Automated Anomaly Detection in Distribution Grids Using μPMU Measurements

Mahdi Jamei*, Anna Scaglione*, Ciaran Roberts†, Emma Stewart†,
Sean Peisert†, Chuck McParland†, Alex McEachern‡,
*School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ, USA
†Lawrence Berkeley National Laboratory, Berkeley, CA, USA
‡Power Standards Laboratory, Alameda, CA, USA

Abstract—The impact of Phasor Measurement Units (PMUs) for providing situational awareness to transmission system operators has been widely documented. Micro-PMUs (μPMUs) are an emerging sensing technology that can provide similar benefits to Distribution System Operators (DSOs), enabling a level of visibility into the distribution grid that was previously unattainable. In order to support the deployment of these high resolution sensors, the automation of data analysis and prioritizing communication to the DSO becomes crucial. In this paper, we explore the use of μPMUs to detect anomalies on the distribution grid. Our methodology is motivated by growing concern about failures and attacks to distribution automation equipment. The effectiveness of our approach is demonstrated through both real and simulated data.

Index Terms—Intrusion Detection, Anomaly Detection, Micro-Phasor Measurement Unit, Distribution Grid

I. INTRODUCTION

The state vectors of the transmission grid are closely monitored and their physical behavior is well-understood [1]. In contrast, Distribution System Operators (DSOs) have historically lacked detailed real-time actionable information about their system. This, however, is set to change. As the distribution grid shifts from a demand serving network towards an interactive grid, there is a growing interest in gaining situational awareness via advanced sensors such as Micro-Phasor Measurement Units (μPMUs) [2].

The deployment of the μPMUs in isolation without additional data driven applications and analytics is insufficient. It is critical to equip DSOs with complimentary software tools that are capable of automatically mining these large data sets in search of useful, actionable information. There has been a lot of work focused on using PMU data at the transmission level to improve Wide-Area Monitoring, Protection and Control (WAMPC) [3], [4]. The distribution grid, however, is lagging in this respect. Due to inherent differences between operational behavior, such as imbalances and increased variability on the distribution and transmission grid, the algorithms derived for WAMPC at the transmission level are generally not directly applicable at the distribution level. Our work is aimed at addressing this issue. We focus on an important application of μPMU data in the distribution system: anomaly detection, i.e., behavior that differs significantly from normal operation of the grid during (quasi) steady-state. An anomaly can take

a number of forms, including faults, misoperations of devices or switching transients, among others, and its root cause can be either a natural occurrence, error or attack. The risk of cyber-physical attacks via an IP network has recently gained significant interest due to the increase in automation of our power gird via two-way communication. This communication is typically carried out on breachable networks that can be manipulated by attackers [5]. Even if an anomaly naturally occurs, it is important to notify the DSO to ensure proper remedial action is taken.

A. Related Work

The majority of published work in anomaly detection using sensor data, primarily SCADA and PMU data, has focused on the transmission grid. The proposed methods are typically data-driven approaches, whereby the measurements are inspected for abnormality irrespective of the underlying physical model. One such example, the *common path* data mining approach implemented on PMU data and audit logs at a central server, is proposed in [6] to classify between a disturbance, an attack via IP computer networks and normal operation. Chen et al., [7] derive a linear basis expansion for the PMU data to reduce the dimensionality of the measurements. Through this linear basis expansion, it is shown how an anomaly, which changes the grid operating point, can be spotted by comparing the error of the projected data onto the subspace spanned by the basis and the actual values. Valenzuela et al., [8] used Principal Component Analysis (PCA) to classify the power flow results into regular and irregular subspaces. Through analyzing the data residing in the irregular subspace, their method determines whether the irregularity is caused by a network attack or not. Jamei et al., [9] propose an intrusion detection architecture that leverages μPMU data and SCADA communication over IP networks to detect potentially damaging activities in the grid. These aforementioned algorithms are all part of the suite of machine learning techniques that the security monitoring architecture will rely on.

B. Our Contribution

μPMUs, due to their high sampling frequency, are a much richer data source in comparison to traditional Distribution Supervisory Control and Data Acquisition (DSCADA). In this

paper, we highlight the capability of μPMUs to detect transient events by proposing a set of rules for anomaly detection. The main advantages of our new approach are:

1) The underlying physical model of the data forms the basis in deriving the detection method; providing an interpretation of the event that is lacking in a model free approach.

2) The distribution grid is modeled allowing unbalanced loading and non-transposed lines. The rules are formulated in such a way that allow for distribution grids with neutral wires, and single-phase or two-phase laterals.

3) Quasi steady-state, rather than steady-state, is considered the norm for grid behavior.

4) Part of the proposed methodology only requires the phasor data stream of a single μPMU and is agnostic of the grid interconnection parameters, while the other part correlates the phasor streams across multiple μPMUs using electrical properties of the grid.

The detection method applicable to the measurements of a single μPMU is particularly attractive from a security perspective because, assuming that the algorithms are programmed onto the sensor itself, no network communication exchange is needed to obtain results. Therefore, the attacker will have to directly compromise the sensor to alter its response and erase evidence of a physical change.

In addition to defining these algorithms, we explore their effectiveness in the field via an actual deployment of μPMUs designed by our partners at Power Standards Lab. These devices output the three phase voltage and current phasors at specific locations on the distribution grid [10] at a rate of 120 Hz. The proprietary filtering implemented within the device, which differs from the options for the filter given in the C.37.118 standard [11], overcomes some of the technical obstacles limiting the deployment of conventional PMUs on the distribution grid. Technical obstacles in real world deployments include the presence of signal noise and smaller voltage angle differences. These devices are also inexpensive, which is a key feature for distribution sensors [12]. We also investigate the sensitivity of our rules with respect to a set of attack scenarios on the μPMU, and the grid connectivity data.

The remainder of the paper is organized as follows. Section II introduces the μPMU data model. Section III presents the π model of a distribution line, and the relationship between the voltage and current phasors in quasi steady-state conditions. Section IV forms the body of our work that concerns itself with the formulation of the rules and tracking the anomalies. The effectiveness of the proposed rules, and their sensitivity to partially-compromised data, are tested using real and simulated data in Section V. The conclusion follows in Section VI.

II. μPMU DATA MODELING

Assuming normal conditions, μPMUs are designed to report 120 samples per sec. of the three phase voltage phasors, denoted as $\mathbf{v}[k] \in \mathbb{C}^{3 \times 1}$, and the current phasors, $\mathbf{i}[k] \in \mathbb{C}^{3 \times 1}$, at specific points on the distribution network. To help understand the effect of transients on μPMU measurements, and

the difference between using phasor information compared to higher resolution time domain samples, we review the notion of a phasor, or complex envelope, in this section and tie its formal derivation to the actual implementation in practical μPMU sensors (including the C.37.118 standard [11]).

The *phasor* is, in signal theory, often referred to as the *complex envelope*, or the *complex baseband equivalent representation*, of an arbitrary signal $s(t)$. In its textbook derivation, it is obtained in two steps. In the first step, the frequency content of the signal at negative frequencies is removed, which leads to a complex signal called an analytic signal, $s^+(t)$. The time domain mapping from $s(t)$ to $s^+(t)$ is as follows[1]:

$$s^+(t) = \frac{1}{2}s(t) + \frac{j}{2\pi t} \star s(t) \qquad (1)$$

The second term in the sum is the *Hilbert transform* of the signal. The complex envelope, or phasor, can then be extracted by shifting down the analytic signal in the frequency domain and scaling it or, equivalently, demodulating and scaling the signal in the time-domain[2]:

$$\tilde{s}(t) = \sqrt{2}s^+(t)e^{-j2\pi f_0 t}. \qquad (2)$$

From (2) and (1) it is easy to see that:

$$s(t) = \sqrt{2}\Re[\tilde{s}(t)e^{j2\pi f_0 t}]. \qquad (3)$$

If power spectral density of the signal $s(t)$ is centered around f_0, then the complex representation is the *smoothest* signal that one can associate to $s(t)$. The mapping is one to one, and therefore, there is no loss of information. Actual μPMUs do not perform the envelope this way, as explained in the following section.

A. Practical Implementations of μPMUs

The Hilbert transform requires implementing a non-causal filter with infinite impulse response, hence, it is purely theoretical. While there are other ways of approximating the Hilbert filter, the simplest implementation of the μPMU is based on the following observations. Replacing the operator that takes the real part of the modulated complex envelope in (3), by the scaled summation of the modulated complex phasor and its conjugate and multiplying both sides of equation (3) by $\sqrt{2}e^{-j2\pi f_0 t}$, we have :

$$\sqrt{2}s(t)e^{-j2\pi f_0 t} = \frac{\tilde{s}(t) + \tilde{s}^*(t)e^{-j4\pi f_0 t}}{2} \qquad (4)$$

This representation is insightful since we can observe that a low-pass filter, $h(t)$, can extract the correct envelope $\tilde{s}(t)$ from the signal $\sqrt{2}s(t)e^{-j2\pi f_0 t}$ if and only if:

$$\tilde{s}(t) \star h(t) = \tilde{s}(t) \qquad (5)$$
$$\tilde{s}^*(t)e^{-j4\pi f_0 t} \star h(t) = 0 \qquad (6)$$

---

[1]The operator $\star$ stands for convolution.
[2]The reason for the $\sqrt{2}$ scaling is that the signal and its envelope have the same power and energy.

For the bandpass signal $s(t)$ centered around f_0, that has a limited support $[f_0 - W_1, f_0 + W_2]$, these two conditions are satisfied if:

$$W/2 \leq f_0, \quad W = W_1 + W_2. \tag{7}$$

Accordingly, the mapping becomes:

$$\tilde{s}(t) = \sqrt{2}(s(t)e^{-j2\pi f_0 t}) \star h(t) \tag{8}$$

This means that the phasor can be extracted without obtaining the analytic signal if (5)-(6) hold, and the filter frequency response $H(f)$ is flat within the bandwidth of the signal and isolates its spectrum, emulating an ideal low pass filter.

B. Information from μPMUs During a Transient

In normal conditions, the mains AC voltage and current are very close approximations of bandpass signals with a narrow support centered around the frequency $f_0 = 50/60$ Hz. In turn, (7) holds in the quasi steady-state scenario and for some types of transients. Because of the normally-small frequency support of the AC signals, the dynamic effects of the electrical wires are not apparent, and their effects can be approximated by a scaling and phase rotation equal to the amplitude and phase of their frequency response at the center frequency. This results in the well-known algebraic equations used in steady-state power systems analysis.

During a severe transient, however, the envelope that is obtained through (2) is not the signal envelope, not just because frequency content is effectively cut by $h(t)$, but also because of the component in (6) that is not zero and the spilling of its tail into the band selected by $h(t)$ distorts the content. In either case, however distorted, the signal that emerges out of the filter $h(t)$ is band-limited, and can therefore be sampled without aliasing at a rate of $2f_0$ Hz.

Our proposed methodology is to monitor whether the instantaneous μPMU measurements of voltage and current phasors belong to a complex hyperplane compatible with the algebraic steady-state version of Ohm's law for a three-phase unbalanced system. In the next section, we introduce the general Multi-Input Multi-Output (MIMO) representation of a distribution line and its quasi steady-state representation. These equations form the cornerstone of our anomaly detection rules.

III. DISTRIBUTION LINE MIMO MODELING

The π model of a distribution line that connects bus i to j is shown in Fig. 1 where $\boldsymbol{Y}_{ij}(f)$ denotes the three phase series admittance of the line (i, j) and $\boldsymbol{Y}_{ij}^{sh}(f)$ is the three phase shunt admittance matrix of that line. The modeling of the self and mutual impedance, rather than using the positive sequence representation, is essential for accurately modeling the distribution grid since it does not impose the assumption of a balanced system nor transposition of the lines. In addition, this representation enables us to include three phase, two phase, and single phase lines with/without neutrals by using their 3×3 phase frame matrix model that is obtained from the modified Carson's equations and Kron reduction [13].
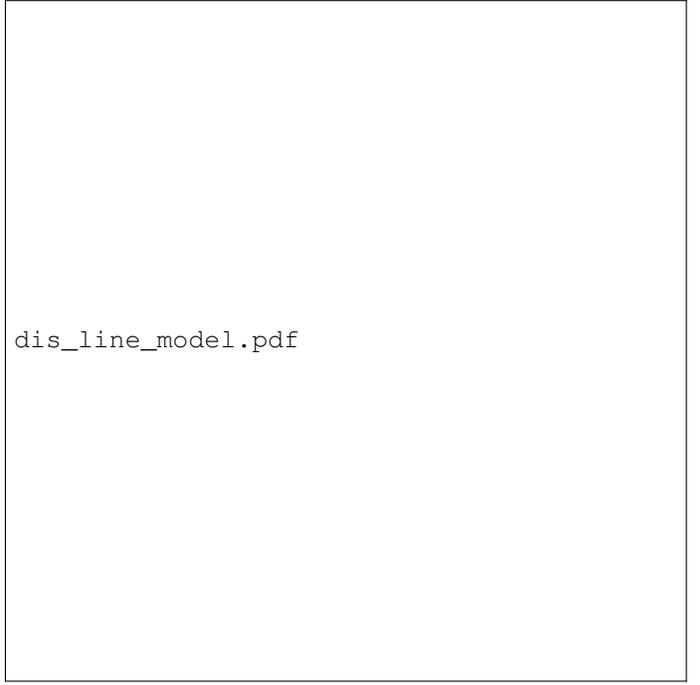

Fig. 1. π Model of the Distribution Line

It is well known in linear systems theory that the relationship between the voltage and the current signals in a passive Linear-Time Invariant (LTI) circuit with a known admittance matrix can be represented as the multiplication in the frequency domain and as a convolution in the time domain, which also takes MIMO form due to the three-phase modeling. The relationship also holds between the complex envelopes of the signals.

Denoting the baseband model of the line admittance matrices frequency response around the center frequency as $\mathbf{y}_{ij}(f) = \boldsymbol{Y}_{ij}^{sh}(f + f_0)H(f)$, $\mathbf{y}_{ij}^{sh}(f) = \boldsymbol{Y}_{ij}^{sh}(f + f_0)H(f)$, we have:

$$\boldsymbol{i}_{ij}(f) = (\mathbf{y}_{ij}^{sh}(f) + \mathbf{y}_{ij}(f))\boldsymbol{v}_i(f) - \mathbf{y}_{ij}(f)\boldsymbol{v}_j(f) \tag{9}$$

$$\mathbf{i}_{ij}(t) = (\boldsymbol{y}_{ij}^{sh}(t) + \boldsymbol{y}_{ij}(t)) * \mathbf{v}_i(t) - \boldsymbol{y}_{ij}(t) * \mathbf{v}_j(t) \tag{10}$$

where $\boldsymbol{y}_{ij}^{sh}(t)$ and $\boldsymbol{y}_{ij}(t)$ are the time-domain equivalence of the baseband shunt and admittance matrices respectively. We will use the following notation for brevity in the remainder:

$$\overline{\boldsymbol{Y}}_{ij}(f) \triangleq \boldsymbol{Y}_{ij}^{sh}(f) + \boldsymbol{Y}_{ij}(f)$$
$$\overline{\mathbf{y}}_{ij}(f) \triangleq \mathbf{y}_{ij}^{sh}(f) + \mathbf{y}_{ij}(f)$$
$$\overline{\boldsymbol{y}}_{ij}(t) \triangleq \boldsymbol{y}_{ij}^{sh}(t) + \boldsymbol{y}_{ij}(t)$$

Using this defined notation, (9) and (10) can be re-written as:

$$\boldsymbol{i}_{ij}(f) = \overline{\mathbf{y}}_{ij}(f)\boldsymbol{v}_i(f) - \mathbf{y}_{ij}(f)\boldsymbol{v}_j(f) \tag{11}$$

$$\mathbf{i}_{ij}(t) = \overline{\boldsymbol{y}}_{ij}(t) * \mathbf{v}_i(t) - \boldsymbol{y}_{ij}(t) * \mathbf{v}_j(t) \tag{12}$$

For the discrete time samples of the output of the μPMU, the counterpart of (12) is:

$$\mathbf{i}_{ij}[k] = \sum_{n=0}^{N-1} \overline{\boldsymbol{y}}_{ij}[n]\mathbf{v}_i[k-n] - \sum_{n=0}^{N-1} \boldsymbol{y}_{ij}[n]\mathbf{v}_j[k-n] \quad (13)$$

where the assumption is that $\boldsymbol{y}_{ij}^{sh}[n]$ and $\boldsymbol{y}_{ij}[n]$ are the samples of $\boldsymbol{y}_{ij}^{sh}(t)$ and $\boldsymbol{y}_{ij}(t)$, respectively, and that they are causal and have a finite support of N.

In the quasi-steady state condition, the fundamental frequency of the voltage and current signals is always changing, albeit slowly and over a very small range, because of load-generation imbalances, active power demand interactions, large generators inertia, and the automatic speed controllers of the generators [14]. As a result, the off-nominal frequency affects the phase angle captured by μPMUs. Fig. 2 shows the three phase unwrapped angle of the voltage phasor data captured by a μPMU installed at our partner utility grid, which clearly shows the grid is working at off-nominal frequency where the frequency drift is not even fixed over time though varies slowly. Mathematically, it is insightful to decompose

Fig. 2. Unwrapped Voltage Phasor Angle During Quasi-Steady State

the phasor $\mathbf{v}_i[k]$ and $\mathbf{i}_{ij}[k]$ as follows:

$$\mathbf{v}_i[k] = \hat{\mathbf{v}}_i[k]e^{j\beta_i[k]k}, \ \mathbf{i}_{ij}[k] = \hat{\mathbf{i}}_{ij}[k]e^{j\beta_i[k]k} \quad (14)$$

where $\hat{\mathbf{v}}_i[k]$ is the voltage phasor that is captured at nominal frequency, and $\hat{\mathbf{i}}_{ij}[k]$ is the current phasor after removing the exponential term due to $\beta_i[k]$ that captures the time-

varying drift in the frequency previously described. We can then write(13) in the following way:

$$\mathbf{i}_{ij}[k] = \sum_{n=0}^{N-1} \overline{\boldsymbol{y}}_{ij}[n]\hat{\mathbf{v}}_i[k-n]e^{j\beta_i[k][k-n](k-n)}$$
$$- \sum_{n=0}^{N-1} \boldsymbol{y}_{ij}[n]\hat{\mathbf{v}}_j[k-n]e^{j\beta_j[k][k-n](k-n)} \quad (15)$$

The variations of the $\hat{\mathbf{v}}_i[k]$ and $\beta[k]$ can be approximately neglected over N samples of the discrete time convolution in (13), i.e., $\hat{\mathbf{v}}_i[k-n] \approx \hat{\mathbf{v}}_i[k]$ and $\beta_i[k-n] \approx \beta_i[k]$ for $n = 0, \ldots, N-1$. Followed by this approximation, we can write:

$$\mathbf{i}_{ij}[k] \approx \sum_{n=0}^{N-1} \overline{\boldsymbol{y}}_{ij}[n]\hat{\mathbf{v}}_i[k]e^{j\beta_i[k](k-n)}$$
$$- \sum_{n=0}^{N-1} \boldsymbol{y}_{ij}[n]\hat{\mathbf{v}}_j[k]e^{j\beta_j[k](k-n)}$$
$$= \left(\sum_{n=0}^{N-1} \overline{\boldsymbol{y}}_{ij}[n]e^{-j\beta_i[k]n}\right)\mathbf{v}_i[k] \quad (16)$$
$$- \left(\sum_{n=0}^{N-1} \boldsymbol{y}_{ij}[n]e^{-j\beta_j[k]n}\right)\mathbf{v}_j[k]$$

Considering the following relationship:

$$\boldsymbol{Y}_{ij}(f+f_0)H(f) = rect(Tf)\sum_{n=0}^{N-1} T\boldsymbol{y}_{ij}[n]e^{-j2\pi nTf} \quad (17)$$

where $T = 1/120$ sec. is the sampling interval of the μPMU, we can introduce the following two matrices:

$$\overline{\boldsymbol{Y}}_{ij}(f_0, k) \triangleq \frac{1}{T}\overline{\boldsymbol{Y}}_{ij}\left(f_0 + \frac{\beta_i[k]}{2\pi T}\right)H\left(\frac{\beta_i[k]}{2\pi T}\right),$$
$$\boldsymbol{Y}_{ij}(f_0, k) \triangleq \frac{1}{T}\boldsymbol{Y}_{ij}\left(f_0 + \frac{\beta_j[k]}{2\pi T}\right)H\left(\frac{\beta_j[k]}{2\pi T}\right), \quad (18)$$

and therefore we can write (16) as follows:

$$\mathbf{i}_{ij}[k] = \overline{\boldsymbol{Y}}_{ij}(f_0, k)\mathbf{v}_i[k] - \boldsymbol{Y}_{ij}(f_0, k)\mathbf{v}_j[k] \quad (19)$$

Equation (19) is Ohm's law in the phasor domain under the quasi-steady state and comprises part of our anomaly detection algorithm derived in the next section. The corresponding equation for the steady state can be simply obtained by setting $\beta[k] = 0$ in (19).

From the analysis above, and specifically equation (19), it is clear that the effect of the quasi-steady state in the phasor domain is the modulation of the admittances (18). The effect is usually modest, as $\beta[k]$ is small. However, during a severe transient with frequency support in the order of 10 Hz, the relationship (19) with the matrices in (18) does not hold anymore, and an indication of transients in the phasor domain is, in reality, a manifestation of the full dynamic behavior in (13).

IV. Anomaly Detection

When the power grid is no longer in quasi-steady state conditions, the relationship between voltage and the current phasors manifests its full dynamic behavior. What we propose is to inspect the validity of the memoryless algebraic equations to flag the existence of transients in the grid. Equation (19) provides the basis for our rule.

A. Single μPMU Metric

Considering the line in Fig. 1, we first assume that two μPMUs are installed at both ends of a line, i.e., bus i and j, which means $\mathbf{i}_{ij}[k]$, $\mathbf{i}_{ji}[k]$, $\mathbf{v}_i[k]$ and $\mathbf{v}_j[k]$ are all available and they can exchange their information locally. This case helps explaining the method that follows, which makes use of data from a single μPMU.

We can cast the two equations that hold between the voltage and current at two ends of the three-phase line as follows:

$$\underbrace{\begin{pmatrix} \mathbf{i}_{ij}[k] \\ \mathbf{i}_{ji}[k] \end{pmatrix}}_{\mathbf{I}_{ij}[k]} = \underbrace{\begin{pmatrix} \overline{\mathbf{Y}}_{ij}(f_0,k) & -\mathbf{Y}_{ij}(f_0,k) \\ -\mathbf{Y}_{ij}(f_0,k) & \overline{\mathbf{Y}}_{ij}(f_0,k) \end{pmatrix}}_{\mathbf{H}_{ij}(f_0,k)} \underbrace{\begin{pmatrix} \mathbf{v}_i[k] \\ \mathbf{v}_j[k] \end{pmatrix}}_{\mathbf{V}_{ij}[k]} \quad (20)$$

Let us define, with $M \geq 6$, the following sample correlation matrices:

$$\mathbf{R}_{IV}[k] = \frac{1}{M-1} \sum_{m=0}^{M-1} \mathbf{I}_{ij}[k-m]\mathbf{V}_{ij}^H[k-m], \quad (21)$$

$$\mathbf{R}_{VV}[k] = \frac{1}{M-1} \sum_{m=0}^{M-1} \mathbf{V}_{ij}[k-m]\mathbf{V}_{ij}^H[k-m]. \quad (22)$$

Assuming that variations of $\mathbf{H}_{ij}(f_0,k)$ is negligible over a window of M samples, equation (20) implies that the following homogeneous equation holds in quasi-steady:

$$\big(\; \mathcal{I}_6 \; \vdots \; -\mathbf{H}_{ij}(f_0,k) \; \big) \underbrace{\begin{pmatrix} \mathbf{R}_{IV}[k] \\ \mathbf{R}_{VV}[k] \end{pmatrix}}_{\mathbf{R}_k} = \mathbf{0} \quad (23)$$

Proposition 1. *Correlation matrix $\mathbf{R}_k$ is approximately rank-1 during the quasi-steady state.*

Proof. During the quasi-steady state along a distribution line, the following assumptions hold with a very good approximation for $n = 0, 1, ..., M-1$:

$$\hat{\mathbf{v}}_i[k-n] \approx \hat{\mathbf{v}}_i[k], \quad \hat{\mathbf{v}}_j[k-n] \approx \hat{\mathbf{v}}_j[k]$$
$$\beta_i[k-n] \approx \beta_i[k], \quad \beta_j[k-n] \approx \beta_j[k], \quad (24)$$
$$\beta_i[k] \approx \beta_j[k]$$

Therefore, we can write:

$$\mathbf{R}_{VV}[k] = \frac{1}{M-1}(\mathbf{V}_{ij}[k] \otimes \mathbf{E}[k])(\mathbf{V}_{ij}^H[k] \otimes \mathbf{E}^H[k])$$
$$= \frac{1}{M-1}(\mathbf{V}_{ij}[k]\mathbf{V}_{ij}^H[k]) \otimes (\mathbf{E}[k]\mathbf{E}^H[k]) \quad (25)$$

where $\mathbf{E}[k]$ is defined as follows and represents the variations due to the off-nominal frequency:

$$\mathbf{E}[k] = \mathbf{1}_{6\times 1} \otimes \big(e^{-j\beta_i[k](M-1)} \quad \dots \quad e^{-j\beta_i[k]} \quad 1 \big) \quad (26)$$

We can then write:

$$\mathbf{E}[k]\mathbf{E}^H[k] = (\mathbf{1}_{6\times 1}\mathbf{1}_{1\times 6}) \otimes (M) = M\mathbf{1}_{6\times 6} \quad (27)$$

and therefore:

$$\mathbf{R}_{VV}[k] = \frac{M}{M-1}(\mathbf{V}_{ij}[k]\mathbf{V}_{ij}^H[k]) \otimes (\mathbf{1}_{6\times 6}) \quad (28)$$

which accordingly means that:

$$rank(\mathbf{R}_{VV}[k]) = rank(\mathbf{V}_{ij}[k]\mathbf{V}_{ij}^H[k]) \times rank(\mathbf{1}_{6\times 6}) = 1$$

Because:

$$rank(\mathbf{R}_k) = rank(\mathbf{R}_k^H \mathbf{R}_k)$$

We analyze the rank of $\mathbf{R}_k^H \mathbf{R}_k$ here, where:

$$\mathbf{R}_k^H \mathbf{R}_k = \mathbf{R}_{IV}^H[k]\mathbf{R}_{IV}[k] + \mathbf{R}_{VV}^H[k]\mathbf{R}_{VV}[k] \quad (29)$$

From the structure of (23) during the quasi-steady state, we have:

$$\mathbf{R}_{IV}[k] = \mathbf{H}_{ij}(f_0,k)\mathbf{R}_{VV}[k] \quad (30)$$

Substituting (30) in (29), we have:

$$\mathbf{R}_k^H \mathbf{R}_k = \mathbf{R}_{VV}^H[k]\boldsymbol{\mathcal{G}}_{ij}(f_0,k)\mathbf{R}_{VV}[k] \quad (31)$$

where:

$$\boldsymbol{\mathcal{G}}_{ij}(f_0,k) = \mathbf{H}_{ij}^H(f_0,k)\mathbf{H}_{ij}(f_0,k) + \mathcal{I}$$

Since the linear transformation of $\mathbf{R}_{VV}[k]$ does not increase its rank, and since we have already shown that $\mathbf{R}_{VV}[k]$ is of rank-1 during the quasi-steady state, we can conclude that:

$$rank(\mathbf{R}_k) = rank(\mathbf{R}_k^H \mathbf{R}_k) =$$
$$\leq rank(\mathbf{R}_{VV}[k]) = 1 \rightarrow \quad (32)$$
$$rank(\mathbf{R}_k) = 1$$

$\square$

Therefore, we can write $\mathbf{R}_k$ as follows:

$$\mathbf{R}_k \approx \sigma_1[k]\boldsymbol{u}_1[k]\boldsymbol{\nu}_1^H[k] \rightarrow \mathbf{R}_k\mathbf{R}_k^H \approx \sigma_1^2[k]\boldsymbol{u}_1[k]\boldsymbol{u}_1^H[k], \quad (33)$$

which means that all columns of $\mathbf{R}_k$ must lie in the same hyperplane. Overall, deviation from this behavior is an indicator that the line is experiencing a transient and/or that the three-phase measurements no longer lie over a single principal component. The detection can be automated by computing the following cost and tracking the fast changes in $x[k]$ [3]:

$$x[k] = \min_{\boldsymbol{u}} ||(\mathcal{I}_{12} - \boldsymbol{u}\boldsymbol{u}^H)\mathbf{R}_k\mathbf{R}_k^H||_F \text{ s.t. } ||\boldsymbol{u}|| = 1. \quad (34)$$

The solution $\boldsymbol{u}$ is the principal subspace of $\mathbf{R}_k$ and it is normally obtained by minimizing the orthogonal projection with respect to $\boldsymbol{u}$ that is expected to ideally go to zero in the stationary balanced case, and be close to zero for stationary unbalanced case.

Assume now that only a single μPMU is available at one end of a line. Using some reasonable approximations, it is still possible to apply this rule using the data stream from a single

---

[3]$||.||_F$ denotes the Frobenious norm.

μPMU. For bus i with a μPMU, the two voltage vectors at the two ends of each incident line to that bus are such that:

$$\mathbf{v}_j[k] = \underbrace{\text{diag}(\boldsymbol{\alpha}[k])}_{\tilde{\boldsymbol{\alpha}}[k]}\mathbf{v}_i[k], \qquad (35)$$

where $\boldsymbol{\alpha}[k]$ is a complex vector that relates the voltage phasors at the two ends. If we define now:

$$\boldsymbol{R}_{\text{iv}}^{(ij)}[k] = \frac{1}{M-1}\sum_{m=0}^{M-1}\mathbf{i}_{ij}[k-m]\mathbf{v}_i^H[k-m], \qquad (36)$$

$$\boldsymbol{R}_{\text{vv}}^{(ji)}[k] = \frac{1}{M-1}\sum_{m=0}^{M-1}\mathbf{v}_j[k-m]\mathbf{v}_i^H[k-m]. \qquad (37)$$

Assuming that $\boldsymbol{\alpha}[k]$ remains constant over a window of M samples during the quasi-steady state, we can write:

$$\boldsymbol{R}_{\text{vv}}^{(ji)}[k] \approx \tilde{\boldsymbol{\alpha}}[k]\boldsymbol{R}_{\text{vv}}^{(ii)}[k] \qquad (38)$$

Assuming that the variation of $\overline{\boldsymbol{Y}}_{ij}(f_0, k)$ is negligible over M samples during the quasi-steady state, we can use (19) to write:

$$\left(\begin{array}{c|c} \mathcal{I}_3 & -\overline{\boldsymbol{Y}}_{ij}(f_0,k) + \boldsymbol{Y}_{ij}(f_0,k)\tilde{\boldsymbol{\alpha}}[k] \end{array} \right) \underbrace{\begin{pmatrix} \boldsymbol{R}_{\text{iv}}^{(ij)}[k] \\ \boldsymbol{R}_{\text{vv}}^{(ii)}[k] \end{pmatrix}}_{\boldsymbol{R}_k^{(i)}} \approx \mathbf{0} \qquad (39)$$

Proposition 2. *Correlation matrix $\boldsymbol{R}_k^{(i)}$ is approximately rank-1 during the quasi-steady state.*

Proof. The proof is very similar to that of Proposition 1, and follows by assuming that $\hat{\mathbf{v}}_i[k-n] \approx \hat{\mathbf{v}}_i[k]$, $\beta_i[k-n] \approx \beta_i[k]$, and $\boldsymbol{\alpha}[k-n] \approx \boldsymbol{\alpha}[k]$ for $n = 0, 1, ..., M-1$, as well as using the structure of (39). $\qquad\square$

Proposition 2 suggests a similar criterion for a single μPMU to flag the exit from a quasi-steady state regime, and can be achieved by tracking the fast changes in $x[k]$ defined as follows for each individual incident line to that bus:

$$x[k] = \min_{\boldsymbol{u}} ||(\mathcal{I}_6 - \boldsymbol{u}\boldsymbol{u}^H)\boldsymbol{R}_k^{(i)}(\boldsymbol{R}_k^{(i)})^H||_F \quad \text{s.t.} \quad ||\boldsymbol{u}|| = 1 \qquad (40)$$

B. Multiple μPMUs Metric

In this section, we correlate the phasor data across multiple μPMUs to detect anomalies in the grid. The applied rule here extends the test of the quasi-steady state equations validity applying it to multiple μPMU measurements scattered over the grid. The rule can be hosted in the Distribution Management System (DMS), where the data from all the μPMUs could be available, or it can be decentralized over a set of anomaly detection engines, where each agent is responsible to check the anomalies on a dedicated part of the grid and sharing the edge information with the other agents.

We assume that knowledge of $\boldsymbol{Y}_{ij}^{sh}(f_0, 0)$ and $\boldsymbol{Y}_{ij}(f_0, 0)$ for each line in the perimeter monitored by that detector engine is available. We can take advantage of the fact that β_k is

small and consider their difference from (18) as a perturbation, which is equivalent to noise in the observation model. For brevity, when introducing the rule in this part, we will use $\boldsymbol{Y}_{ij}^{sh}$ and $\boldsymbol{Y}_{ij}$ to refer to $\boldsymbol{Y}_{ij}^{sh}(f_0, 0)$ and $\boldsymbol{Y}_{ij}(f_0, 0)$.

A natural way to relate the measurements across multiple devices is through the grid interconnection. We represent the vector of three-phase current injection and bus voltage phasors in the whole grid by $\mathbf{I}[k]$ and $\mathbf{V}[k]$, respectively, each vector contains $3B$ elements where B is the number of buses. We also define the vector $\mathbf{d}$ as follows:

$$\mathbf{d}[k] = \begin{pmatrix} \mathbf{I}[k] \\ \mathbf{V}[k] \end{pmatrix} \qquad (41)$$

The following set of algebraic equations are homogeneous during the steady-state and should be close to homogeneity during the quasi-steady state:

$$\mathbf{H}\mathbf{d}[k] = \mathbf{0}, \quad \mathbf{H} = \left(\begin{array}{c|c} \mathcal{I}_{3B} & -\boldsymbol{Y}_{3(B\times B)} \end{array} \right) \qquad (42)$$

where $\boldsymbol{Y}$ is the admittance matrix of the grid that connects the current injection to the bus voltages, and is constructed from the 3×3 line shunt and series admittance matrices introduced in Section III.

It should be noted that the elements of $\mathbf{d}$ in (42) are not all independent variables. The challenge is that, in general, we will have a very limited number of measurements of $\mathbf{d}$ from μPMUs, due to the cost limitations of deploying these devices. Let K denote the number of μPMUs that are available. We assume that each μPMU device has enough channels to measure the voltage and all incident current measurements of the bus on which it is installed. Hence, having a μPMU at bus i means that the following three phase voltage phasors and three phase current injection phasors for that bus are available:

$$[\mathbf{V}[k]]_i = \mathbf{v}_i[k], \quad [\mathbf{I}[k]]_i = \sum_{j:i\sim j}\mathbf{i}_{ij}[k] \qquad (43)$$

where $i \sim j$ denotes that bus i and j are connected. We can define a permutation matrix $\mathbf{T}$ that parses the vector $\mathbf{d}[k]$ into two sub-vectors corresponding to the non-available measurements, $\mathbf{d}_u[k]$, and the available measurements, $\mathbf{d}_a[k]$, that is:

$$\mathbf{T} = \begin{pmatrix} \mathbf{T}_u \\ \mathbf{T}_a \end{pmatrix} \ \rightarrow \ \mathbf{T}\mathbf{d} = \begin{pmatrix} \mathbf{d}_u \\ \mathbf{d}_a \end{pmatrix}, \ \mathbf{H}\mathbf{T}^T = \left(\begin{array}{c|c} \mathbf{H}_u & \mathbf{H}_a \end{array} \right) \quad (44)$$

where

$$\mathbf{T}_u \in \{0,1\}^{6(K'\times B)}, \ K' = B - K$$
$$\mathbf{T}_a \in \{0,1\}^{6(K\times B)}$$

Since $\mathbf{T}^T\mathbf{T} = \mathcal{I}$, we can rewrite (42) in the following form:

$$\mathbf{H}_u\mathbf{d}_u[k] + \mathbf{H}_a\mathbf{d}_a[k] = \mathbf{0}. \qquad (45)$$

Even though the equation is not exactly homogeneous, primarily due to the frequency drift discussed in Section III, it suggests that an estimate of $\mathbf{d}_u$ can be found through the following minimization:

$$\hat{x}[k] = \min_{\mathbf{d}_u} ||\mathbf{H}_u\mathbf{d}_u + \mathbf{H}_a\mathbf{d}_a[k]||_2^2. \qquad (46)$$

This is a least-square problem, with the well known solution:

$$\mathbf{d}_u^{\text{opt}}[k] = -\mathbf{H}_u^\dagger \mathbf{H}_a \mathbf{d}_a[k] \qquad (47)$$

where $(.)^\dagger$ denotes the pseudo-inverse operator. Consequently:

$$\hat{x}[k] = ||(\mathcal{I} - \mathbf{H}_u \mathbf{H}_u^\dagger)\mathbf{H}_a \mathbf{d}_a[k]||_2^2, \qquad (48)$$

In fact, $\hat{x}[k]$ can be interpreted as follows. If we pre-multiply both sides of (45) by the orthogonal projector onto the left null-space of $\mathbf{H}_u$ i.e., $(\mathcal{I} - \mathbf{H}_u \mathbf{H}_u^\dagger)$, by definition, for any $\mathbf{d}_u[k]$, we have:

$$(\mathcal{I} - \mathbf{H}_u \mathbf{H}_u^\dagger)\mathbf{H}_u \mathbf{d}_u[k] = \mathbf{0} \qquad (49)$$

and therefore, if (45) holds, it must also hold that:

$$(\mathcal{I} - \mathbf{H}_u \mathbf{H}_u^\dagger)\mathbf{H}_a \mathbf{d}_a[k] = \mathbf{0} \qquad (50)$$

which explains why the cost in (48) should be small when the system is in quasi-steady state. What we propose is to track the fast changes of the normalized $\hat{x}[k]$ defined as follows to detect anomalies in the data.

$$x[k] = \frac{||(\mathcal{I} - \mathbf{H}_u \mathbf{H}_u^\dagger)\mathbf{H}_a \mathbf{d}_a[k]||_2^2}{||\mathbf{d}_a||_2^2} \qquad (51)$$

However, (50) becomes trivial if $(\mathcal{I} - \mathbf{H}_u \mathbf{H}_u^\dagger) = 0$. This is the case when $\mathbf{H}_u$ is a full rank square matrix or a fat matrix, i.e., $K < \frac{B}{2}$, which has full row rank. This is actually the most common case because the number of μPMUs is going to be generally very small relative to the grid size. However, we can rely on the fact that the matrix $\mathbf{H}_u \mathbf{H}_u^H$ has a high condition number, due to the weak connectivity of the radial or weakly meshed networks, and relative homogeneity of the line parameters. Considering the singular value decomposition of the matrix $\mathbf{H}_u$:

$$\mathbf{H}_u = \mathbf{U}_u \mathbf{S}_u \mathbf{V}_u^H \qquad (52)$$

We define $\mathbf{u}_{u,2}$ to denote the last column of $\mathbf{U}_u$, representing the left singular vector that corresponds to the smallest singular value of the matrix $\mathbf{H}_u$. If (45) holds, we can expect that the $x[k]$ defined as follows should be small with smooth variations during the quasi-steady state:

$$x[k] = \frac{||\mathbf{u}_{u,2}^H \mathbf{H}_a \mathbf{d}_a||_2^2}{||\mathbf{d}_a||_2^2} \qquad (53)$$

Therefore, we propose $x[k]$ as the quantity to track its sudden changes in order to detect that a transient is present and apparent through the measurements $\mathbf{d}_a[k]$ or equivalently to spot when (45) no longer holds.

C. Fast Change Detection Method

The quantities defined in the last section for our metrics are tracked for fast changes, since their sudden variations are signatures of an anomaly. The quantities consist of the optimum cost functions defined in (34) and (40) for the steady-state criterion using the data from two adjacent μPMUs and a single μPMU, respectively, and finally, the value defined in (53) using the data from multiple μPMUs.

Using real and simulated data, we have confirmed that changes in the mean value during the quasi steady-state regime are extremely smooth while deviations from this mean value are minimal. This observation motivated us to consider changes in their mean value as the common statistical trademark of anomalies in all of these quantities.

To achieve fast detection of such changes, we propose to apply the sequential two-sided Cumulative Sum (CUSUM) algorithm [15]–[17]. To cast our problem in the CUSUM frame, we approximate the samples $x[k]$ for all the aforementioned quantities as outcomes of a Gaussian non-zero mean process $X[k]$:

$$X[k] = \mu_x[k] + \omega_x[k] \qquad (54)$$

where $\mu_x[k]$ is the mean of the process, and $\omega_x[k] \sim \mathcal{N}(0, \sigma_x^2[k])$ is zero-mean, additive Gaussian noise in the measurements with variance $\sigma_x^2[k]$. We assume that $\omega_x[k]$ are independent random variables for each k. Our intention is to detect sudden changes in the mean of the process, $\mu_x[k]$. Although there is temporal correlation among the observations, our objective, i.e., fast change detection in mean, justifies the relaxation that the random process has independent observation samples.

The algorithm decides between two possible hypotheses at time k: $\mathcal{H}_0$: no change is detected in the mean, $\mathcal{H}_1$: change is detected in the mean. The decision in the CUSUM algorithm is based on two *"instantaneous log-likelihood ratios"*, corresponding to upward and downward change of the mean, defined as follows:

$$\lambda_X^u[k] = +\frac{|\hat{\delta}_x|}{\hat{\sigma}_x^2[k]}\left(x[k] - \hat{\mu}_{0,x}[k] - \frac{|\hat{\delta}_x|}{2}\right) \qquad (55a)$$

$$\lambda_X^d[k] = -\frac{|\hat{\delta}_x|}{\hat{\sigma}_x^2[k]}\left(x[k] - \hat{\mu}_{0,x}[k] + \frac{|\hat{\delta}_x|}{2}\right) \qquad (55b)$$

where the superscripts u and d represent the variables corresponding to the "upward" and "downward" change detection respectively. $\hat{\delta}_x$ is the mean change estimate, and is initialized based on *"a priori"* knowledge. $\hat{\sigma}_x^2[k]$ is the random process variance estimate that is assumed to remain constant during the change and $\hat{\mu}_{0,x}[k]$ is the mean estimate. The mean estimate is obtained adaptively from normal ensembles using the exponential window as follows:

$$\hat{\mu}_{0,x}[k] = \rho\,\hat{\mu}_{0,x}[k-1] + (1-\rho)x[k] \qquad (56)$$

where $0 \le \rho \le 1$ determines the dependency of the mean estimator on the past samples compared to the current sample respectively. Accordingly, two cumulative sums, $M_X^u[k]$ and $M_X^d[k]$, and two decision functions, $G_X^u[k]$ and $G_X^d[k]$, are derived as follows.

$$M_X^u[k] = M_X^u[k-1] + \lambda_X^u[k] \qquad (57a)$$

$$M_X^d[k] = M_X^d[k-1] + \lambda_X^d[k] \qquad (57b)$$

$$G_X^u[k] = \max(G_X^u[k-1] + \lambda_X^u[k], 0) \qquad (57c)$$

$$G_X^d[k] = \max(G_X^d[k-1] + \lambda_X^d[k], 0) \qquad (57d)$$

The decision functions are then compared to a user-defined threshold, α_x. Then, hypothesis $\mathcal{H}_1$ for the upward or downward change in the mean is chosen if either $G_X^u[k] > \alpha_x$ or $G_X^d[k] > \alpha_x$, respectively. Depending on the change direction, the estimate of the anomaly start time is:

$$\widehat{k}_{c,x} = \underset{k_{0,x} \leq k_{c,x} \leq k-1}{\text{argmin}} M_X^{u/d}[k_{c,x}] \qquad (58)$$

where $k_{0,x}$ is the last detected change time index.

During an event, we expect to see multiple change points. Detection of multiple changes is done by resetting the decision functions and cumulative sums to zero after the change is detected, and continuing the dynamic rule for upcoming samples. The fast change anomaly is completed if no new changes are detected for a defined window of time.

V. NUMERICAL RESULTS

A. Field Recorded Data

Field recorded data is obtained from μPMUs that are installed in our partner utility medium voltage (12.47 kV) grid. The window of data that we use to validate our detection rules contains a voltage sag event recorded on the network. The captured three phase voltage phasor magnitude by one of the μPMUs during this event is shown in Fig. 3.

vol_sag.pdf

Fig. 3. Three Phase Voltage Phasor Magnitude Captured by a μPMU During the Voltage Sag

The μPMUs used in this use case are installed in pairs at opposite ends of a line. We first show, for a specific line the change detection using the optimum cost function defined in (34) with two μPMUs data, and then demonstrate that it would be possible to detect the changes with a single μPMU using (40). The size of the null space using two and one μPMU

during the voltage sag event is illustrated in Fig. 4(a) and Fig. 4(b) with $M = 32$ samples, respectively [4]. It can be observed that our metrics are able to effectively detect the anomaly as changes in the null-space, even when there is only one μPMU installed. The red markers are pointing to the detected anomaly start time, $\hat{k}_c$. Although the behavior in double and single μPMU metric is very similar, the change in the double-μPMU case is more pronounceable (noticing the scale on the vertical axis), obviously because it is augmenting the effects of the voltage sag on two μPMUs rather than one. It should be noted that since the value of the optimum cost
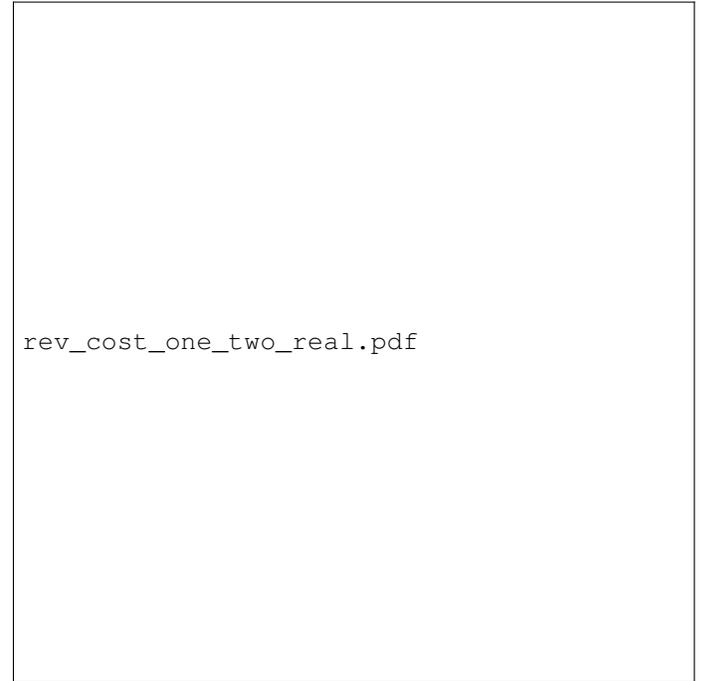

rev_cost_one_two_real.pdf

Fig. 4. Double and Single μPMU Metric Using Real Data for a Sag Event

function at time k depends on the last M phasor samples, there is a delay in the appearing and disappearing of the event in the cost functions.

B. Simulated Data

The IEEE 34-bus system is simulated in this section to test the rules on simulated data. The single-line diagram is shown in Fig. 5, where the bus numbers are restarted from 1, compared to original feeder, for simplicity. The test feeder data can be found at [18]. The simulation is performed in the time-series simulation environment of DIgSILENT [19] and the time-domain waveforms are forwarded to our simulated μPMU model to obtain the phasor representation. We used the two-cycle P class filter in C.37.118 standard [11] in our simulated μPMU since we did not have access to the proprietary filters implemented within the μPMU. This allows us to more closely mirror what a μPMU actually outputs in

[4] The voltage and current phasors are first converted to per-unit system assuming $S_b = 1$ MVA.

Fig. 5. IEEE 34-bus Test Feeder Single-Line Diagram

Fig. 6. Single μPMU Metric Using Simulated Data for SLG Fault

comparison to simulations that use FFT to extract the phasors. We consider a SLG fault scenario to assess how our rules perform. We assume that three μPMUs, i.e., $K = 3$, are available and are placed at bus 9, 19, and 31. The μPMU placement is done aiming to achieve the maximum change in (53) during a transient. The detailed formulation of the placement problem will be given in our future work.

A temporary SLG fault at 50% on phase A of line $(16, 17)$ occurs at $t = 0.5$ sec and is cleared at $t = 0.52$ sec, before the recloser opens. Fig. 6 shows the single μPMU metric derived in (40) using the current on lines $(9, 13)$, $(19, 20)$ and $(31, 32)$ with $M = 6$ [5]. The detected start time of the changes are marked for these three μPMUs during the event using the same detection threshold and "a priori" knowledge about the magnitude of the change. As it can be observed, the number of changes found in each line's metric is correlated with how severely the event affects that particular μPMU measurement. The metric can be formed for all the incident lines to the buses with μPMUs, however, we just show the cost relating to the aforementioned lines. Clearly, in each case $x[k]$, calculated for each line, experiences a sudden change due to the fault event being sufficiently severe to shift all of the lines into a transient state. It is important to note that since phasor calculation is based on the past two-cycles of time-domain data, and since the cost function is affected by the $M - 1$ previous phasor samples, as also the case for the real data, the anomaly appears and disappears with a delay in the cost function. To reduce the delay, the phasor can be estimated

[5]The voltage and current phasors are first converted to per-unit system assuming $S_b = 1$ MVA.

with less samples, e.g., one cycle instead of two cycles, but the associated accuracy decreases. In addition, M can be set to the minimum possible value, which was the case here. The metric across multiple μPMUs in (53) is also tested for this scenario and the corresponding metric, as well as the start time of the detected changes, are shown in Fig. 7. As we expected, the fault manifests itself as a sudden change across the μPMUs; a signature of an anomaly. The delay here is solely due to the two-cycle calculation of the phasor, therefore less in comparison to the single μPMU metric.

Metrics Sensitivity to μPMU Data Manipulation: An advantage of our metric across multiple μPMUs is that it is robust to some degree to the μPMU data manipulation. In fact, as long as some of μPMU data are not compromised, the metric reveals the presence of a transient. What matters here is whether the change is sufficiently severe to trigger the change detector. If the detector is set to be too sensitive, the "false alarms" increase, so there is an associated trade-off that should be considered when the detector is designed. However, this study is beyond the scope of this paper.

To test the performance of our metric in this situation, we consider again the single-line to ground fault previously explained for three cases where the attacker manipulates the data samples of the μPMU at bus 9 for the first case, the μPMU at bus 19 in the second case, and finally μPMUs at buses 9 and bus 19 together for the third case. The attacker manipulates the data during the fault by pointing to the last sample before the fault starts. Fixing the detector parameters for all the three cases, Fig. 8 shows the detected changes in the metric across multiple μPMUs. It can be observed that the metric shows sudden changes in all the three cases. For the first

Fig. 7. Detected Changes in the Metric Across Multiple μPMUs for Single-Line to Ground Fault


Fig. 8. Detected Changes in the Metric Across Multiple μPMUs for Single-Line to Ground Fault under Manipulated μPMU Data

and the second case, a few changes can still be detected by our CUSUSM detector. However, when both μPMU 9 and 19 are manipulated (case 3), the detector fails to spot the transient.

It is also clear that our single μPMU metric can only flag an anomaly if the corresponding μPMU data is not compromised. However, it should be mentioned that injecting false data at the device level is not an easy task for the attacker since the μPMU devices are designed to be read-only. In fact, man-in-the-middle attacks are more likely, which in this case, our single μPMU rule that is checked next to each device can detect the anomaly, even when the multiple μPMU rule is compromised.

Metrics Sensitivity to Grid Connectivity Manipulation: The attacker can falsify the grid connectivity data to mislead our intrusion detector. Our single μPMU rule is set up to be agnostic about the grid interconnection, and therefore is not affected by such an attack. However, the grid connectivity becomes important for our rule using all the μPMU data.

Case–1: We assume that grid connectivity data is manipulated, indicating that line $(25, 26)$ is out of service, while it is in service. The reason for choosing line $(25, 26)$ in our scenario is that no μPMU is installed on the lateral from bus 25 to 29 so we cannot confirm the line outage by just looking at the magnitude of the current phasor of that μPMU. However, if a line outage is the ground truth, we expect to see a sudden change in our metrics because of the transient induced by the line switching before the step variation due to the change of topology. This expected transient is not observable in our metric, as can be seen in Fig. 9, and therefore this indication of a line outage can be flagged as an intrusion.
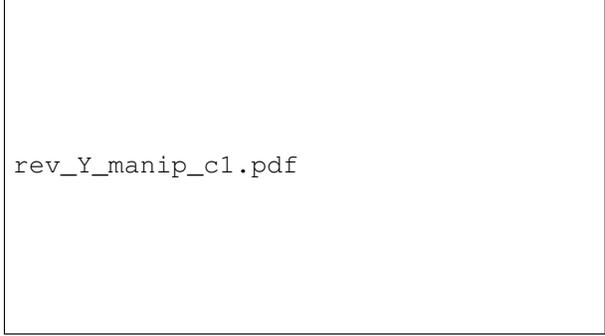
Fig. 9. Multiple μPMU Metric under Grid Topology Data Attack-Case 1

Case–2: In this case, we assume that a three-phase fault occurs on line $(25, 26)$ at $t = 0.4s$, resulting in the outage of the line due to the fuse operation at $t = 0.46s$. Therefore the lateral $25 - 29$ is deenergized. The attacker falsifies the data, indicating that the line $(25, 26)$ is still in service after the transient finishes. Fig. 10 shows the correct and the compromised metric. An attack is undetectable by our multiple μPMU metric if the attacker is able to compromise some data, while not changing the value of the metric significantly before and after the fault. In our case, as we can observe in Fig. 10, the value of the compromised metric (post-fault value) is close to the ground truth (pre-fault value), and therefore this rule is not able to detect the associated data manipulation in this case.

Fig. 10. Multiple μPMU Metric under Grid Topology Data Attack-Case 2

VI. Conclusion

In this paper, we have formulated two types of metrics for anomaly detection using μPMU measurements for the distribution grid. The first metric requires phasor data of a single device and is agnostic about the grid model. The second metric correlates data across multiple μPMUs to flag events. Anomalies are detected applying the sequential CUSUM algorithm in search for a sudden change in their mean value. To be more realistic, the rules have been formulated recognizing that the grid is in the quasi-steady state during the normal operation rather than operating at nominal frequency. In addition, the distribution grid has been modeled allowing for unbalanced load, absent of the assumption of transposed lines and considering the fact that there may be two-phases or single phase laterals in the grid. The proposed rules have been verified using the real and synthetic data.

For future work, we will extend our rules to robustify the intrusion detector. For example, case-2 in the grid connectivity data manipulation can be identified as anomaly by checking the post and pre-fault power flow measured by the μPMU upstream of the event location. We will also present the optimal placement of the μPMU sensors with respect to our multiple μPMU metric to achieve the maximum change in the metric when an anomaly happens, as well as methods to localize the events that caused the transient. We will also propose a formal architecture for the implementation of our anomaly detectors in the context of cyber-physical security, where the analysis results of the μPMU data will be tied with the monitored DSCADA traffic for intrusion detection.

References

[1] A. Abur and A. G. Exposito, *Power system state estimation: theory and implementation*. CRC press, 2004.
[2] J. H. Eto, E. M. Stewart, T. Smith, M. Buckner, H. Kirkham, F. Tuffner, and D. Schoenwald, "Scoping study on research and priorities for distribution-system phasor measurement units," 2015.
[3] A. Phadke and R. M. de Moraes, "The wide world of wide-area measurement," *Power and Energy Magazine, IEEE*, vol. 6, no. 5, pp. 52–65, 2008.
[4] V. Terzija, G. Valverde, D. Cai, P. Regulski, V. Madani, J. Fitch, S. Skok, M. M. Begovic, and A. Phadke, "Wide-area monitoring, protection, and control of future electric power networks," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 80–93, 2011.
[5] J. Slay and M. Miller, *Lessons learned from the maroochy water breach*. Springer, 2008.
[6] S. Pan, T. Morris, and U. Adhikari, "Developing a hybrid intrusion detection system using data mining for power systems," *Smart Grid, IEEE Transactions on*, vol. 6, no. 6, pp. 3104–3113, 2015.
[7] Y. Chen, L. Xie, and P. R. Kumar, "Dimensionality reduction and early event detection using online synchrophasor data," in *Power and Energy Society General Meeting (PES), 2013 IEEE*. IEEE, 2013, pp. 1–5.
[8] J. Valenzuela, J. Wang, and N. Bissinger, "Real-time intrusion detection in power system operations," *Power Systems, IEEE Transactions on*, vol. 28, no. 2, pp. 1052–1062, 2013.
[9] M. Jamei, E. Stewart, S. Peisert, A. Scaglione, C. McParland, C. Roberts, and A. McEachern, "Micro synchrophasor-based intrusion detection in automated distribution systems: Towards critical infrastructure security," *IEEE Internet Computing*, vol. 20, no. 5, 2016.
[10] "PQube3 information," http://www.powersensorsltd.com/PQube3.php, accessed: 2016-01-15.
[11] "Ieee standard for synchrophasor measurements for power systems," *IEEE Std C37.118.1-2011 (Revision of IEEE Std C37.118-2005)*, pp. 1–61, Dec 2011.
[12] A. von Meier, D. Culler, A. McEachern, and R. Arghandeh, "Micro-synchrophasors for distribution systems," in *Proc. IEEE PES Innovative Smart Grid Technologies Conference (ISGT)*, 2014, pp. 1–5.
[13] W. H. Kersting, *Distribution system modeling and analysis*. CRC press, 2012.
[14] A. G. Phadke and J. S. Thorp, *Synchronized phasor measurements and their applications*. Springer Science & Business Media, 2008.
[15] E. Page, "Continuous inspection schemes," *Biometrika*, pp. 100–115, 1954.
[16] M. Basseville, I. V. Nikiforov *et al.*, *Detection of abrupt changes: theory and application*. Prentice Hall Englewood Cliffs, 1993, vol. 104.
[17] T. L. Lai, "Sequential changepoint detection in quality control and dynamical systems," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 613–658, 1995.
[18] http://ewh.ieee.org/soc/pes/dsacom/testfeeders/index.html.
[19] D. P. F. Manual and D. PowerFactory, "Version 14.0," *DIgSILENT GmbH, Gomaringen, Germany*, 2009.