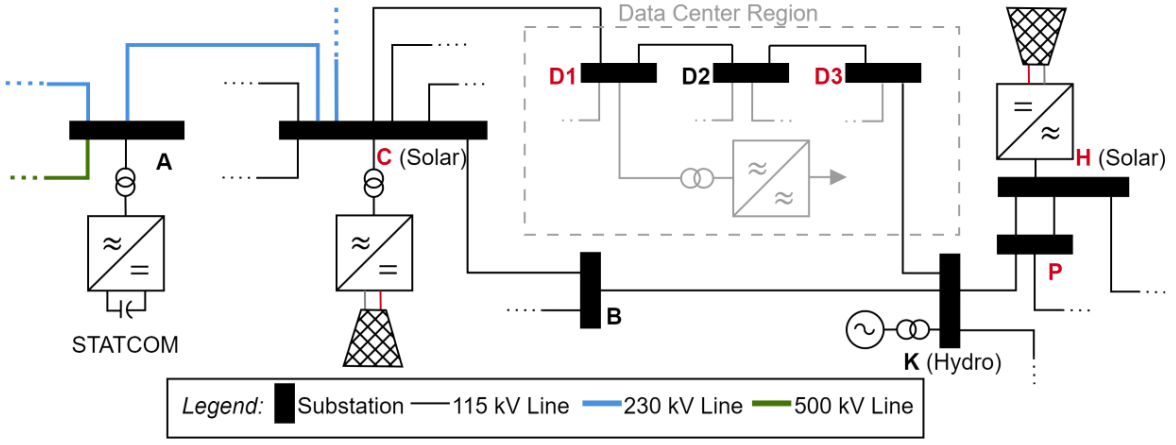

Figure 1. Region of Dominion Energy's Power Network.

3.2. Measurement Data

The synchrophasor data used in this study is reported at 30 Hz by Digital Fault Recorders (DFRs). The DFRs perform 4800 Hz sampling of analog signals, which are then filtered and down sampled to 960 Hz from where synchrophasors at 60 Hz are estimated. The estimated phasors are down sampled to 30 Hz before reporting them to Dominion Energy's Phasor Data Concentrator (PDC). Measurements are available at all the substations in Fig. 1 that are labelled with red letters, i.e., C, D1, D3, H and P. Note that there are no measurements at D2 and other unlabelled substations.

3.3. Event Description

Personnel at substations B, C, and D2 noticed flickering during a summer day and reported it to operations, this was the first instance when the behaviour shown in Figure 2 became of concern. Figure 2 shows that the largest magnitude oscillation was observed in the data center region. The oscillations were preceded by four hydro units at K ramping off and being taken offline at 2:05 AM, labelled A in Figure 2. At around 3:58 AM, labelled B in Figure 2, the STATCOM setpoint was changed from 0 MVAR to -71 MVAR, that is, switching it to reactor operation, resulting in the oscillations vanishing shortly after.

To adequately determine the inception and emergence of the identified phenomena, a rigorous characterization of the oscillation was needed. This case also required to understand the role of all other dynamic devices and processes, to determine if they were involved and to analyse historical data that would help better explained the observed phenomena. In the next section, we attempt to address all these issues through a data-driven analysis, i.e., using measurements only.

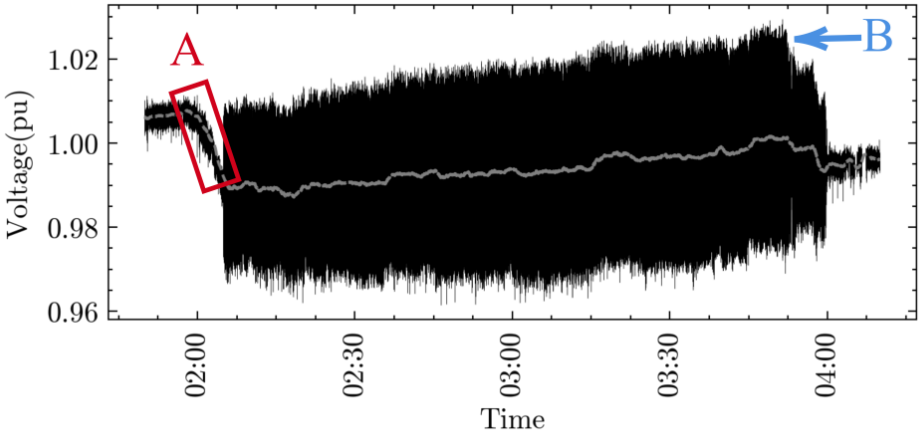

Figure 2 Voltage Magnitude at Substation D3. Label A indicates when four hydro units where ramped down preceding the oscillation. Label B indicates when the STACOM setpoint was changed from 0 to -71 MVAR.

4. Analysis

In the sequel, we present an analysis of the oscillatory phenomena presented above and verify it through historical data analysis of other similar instances. The analysis is divided in three parts, as shown in Fig. 3. First, the oscillation in Fig. 3 is characterized in the frequency domain to determine its frequency range. Once this has been determined, we analyse the mode shape of reactive and active power flows within the identified frequency range to make a more granular localization and to determine its governing control loop, i.e., whether it's related to a controller that regulates active or reactive power. Second, we expand our analysis to understand the underlying mechanism that gives rise to its inception of the oscillation. This is achieved by exploiting the presence of periodic voltage sags in the signal, which are exploited as a form of exogenous excitation to the system serving to aid in the characterization of the inception mechanism observed. Lastly, the last part performs historical data analysis serving to better explain both the inception process and the emergence of the oscillation for multiple other instances of the phenomena analysed for a single

instance. This last section serves both to verify and generalize our analysis of the original phenomena by considering multiple instances of it.

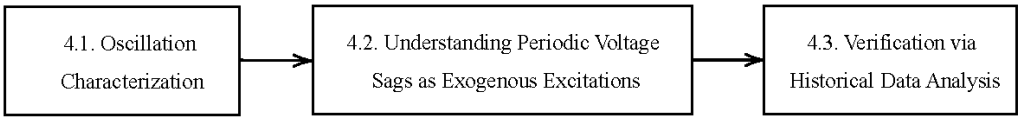

Figure 3. Analysis Process Diagram

4.1. Oscillation Characterization

4.1.1. Spectral Analysis

The first step to identifying the potential source of an oscillation is to characterize the oscillation itself. To do so, we start by contrasting the measurement data pre-and-during the event by computing the voltage magnitude spectrum at substation D3, as shown in Fig. 4. Each graph was obtained by computing spectral estimates of 10 min windows, each using Welch's method [26] with a 2 min FFT window. The dark black line is the average across individual estimates.

Observing both Fig. 4 a and b, it can be noted that there is a significant variability below 2 Hz range in both pre-and-during event time periods (labeled A), which indicates nonstationary behavior. Meanwhile, the behavior is largely stationary in roughly the 2-12 Hz range (labeled B). Furthermore, there is a family of spectral spikes at 1 Hz and their harmonics (indicated by a dashed square labeled C) in both pre-and-during event periods.

Next, in Fig. 4a's it can be observed that the PSD in the 10–12 Hz range most (labeled D) has a relatively higher value compared to that during the event. Meanwhile, observe in Fig. 4b, that there is an increase in the PSD in the 12-14 Hz frequency (labeled E), which significantly varies during the oscillation as compared to Fig. 4a.

Finally, observing Fig. 4 b, the PSD during the event is characterized by a prominent, narrow band spectral peak(s) in the 14-15 Hz range (labeled F), which is identified to be the sustained oscillation of interest. In addition, that for the same frequency range in Fig. 4 a, that is prior to the event, there is a substantial reduction in the PSD.

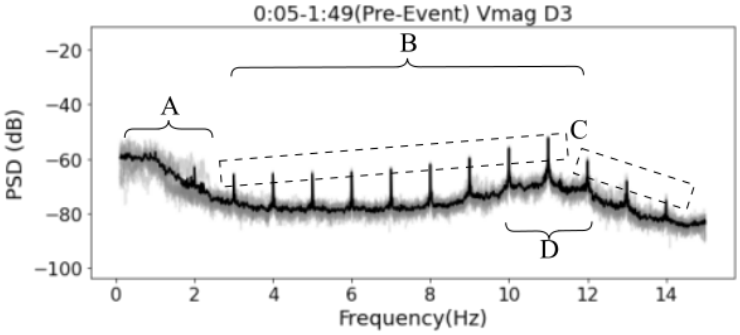

a. Pre-Event Spectrum

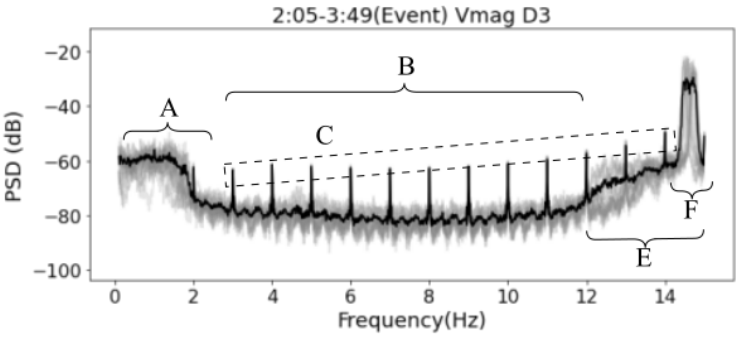

b. During-the-Event Spectrum

Figure 4. Sub D3 Voltage Magnitude Spectrum Pre-and-During the Event

Note that the spectrum peak corresponding to the oscillation (Fig. 4 b, labeled F) does not appear to be sharp, e.g., like an undamped sinusoid, but rather appears to have a flat 'cupola' with two distinct peaks. This could be either the result of two closely spaced oscillatory components or a single time-varying/nonstationary oscillation. To determine what kind of oscillation it is, the time-frequency plot in Figure 5 is obtained using Welch's method with a window length of 5 min. The oscillation of interest appears as a bright horizontal streak with its fundamental frequency gradually declining from 14.8 Hz to 14.6 Hz, which explains the flatness of average spectral estimate in Fig. 4 b. This nonstationary behavior is correlated with the gradually recovering operating voltage as shown by the gray trace in Fig. 2. In Fig. 5, one can also see a peculiar spread of energy from the oscillation frequency down to 13 Hz, which corresponds to the region labeled E in Figure 4 b. While it is difficult to explain such response without a model, it is not uncommon in practice in nonlinear systems [27] to see such cross-frequency couplings.

Next, we further study this frequency range by analyzing the mode shape of reactive and active power flows within to make a more granular localization and to determine its governing control loop, i.e., whether it's related to a controller that regulates active or reactive power.

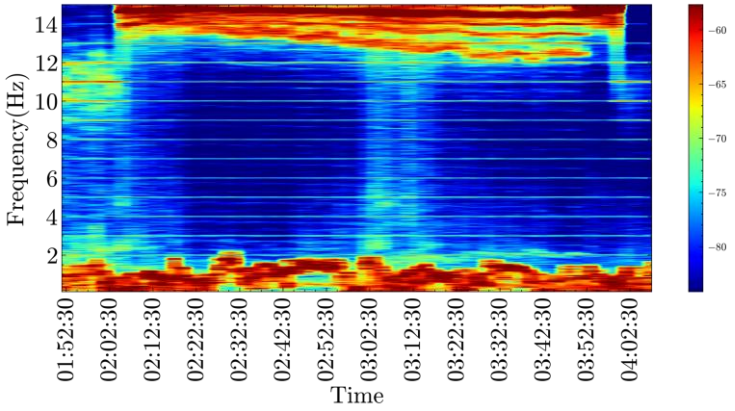

Figure 5. Spectrogram Substation D3 Voltage During Event.

4.1.2. Validation Against Waveform (Point-on-Wave) Measurement Data

Because the PMU data used for the analysis above was down sampled to 30 Hz (from original 60 Hz) by the storage software, there is a risk of aliasing [2], which makes the higher frequencies (15-30 Hz in this case) indistinguishable from the base band (0-15 Hz). In our case, this would result in the 14.7 Hz oscillation being at 15.3 Hz. To determine if this is the case, it is necessary to analyze waveform data (point-on-wave data) from the oscillation period, which is typically not stored. Luckily in this case, the DFR at D3 triggered and stored a 6 min window of data around the oscillation inception, which is plotted in Figure 6. Here, it is possible to observe that the amplitude of the waveform increases slightly in the last 1 minute of the window (approximately), which coincides with the beginning of

oscillation, which we will analyze next. To compare this data with our previous observations in PMU data, we need to transform it. Note that the 60 Hz component in waveform data acts as a carrier signal with the oscillation modulated into its amplitude and/or phase, which need to be extracted for analysis. For this, we use Hilbert transform [28] of the signal from the last 1 min window containing the oscillation and plot its spectral estimate in Figure 7. Here we can see that the oscillation component is indeed at 14.7 Hz and not displaced due to aliasing.

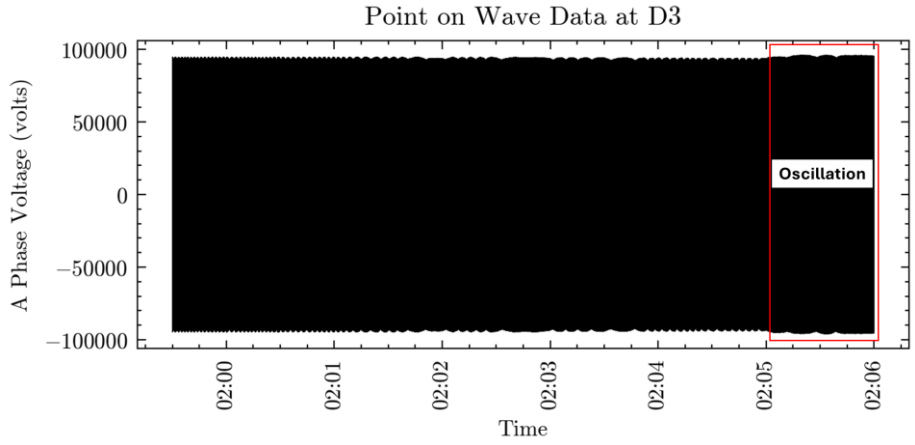

Figure 6 Waveform Data Phase A Voltage at D3

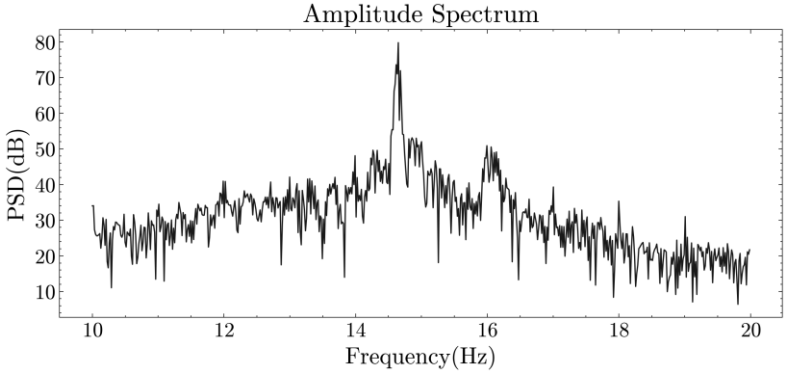

Figure 7 Amplitude Spectrum of Phase A Voltage at D3

4.1.3. Mode Shape Analysis for the 14.7-14.8 Hz Frequency

By analyzing the relative observability of the 14.7-14.8 Hz oscillation across the neighboring substations shown in Figure 1, we aim to localize the source more granularly by applying the approach in [29]. This will be carried out by concentrating on a 30-min. window between 2:10 and 2:40 AM, where the oscillation is in the range of 14.7 to 14.8 Hz. The mode shape was estimated at 14.72 Hz* using Frequency Domain Decomposition (FDD)[20] for which results are shown in Figure 8 and Figure 9.

Figure 8 presents the reactive power (Q) mode shape, where the biggest component of Q oscillation flows outwards from the data center region, i.e., along D1 to C and D3 to K. Based on Q mode shape magnitudes, there is very little flow from D3 to D1 through D2, narrowing the localization only to two possible ones. Although measurements are

* It should be noted that the mode shapes were determined along the 14.6–14.9 Hz range using a modal assurance criterion of 0.95; therefore, the oscillation's modest nonstationary nature does not result in issues in mode shape estimation.

unavailable at D2, it is possible to observe that there is only a small difference between the Q oscillation entering and leaving D2; consequently, for all intents and purposes, it plays little to no role in this event. Therefore, these results point towards both D1 and D3 being the substations where the oscillation originate from, that is, the data center region. Finally, note that the two solar PV plants (at C and H) do not contribute to the oscillation, which is expected given that this oscillation occurred at night (12 AM local time) and these plants do not have night VAR capability.

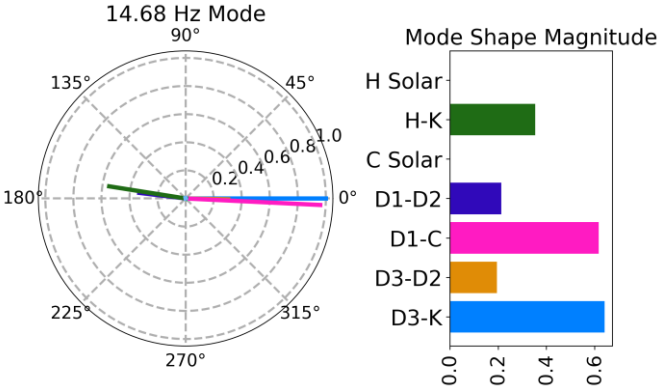

Figure 8. Q Flow Relative Mode Shape (left) and Mode Shape Magnitude (right).

As the above analysis carried out using Q flow mode shapes, it is necessary to verify if indeed the reactive power dominates over other potential variables being affected by the data center region. To make this verification, we obtain the mode shapes for the data center region as shown in Figure 9. The mode shapes shows that the oscillation has a relatively larger participation in Q as compared to P and voltage magnitude, which confirms the analysis above.

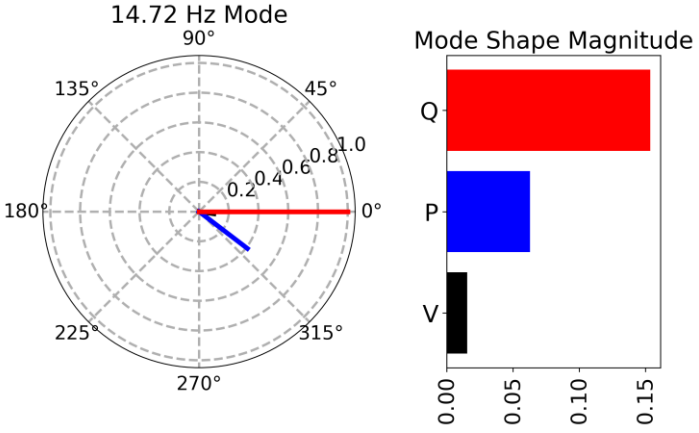

Figure 9. Data center mode shapes of Q and P flow and voltage magnitude

4.2. Understanding Periodic Voltage Sags as Exogenous Excitations

Next, we study the dynamics appearing as spectral peaks at 1 Hz and its multiples (labelled C in Figure 4). The closest known signal type that could result in such a spectrum is a pulse train. In our previous work [30], we observed a similar phenomenon at 0.016 Hz and its multiples resulting from a voltage sag that occurs once a minute, and can be observed throughout the power network. However, in the present case, the fundamental frequency is 1 Hz and, therefore, it should be happening once a second. To confirm this, we plot the voltage magnitude at D1 in Figure 10. One can see a sharp drop in voltage every second (indicated within the dashed box labelled A), which represents periodic voltage sags. Such sags are visible for more than a minute in Figure 10, as indicated by the scale labelled B.

Furthermore, the size of the drop increases significantly for roughly 10 sec. every minute, labelled C in Figure 10. This could be due to the duty cycle of an data center. Although such exogenous excitation is undesirable, it can help analyse the response of the system, which is otherwise difficult to solely estimate from ambient conditions.

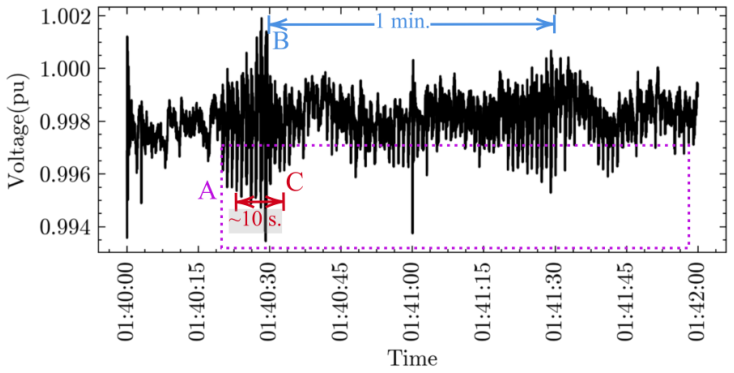

Figure 10. Voltage Magnitude at D1 Showing a Periodic Voltage Sag Every Second

We take a closer look at the dynamics in the 10-12 Hz range (labelled C in Fig. 4), using the time-frequency methods summarized in Section 3. Recall that these dynamics vanish during the oscillation event and, therefore, we need to study the periods where the oscillation is not present. Here is where the exogenous excitation from the periodic voltage sags can be exploited. Figure 11 shows a time-frequency plot of voltage magnitude at sub D1 in the 5-15 Hz range. One can see that the relevant dynamics are concentrated in 10-11 Hz as confirmed by the spectrum plot in Fig. 4. However, what cannot be observed in Figure 4 is the temporal information that Figure 11 provides, which shows that the mode is only observable for a short duration of ~ 10 s, separated by precisely a minute. These periods interestingly align those when voltage sags become large. Zooming into one of these ~10 s periods in Figure 12 confirms this, where it is seen that the voltage sag on top of each second triggers a ringdown response. This is what we were alluding to earlier where exogenous excitations expose otherwise difficult to observe underlying system behaviour. Now, the question arises as to why this mode vanishes when the forced oscillation is happening and whether it has a relationship to it, which is addressed next through historical data analysis.

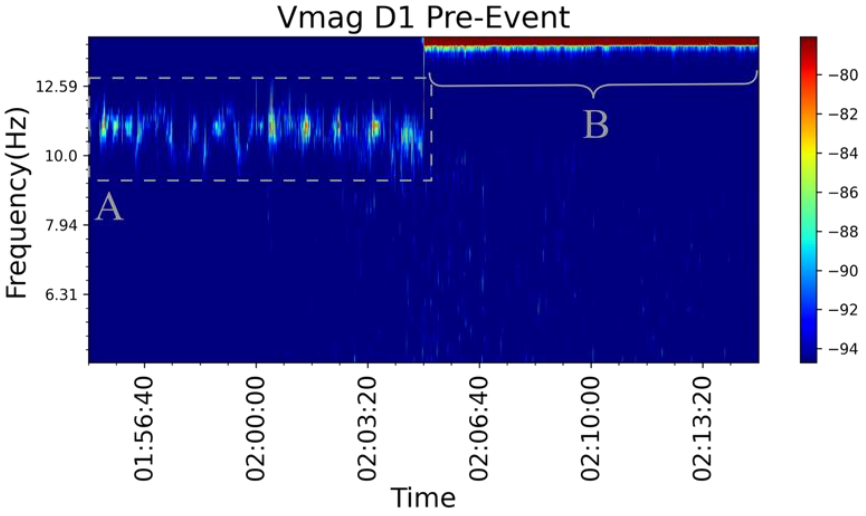

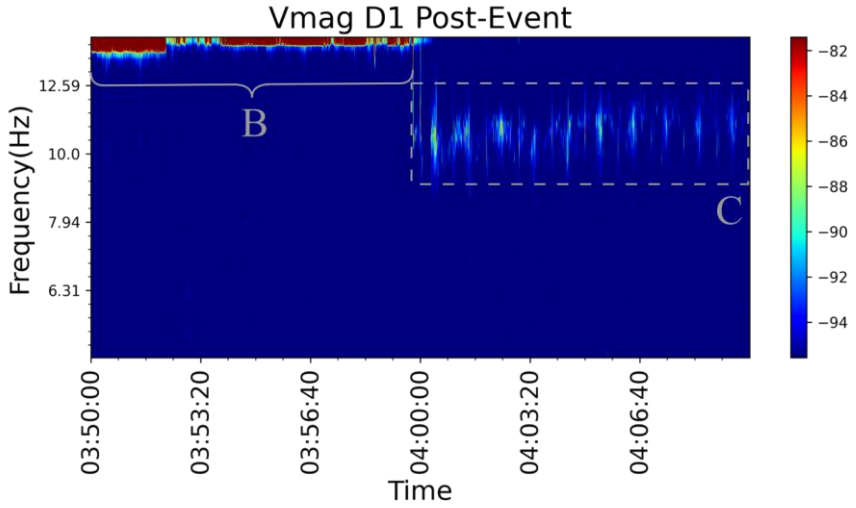

Figure 11. 10-11 Hz Time-Frequency Representation. A: indicates energy content at ~10.5 Hz before the event; B: indicates the oscillation event period; C: indicates energy content at ~10.5 Hz after the event.

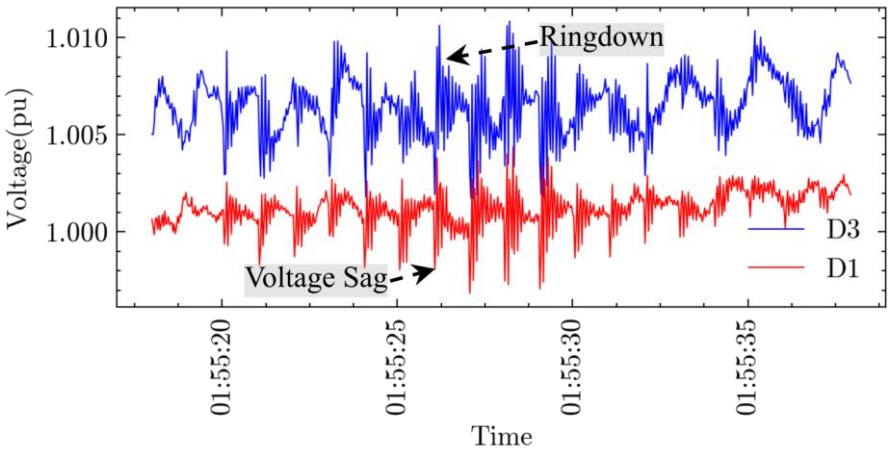

Figure 12. Voltage Sag Exciting 10-11 Hz Dynamics.

4.3. Verification via Historical Data Analysis

Given the fact that data from similar events was found, this provided a unique opportunity to verify the conclusions made regarding the oscillation's inception and emergence. Next, we explain how such data was found, how historical data analysis was used to provided further insight into the incipient behaviour prior to the emergence of the oscillation and to verify the previously made observations.

4.3.1. Detecting Similar Event Instances

According to the analysis so far, the oscillation is concentrated between 13 and 15 Hz. Therefore, high PSD values in this frequency band characterize key periods in the historical data, which can be utilized as a metric to detect them, similar to what is done in RMS energy detectors [31]. Let us denote this metric by

$$S_{event}(t) = \max(\{S(t,f) \forall f \in [13,15]\}), \quad (5)$$

where $S(t,f)$ is the PSD at time t as defined before. During system transients that are a result of faults, large load changes, etc., the PSD will have a broadband frequency domain characteristic [32]. This means that such transients also cause a significant increase in the PSD value in this range of interest (along with all other frequencies), which could lead to false positives when attempting to detect similar events. To address this issue, it is necessary to single out and filter the periods when this occurs, for which we use

$$S_{transient}(t) = \text{median}(\{S(t,f) \forall f \in [5,15]\}). \tag{6}$$

Meanwhile, periods when the oscillation appears are characterized by

$$S_{event}(t) \gg S_{transient}(t). \tag{7}$$

Next, considering ambient conditions and in the absence of any other dynamic behavior in 13-15 Hz range, we have that

$$S_{event}(t) \approx S_{transient}(t). \tag{8}$$

as can be seen in the in Figure 4a. Finally, since we are working with an estimate of $S(t,f)$ denoted by $\hat{S}(t,f)$, taking a log transform is advantageous in that it makes the variance of $\hat{S}(t,f)$ approximately stationary [19], thereby allowing to work across t and f,leading to the following candidate for the detector,

$$d(t) = log(\hat{S}_{event}(t)) - log\left(\hat{S}_{transient}(t)\right) > \text{threshold}. \tag{9}$$

Using this detection on the available historical data, the plot in Figure 13 was obtained. In the figure, there are two major groups of operating conditions labelled A and B, respectively. Here, the distance from the black line ($y = x$) represents the metric d. It was discovered that Cluster A belongs to windows with a considerable amount of missing data and therefore a low value of signal power. Cluster B can be further divided into two clusters, depending on the distance from the black line. Here, Grubbs Test [28] can be applied on $d(t)$ in Cluster B with α=0.05 to detect oscillation events as a form of outliers (red dots labelled C). Note that the Grubbs test stipulates that the points must be i.i.d. normally distributed, however, this is not always the case because the grid's operational conditions affect $S(t,f)$ and vary continuously in time t, therefore having a dependence. Regardless, gross outliers are still identifiable it can clearly by observed in Cluster C shown in Figure 13.

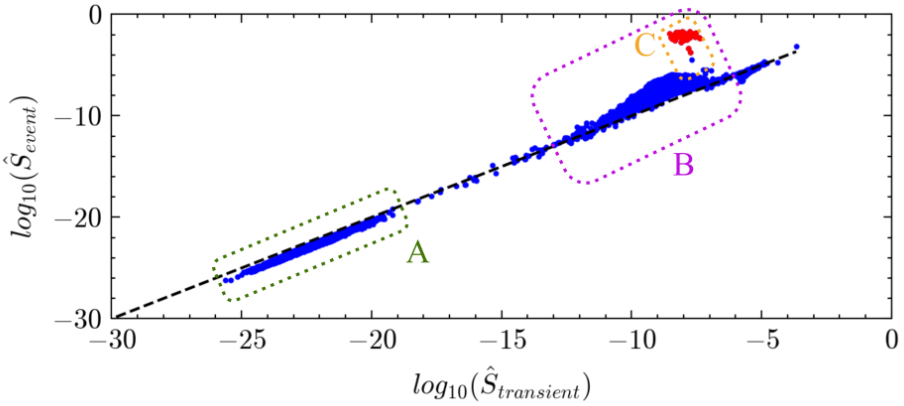

Figure 13. Historical Spectral Estimates. A: cluster with significant missing data; B: cluster with potential events; C: events identified as a form of outlier using Grubbs Test.

4.3.2. Analysis of Multiple Oscillation Events

Next, we take a closer look at a few of the events identified (shown as red dots in Figure 13) using the methods described in Section 3. Figure 14 shows a 2 min window of 7 of the identified events in the time domain (left), as well

as the time-frequency plane (right) obtained using the synchrosqueezed CWT. Here, the oscillation event under study is at the top, which is what we have been concentrating on so far. The previously examined 10–11 Hz dynamics which denote the inception of the ~ 14 Hz oscillation over different time periods characterizes all of these. In addition, oscillations frequently emerge and disappear without the system undergoing any noticeable changes, such as line outages, faults, etc. Note that the short-lived oscillations, which contain vital information on the underlying inception mechanism behind the oscillation, are relatively difficult to capture due to the time-frequency resolution trade-off.

The dynamics between 10 and 11 Hz, which frequently lead to the ~14.7 Hz oscillation, are carefully examined next. As was previously observed, voltage sags occur precisely once every second for about 10 s every minute allow us to examine the system response. For this analysis, we select the window in each of the previously discussed cases (see Figure 14) before the oscillation event. The analysis is carried out as follows:

- First, we divide the window into 10 segments highlighted grey in Figure 15 (left column plots).
- Next, for each segment, we use voltage magnitudes at substations D1 and D3, detrend, and band pass filter them to 9-15 Hz and fit a second-order model to estimate the dominant mode using Dynamic Mode Decomposition [33].
- Finally, the trajectory of the estimated mode is plotted in the damping-frequency plane in Figure 15. As can be observed, a critical mode in the 10- to 11-Hz range approaches dangerously low damping in all cases, starting as a small signal stability problem followed by a sudden increase in frequency when settling on a limit cycle (which are the oscillation events). It is worth noting that the nonlinear system response makes the linear and low-order model employed to estimate the damping of this mode unsuitable for analysis at unstable periods (when the damping < 0), and as a result, the mode estimates obtained under such conditions are unreliable.

Thus, we can see that a 10–11 Hz mode local to data center becomes unstable with changing operating conditions, which is the inception of the oscillation at 14.7 Hz. Observe that this oscillation is different from forced oscillations [34], which are characterized by an external oscillatory input usually due to an equipment malfunction. Since state variables are not measured, explaining why an instability resulted in an oscillation will require deductive reasoning using observations in output measurements. Following an equilibrium becoming unstable, the system trajectory can do one of two things:

- **Go unbounded** – In all of the oscillation events in Figure 14, we can see that once the original equilibrium re-stabilizes due to organically changing operating conditions, which took anywhere from a few seconds to hours, the trajectory immediately returns to an equilibrium marked by the oscillation vanishing abruptly. For this to happen, the trajectory must be bounded within the stability region [35] of the re-stabilized equilibrium, in some cases for as long as 2 Hrs, which makes the unbounded case highly unlikely.
- **Enter some invariant set** other than the original equilibrium (limit cycle, equilibrium, etc). In this case, the trajectory remains bounded, which makes it a more likely scenario in our case. Observe that the trajectory does not settle to a new equilibrium point, as the output measurement changing continuously (or rather, it oscillates). Given that the response behaves as a periodic motion, a limit cycle [36] (which is an isolated periodic orbit), the most likely scenario is that what we observe is in fact a limit cycle.

Proceeding with the limit cycle hypothesis, we would like to understand the mechanism that led to its creation. Noting that the creation of limit cycle was preceded by a pair of complex eigen values (oscillatory mode) crossing the imaginary axis, the initial assumption is that we are observing a supercritical Hopf bifurcation. Such kind of bifurcation involves a creation of stable limit cycle at the instance of a stable equilibrium destabilizing, thereby catching the system trajectory. Note that a distinguishing characteristic of this mechanism is that the limit cycle created starts off with essentially a zero radius and grows as the system conditions change (for the worse). However, in all the events analyzed, as soon as the system becomes unstable, the oscillation starts off at a large radius even though the operating point is changing gradually (particularly in the June 6^{th} events), making the validity of this hypothesis questionable. A more common mechanism behind creation of limit cycles from controller instability, stems from the fact that most real world controllers (particularly those in power electronic-based devices) are equipped with hard limiters (such as under/over excitation limits in voltage controllers) to protect the equipment. [37] and [38] presented a mechanism where the hard limits chop the state space in a way that creates a cyclic motion (i.e., a limit cycle), once

the trajectory goes unstable and hits those limits. These limits are usually away from the original normal operating point (equilibrium), thereby resulting in a relatively large amplitude oscillation, governed by the limits. In other words, the trajectory grows and enters a cyclic motion, which seems to better explain our oscillations stemming from an unstable UPS in the data center.

Finally, the long-term behaviour of the critical mode of 10-11 Hz critical mode, captured by

$$S_{mode}(t) = \max(\{S(t,f) \forall f \in [10,11]\}) \quad (10)$$

and $S_{transient}(t)$ (see Equation (11)), is shown in Figure 16. Periods with a disproportionately high value are labelled A. These are consistent with a local transmission line outage in the vicinity of the data center, which reduces the damping of the critical mode and raises the PSD value (and hence, $S_{mode}(t)$), as explained above. The significant increase that begins in May 2022 (labelled B) and continues until the day before the oscillation events in June is important to highlight. This suggests that, in terms of data center performance, the grid may be operated regularly under more stressed conditions, that is, with a decrease in stability margin. This is in line with our analysis in Section 4.1, and further points to the data center region as the culprit.

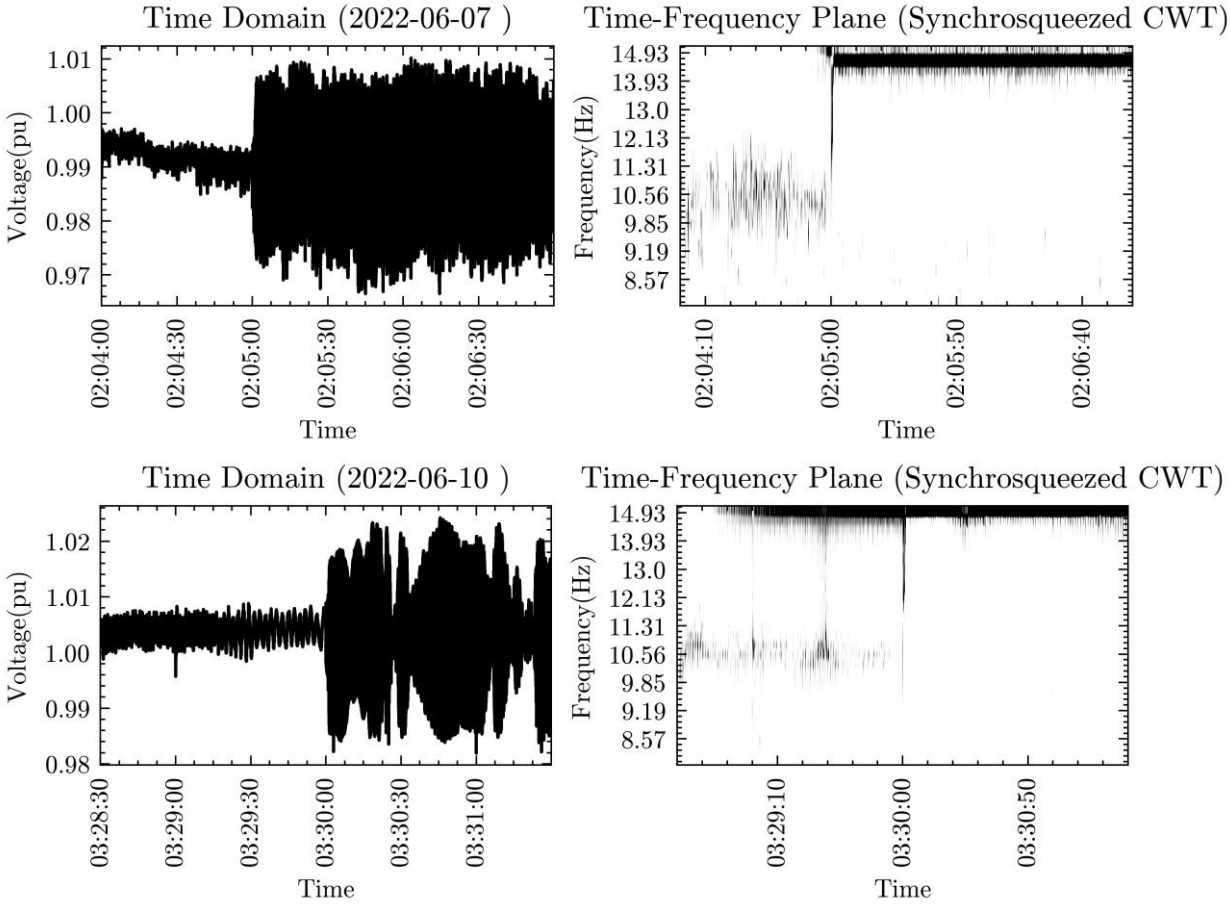

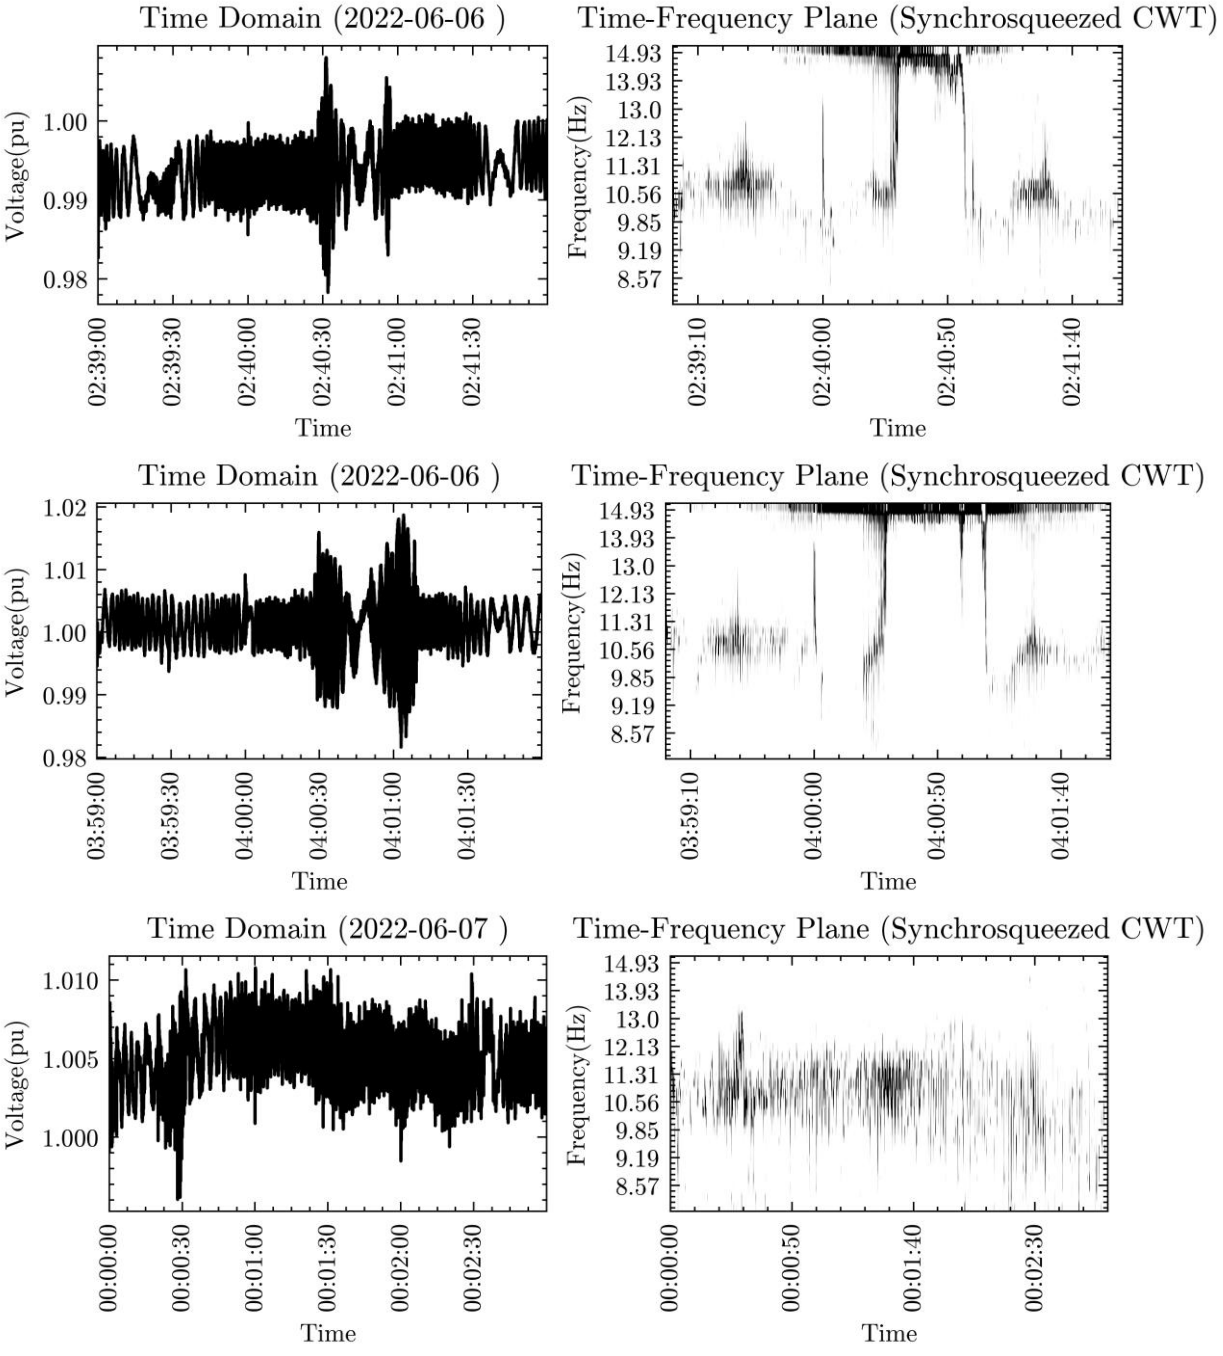
Time Domain (2022-06-06)
Time-Frequency Plane (Synchrosqueezed CWT)
Voltage(pu)
1.00
0.99
0.98
02:39:00
02:39:30
02:40:00
02:40:30
02:41:00
02:41:30
Time
Frequency(Hz)
14.93
13.93
13.0
12.13
11.31
10.56
9.85
9.19
8.57
02:39:10
02:40:00
02:40:50
02:41:40
Time
Time Domain (2022-06-06)
Time-Frequency Plane (Synchrosqueezed CWT)
Voltage(pu)
1.02
1.01
1.00
0.99
0.98
03:59:00
03:59:30
04:00:00
04:00:30
04:01:00
04:01:30
Time
Frequency(Hz)
14.93
13.93
13.0
12.13
11.31
10.56
9.85
9.19
8.57
03:59:10
04:00:00
04:00:50
04:01:40
Time
Time Domain (2022-06-07)
Time-Frequency Plane (Synchrosqueezed CWT)
Voltage(pu)
1.010
1.005
1.000
00:00:00
00:00:30
00:01:00
00:01:30
00:02:00
00:02:30
Time
Frequency(Hz)
14.93
13.93
13.0
12.13
11.31
10.56
9.85
9.19
8.57
00:00:00
00:00:50
00:01:40
00:02:30
Time

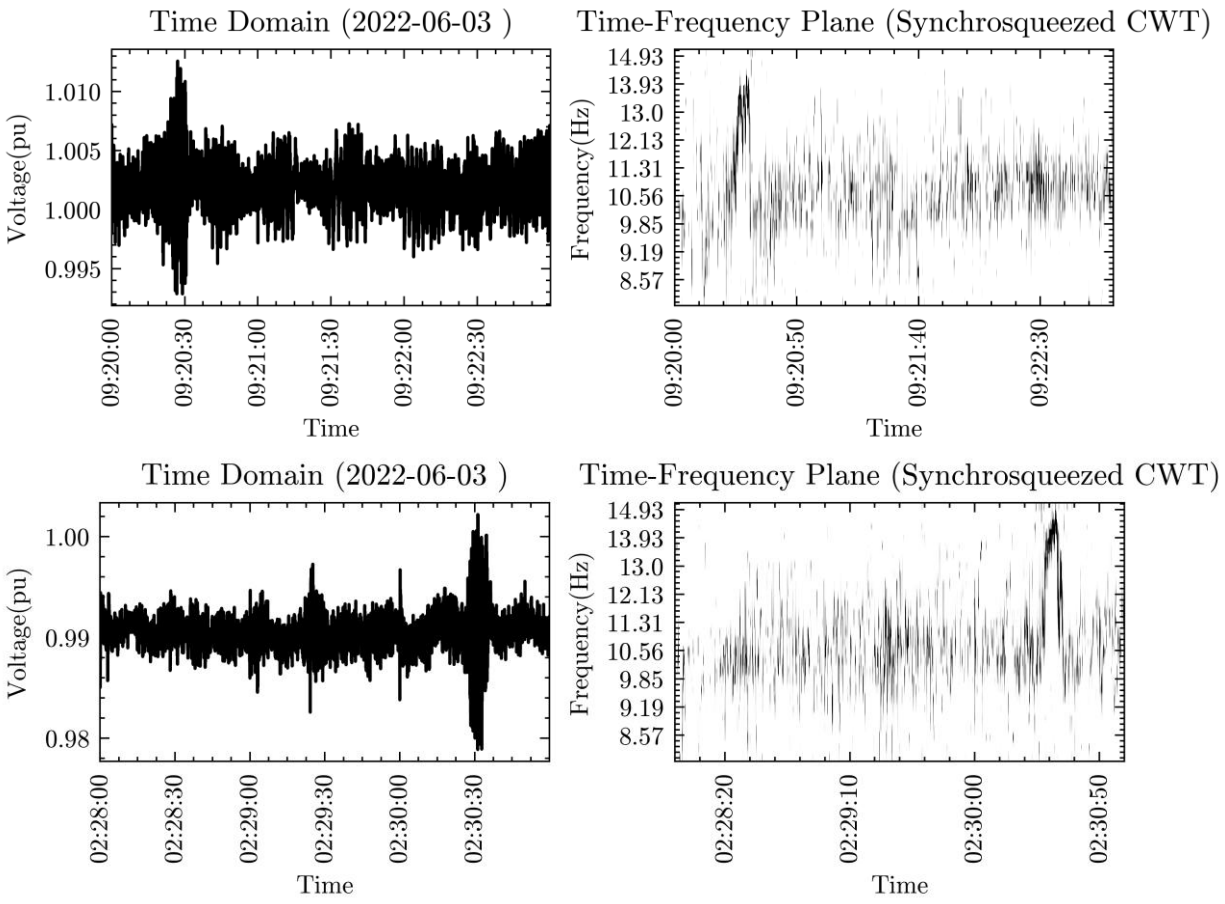

Figure 14. Identified Oscillations, Time Domain (Left), Time-Frequency Plane (Right)

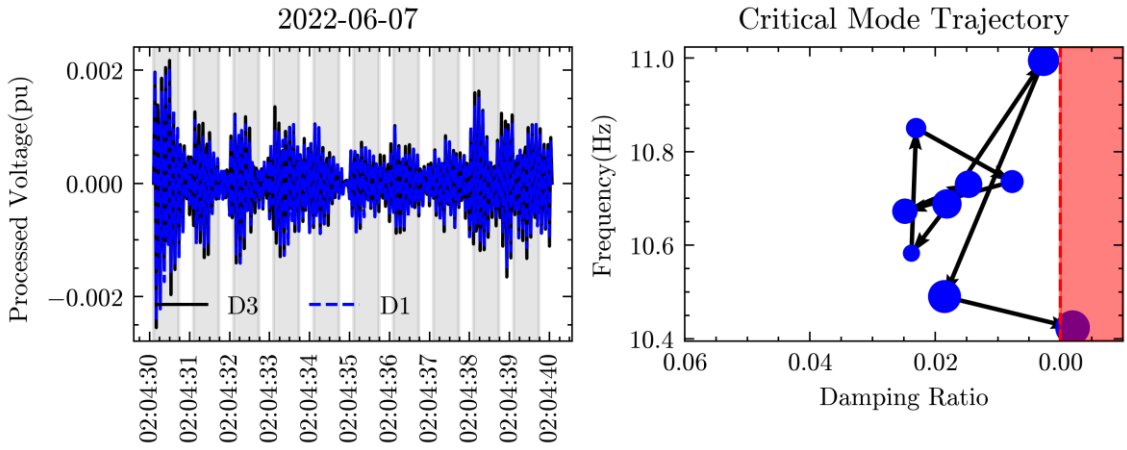

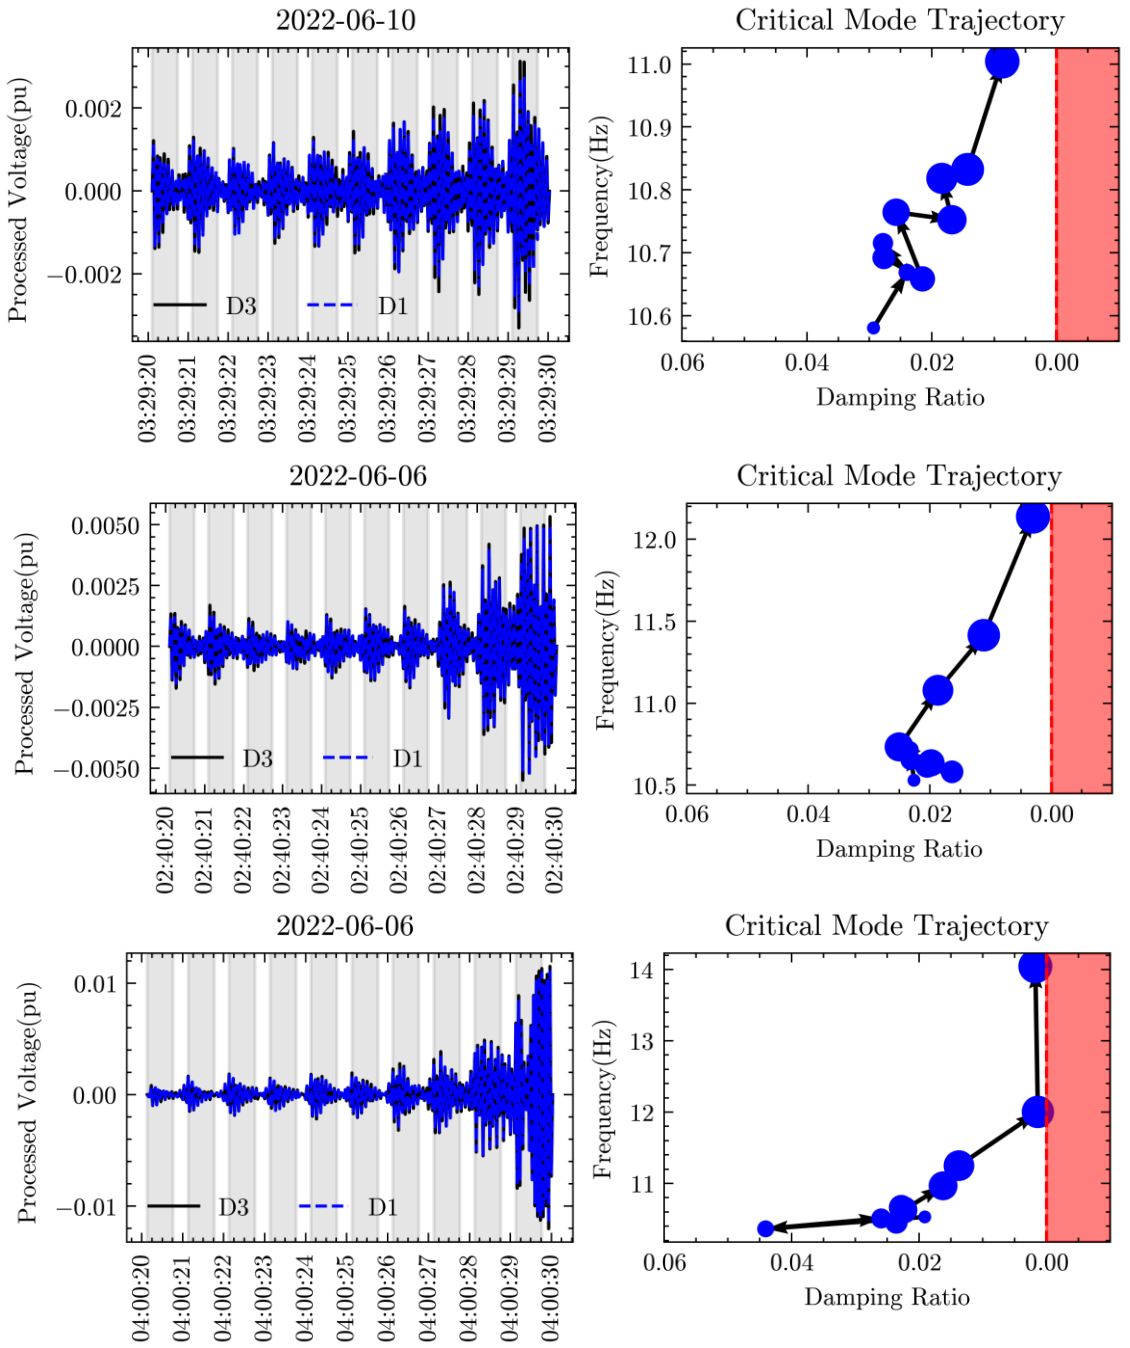
2022-06-10
Processed Voltage(pu)
0.002
0.000
−0.002
D3
D1
03:29:20
03:29:21
03:29:22
03:29:23
03:29:24
03:29:25
03:29:26
03:29:27
03:29:28
03:29:29
03:29:30
Critical Mode Trajectory
Frequency(Hz)
11.0
10.9
10.8
10.7
10.6
0.06
0.04
0.02
0.00
Damping Ratio
2022-06-06
Processed Voltage(pu)
0.0050
0.0025
0.0000
−0.0025
−0.0050
D3
D1
02:40:20
02:40:21
02:40:22
02:40:23
02:40:24
02:40:25
02:40:26
02:40:27
02:40:28
02:40:29
02:40:30
Critical Mode Trajectory
Frequency(Hz)
12.0
11.5
11.0
10.5
0.06
0.04
0.02
0.00
Damping Ratio
2022-06-06
Processed Voltage(pu)
0.01
0.00
−0.01
D3
D1
04:00:20
04:00:21
04:00:22
04:00:23
04:00:24
04:00:25
04:00:26
04:00:27
04:00:28
04:00:29
04:00:30
Critical Mode Trajectory
Frequency(Hz)
14
13
12
11
0.06
0.04
0.02
0.00
Damping Ratio

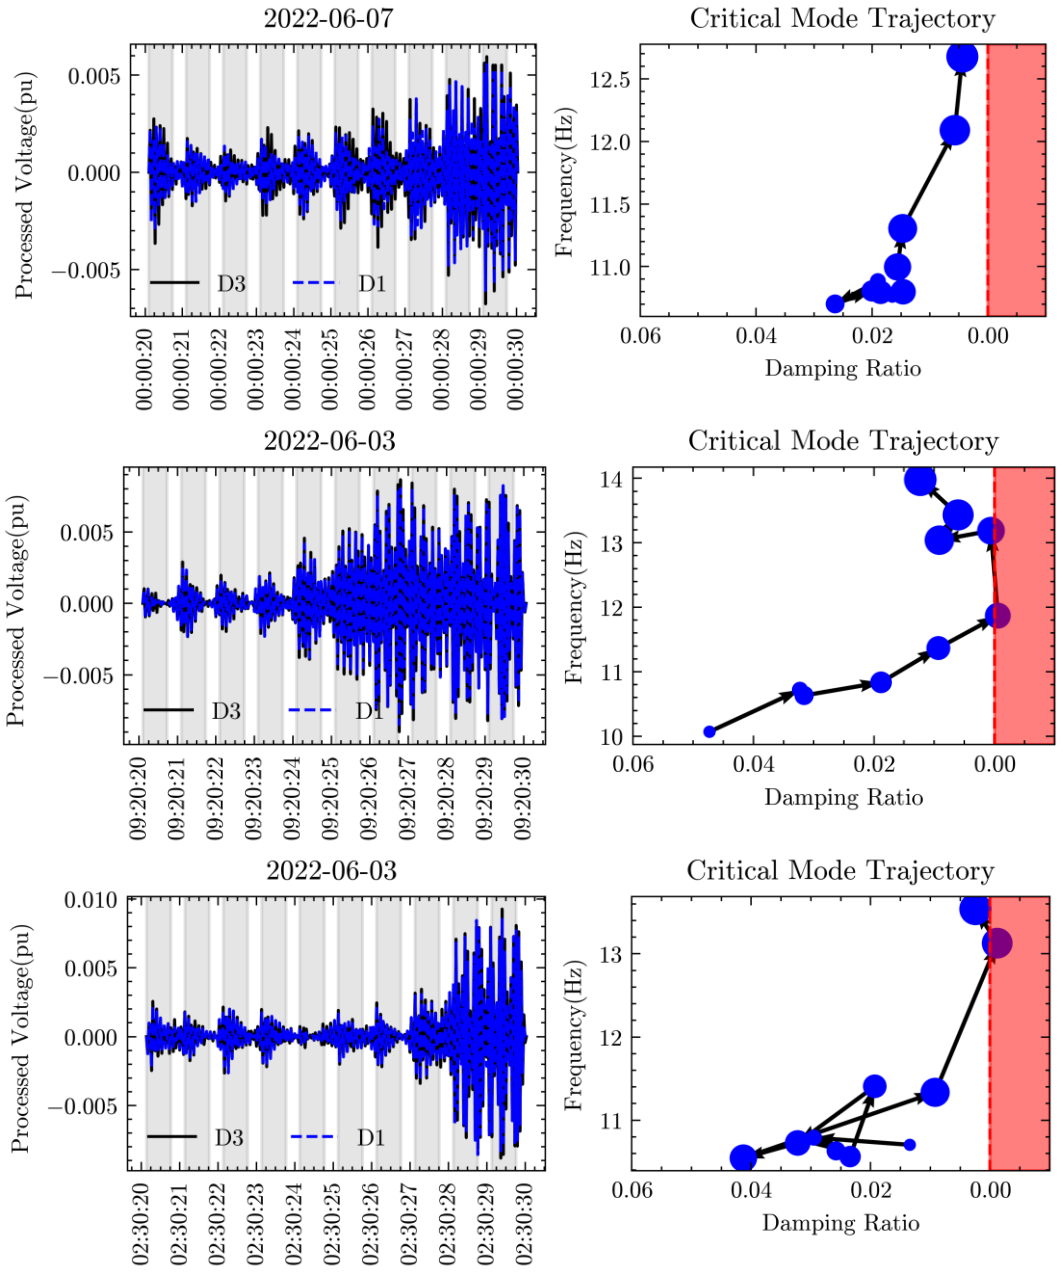

Figure 15. Evolution of the Damping of the 10-11 Hz Dynamics

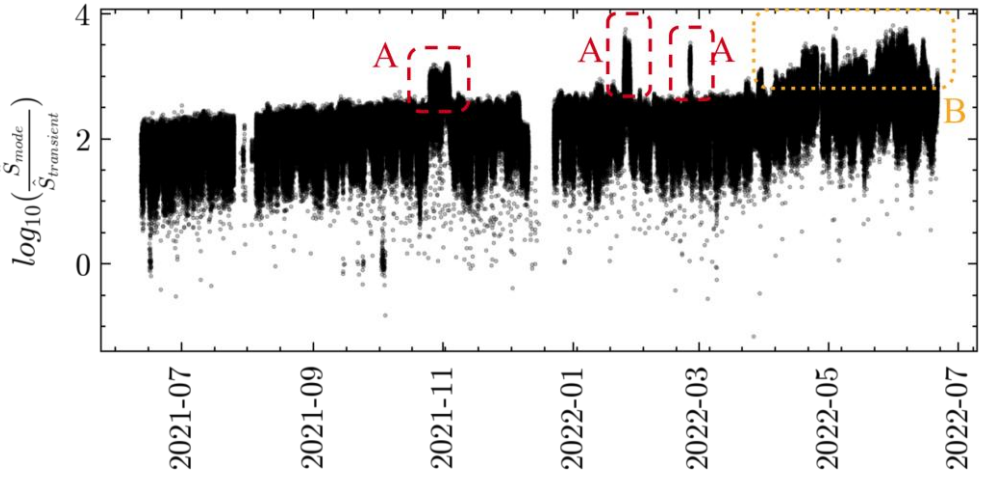

Figure 16. 10-11 Hz Signal Power of the Voltage Magnitude at D3.
A: Periods with disproportionally high value in $S_{mode}(t)$. B: Significant increase in $S_{mode}(t)$.

5. Conclusion

This paper investigated a real-world 14.7 Hz forced oscillation in the Dominion Energy service area that originates from a data center-rich region using measurement-based analysis. The distinctive feature of this phenomena rests in the fact that the oscillation frequency differs significantly from the underlying mode, where its association is inferred via time-frequency analysis. The inception of this unforeseen dynamic behaviour involved a once-a-second voltage sag, triggering a response of a mode in the 10-11 Hz range. Taking advantage of this voltage sag as a form of periodic exogenous excitation, it was possible to track the loss of damping of the 10-11 Hz mode and its transition to a ~14.7 Hz limit cycle, which would otherwise be challenging to estimate from ambient data.

Finally, in this work, it was crucial for the data center owner to verify the analysis results, as it is not owned by the utility. The owner of the asset confirmed that the oscillations were caused by the input stage of the UPS units that are deploy at the data center, which have different ancillary service capabilities [4] [5]. The correct characterization of this complex dynamic process resulted in minimizing the liability on its emergence for the utility, a benefit that will have lasting effects in addressing growing stability issues at the transmission level due to the increase of converter-interfaced loads and energy sources.

References

[1] Y. Cheng *et al.*, "Real-World Subsynchronous Oscillation Events in Power Grids With High Penetrations of Inverter-Based Resources," *IEEE Trans. Power Syst.*, vol. 38, no. 1, pp. 316–330, Jan. 2023, doi: 10.1109/TPWRS.2022.3161418.

[2] C. Wang, C. Mishra, K. D. Jones, R. M. Gardner, and L. Vanfretti, "Identifying Oscillations Injected by Inverter-Based Solar Energy Sources," in *2022 IEEE Power & Energy Society General Meeting (PESGM)*, Jul. 2022, pp. 1–5. doi: 10.1109/PESGM48719.2022.9916830.

[3] C. Mishra, L. Vanfretti, J. Delaree, and K. D. Jones, "Analyzing a Non-Sinusoidal Response from a Real-World Solar PV," *IEEE Trans. Power Syst.*, vol. 39, no. 2, pp. 4771–4774, Mar. 2024, doi: 10.1109/TPWRS.2024.3350377.

[4] I. Alaperä, S. Honkapuro, and J. Paananen, "Data centers as a source of dynamic flexibility in smart girds," *Appl. Energy*, vol. 229, pp. 69–79, Nov. 2018, doi: 10.1016/j.apenergy.2018.07.056.

[5] I. Alaperä, J. Paananen, K. Dalen, and S. Honkapuro, "Fast frequency response from a UPS system of a data center, background, and pilot results," in *2019 16th International Conference on the European Energy Market (EEM)*, Sep. 2019, pp. 1–5. doi: 10.1109/EEM.2019.8916344.

[6] S. Sheehan and A. Rakow, "Evolving a Data Center Into a Microgrid: Industry perspectives and lessons learned," *IEEE Electrification Mag.*, vol. 11, no. 3, pp. 16–25, Sep. 2023, doi: 10.1109/MELE.2023.3291193.

[7] E. Oró, V. Depoorter, A. Garcia, and J. Salom, "Energy efficiency and renewable energy integration in data centres. Strategies and modelling review," *Renew. Sustain. Energy Rev.*, vol. 42, pp. 429–445, Feb. 2015, doi: 10.1016/j.rser.2014.10.035.

[8] S. Naganandhini, C. Sundar, R. Sathya, and G. Venkatesan, "Towards Energy-Efficient Data Centres: A Comprehensive Analysis of Cooling Strategies for Maximizing Efficiency and Sustainability," in *2023 Intelligent Computing and Control for Engineering and Business Systems (ICCEBS)*, Dec. 2023, pp. 1–6. doi: 10.1109/ICCEBS58601.2023.10448782.

[9] C. Guo, F. Luo, Z. Cai, and Z. Y. Dong, "Integrated energy systems of data centers and smart grids: State-of-the-art and future opportunities," *Appl. Energy*, vol. 301, p. 117474, Nov. 2021, doi: 10.1016/j.apenergy.2021.117474.

[10] X. Chen *et al.*, "Energy-saving superconducting power delivery from renewable energy source to a 100-MW-class data center," *Appl. Energy*, vol. 310, p. 118602, Mar. 2022, doi: 10.1016/j.apenergy.2022.118602.

[11] Í. Goiri, W. Katsak, K. Le, T. D. Nguyen, and R. Bianchini, "Designing and Managing Data centers Powered by Renewable Energy," *IEEE Micro*, vol. 34, no. 3, pp. 8–16, May 2014, doi: 10.1109/MM.2014.6.

[12] B. Acun *et al.*, "Carbon Explorer: A Holistic Framework for Designing Carbon Aware Datacenters," in *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, in ASPLOS 2023. New York, NY, USA: Association for Computing Machinery, Jan. 2023, pp. 118–132. doi: 10.1145/3575693.3575754.

[13] D. Kosterev *et al.*, "Implementation and operating experience with oscillation detection application at bonneville power administration," in *Proceedings of CIGRE 2016 Grid of the Future Conference*, 2016, pp. 1–12.

[14] D. J. Sullivan, B. K. Buterbaugh, and R. L. Allison, "Installation and Commissioning of Mitsubishi Electric's MMC STATCOM (SVC Diamond™) at Dominion Energy's Colington Substation," in *2018 Grid of the Future Symposium*, 2018.

[15] C. Mishra *et al.*, "Analysis of STATCOM Oscillations using Ambient Synchrophasor Data in Dominion Energy," in *2022 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*, Apr. 2022, pp. 1–5. doi: 10.1109/ISGT50606.2022.9817489.

[16] C. Mishra, L. Vanfretti, J. Jr, and K. Jones, "Automatically Discerning Power System Dynamics in Synchrophasor Measurements Data Spectra," presented at the 2023 IEEE Power & Energy Society General Meeting, Jul. 2023.

[17] X. Xu *et al.*, "Fast Oscillation Detection and Labeling via Coarse Grained Time Series Data for ML Applications," in *2022 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*, Apr. 2022, pp. 1–5. doi: 10.1109/ISGT50606.2022.9882712.

[18] J. W. Pierre, D. J. Trudnowski, and M. K. Donnelly, "Initial results in electromechanical mode identification from ambient data," *IEEE Trans. Power Syst.*, vol. 12, no. 3, pp. 1245–1251, Aug. 1997, doi: 10.1109/59.630467.

[19] P. Stoica and R. Moses, *Introduction to Spectral Analysis*, 1st edition. Upper Saddle River, N.J: Prentice Hall, 1997.

[20] R. Brincker, L. Zhang, and P. Andersen, "Modal identification of output-only systems using frequency domain decomposition," *Smart Mater. Struct.*, vol. 10, no. 3, p. 441, Jun. 2001, doi: 10.1088/0964-1726/10/3/303.

[21] D. J. Trudnowski, "Estimating Electromechanical Mode Shape From Synchrophasor Measurements," *IEEE Trans. Power Syst.*, vol. 23, no. 3, pp. 1188–1195, Aug. 2008, doi: 10.1109/TPWRS.2008.922226.

[22] K. Gröchenig, *Foundations of Time-Frequency Analysis*. in Applied and Numerical Harmonic Analysis. Boston, MA: Birkhäuser, 2001. doi: 10.1007/978-1-4612-0003-1.

[23] M. B. Priestley, "Non-linear and non-stationary time series analysis," *Lond. Acad. Press*, 1988.

[24] I. Daubechies, "The wavelet transform, time-frequency localization and signal analysis," *IEEE Trans. Inf. Theory*, vol. 36, no. 5, pp. 961–1005, Sep. 1990, doi: 10.1109/18.57199.

[25] I. Daubechies, J. Lu, and H.-T. Wu, "Synchrosqueezed wavelet transforms: An empirical mode decomposition-like tool," *Appl. Comput. Harmon. Anal.*, vol. 30, no. 2, pp. 243–261, Mar. 2011, doi: 10.1016/j.acha.2010.08.002.

[26] 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 Electroacoustics*, vol. 15, no. 2, pp. 70–73, Jun. 1967, doi: 10.1109/TAU.1967.1161901.

[27] "Bicoherence analysis of nonstationary and nonlinear processes," EUROfusion Scientific Publications. Accessed: Sep. 06, 2022. [Online]. Available: https://scipub.euro-fusion.org/archives/eurofusion/bicoherence-analysis-of-nonstationary-and-nonlinear-processes

[28] *The Hilbert-Huang Transform in Engineering*. CRC Press, 2005. doi: 10.1201/9781420027532.

[29] Y. Chompoobutrgool and L. Vanfretti, "Identification of Power System Dominant Inter-Area Oscillation Paths," *IEEE Trans. Power Syst.*, vol. 28, no. 3, pp. 2798–2807, Aug. 2013, doi: 10.1109/TPWRS.2012.2227840.

[30] X. Xu *et al.*, "Tracking Periodic Voltage Sags via Synchrophasor Data in a Geographically Bounded Service Territory," in *2023 IEEE PES Grid Edge Technologies*, San Diego, Apr. 2023.

[31] J. F. Hauer and F. Vakili, "An oscillation detector used in the BPA power system disturbance monitor," *IEEE Trans. Power Syst.*, vol. 5, no. 1, pp. 74–79, Feb. 1990, doi: 10.1109/59.49089.

[32] E. Segerstrom, L. Vanfretti, C. Mishra, X. Xu, K. D. Jones, and R. M. Gardner, "Using Spectral Flatness to Detect and Label Power System Oscillations in the Presence of Intermittent Broadband Noise," in *2022 IEEE Power & Energy Society General Meeting (PESGM)*, Jul. 2022, pp. 1–5. doi: 10.1109/PESGM48719.2022.9917166.

[33] M. de Castro, L. Vanfretti, C. Mishra, X. Xu, and K. D. Jones, "Experiences with Dynamical Mode Decomposition for Wide-Area Mode Estimation," in *2022 10th Workshop on Modelling and Simulation of Cyber-Physical Energy Systems (MSCPES)*, May 2022, pp. 1–6. doi: 10.1109/MSCPES55116.2022.9770157.

[34] S. A. N. Sarmadi and V. Venkatasubramanian, "Inter-Area Resonance in Power Systems From Forced Oscillations," *IEEE Trans. Power Syst.*, vol. 31, no. 1, pp. 378–386, Jan. 2016, doi: 10.1109/TPWRS.2015.2400133.

[35] C. Mishra, A. Pal, J. S. Thorp, and V. A. Centeno, "Transient Stability Assessment (TSA) of Prone-to-Trip Renewable Generation (RG)-Rich Power Systems using Lyapunov's Direct Method," *IEEE Trans. Sustain. Energy*, pp. 1–1, 2019, doi: 10.1109/TSTE.2019.2905608.

[36] S. Wiggins, *Introduction to Applied Nonlinear Dynamical Systems and Chaos*, 2nd ed. in Texts in Applied Mathematics. New York: Springer-Verlag, 2003. Accessed: Sep. 10, 2018. [Online]. Available: //www.springer.com/us/book/9780387001777

[37] X. Jiang, H. Schattler, J. Zaborszky, and V. Venkatasubramanian, "Hard limit induced oscillations," in *Proceedings of ISCAS'95 - International Symposium on Circuits and Systems*, Apr. 1995, pp. 146–150 vol.1. doi: 10.1109/ISCAS.1995.521472.

[38] P. B. Reddy and I. A. Hiskens, "Limit-induced stable limit cycles in power systems," in *2005 IEEE Russia Power Tech*, Jun. 2005, pp. 1–5. doi: 10.1109/PTC.2005.4524706.