Real Time Stream Mining based Attack Detection in Distribution Level PMUs for Smart Grids

Prithwiraj Roy[†], Shameek Bhattacharjee[‡], and Sajal K. Das[†]
[†]Department of Computer Science, Missouri University of Science and Technology, Rolla, USA
[‡]Department of Computer Science, Western Michigan University, Kalamazoo, USA
E-mail: przhr@mst.edu; shameek.bhattacharjee@wmich.edu; sdas@mst.edu

***Abstract*—Reliable automation of smart grids depends on decisions based on situational awareness extracted via real time system monitoring and accurate state estimation. The Phasor Measurement Units (PMU) at distribution and transmission layers of the smart grid provide high velocity real time information on voltage and current magnitudes and angles in a three phase electrical grid. Naturally, the authenticity of the PMU data is of utmost operational importance. Data falsification attacks on PMU data can cause the Energy Management Systems (EMS) to take wrong decisions, potentially having drastic consequences on the power grid's operation. The need for an automated data falsification attack detection and isolation is key for EMS protection from PMU data falsification. In this paper, we propose an automated distributed stream mining approach to time series anomaly based attack detection that identifies attacks while distinguishing from legitimate changes in PMU data trends. Specifically, we provide a real time learning invariant that reduces the multi-dimensional nature of the PMU data streams for quick big data summarization using a Pythagorean means of the *active power* from a cluster of PMUs. Thereafter, we propose a methodology that learns thresholds of the invariant automatically, to prove the predictive power of distinguishing between small attacks versus legitimate changes. Extensive simulation results using real PMU data are provided to verify the accuracy of the proposed method.**

***Index Terms*—Smart Grid Security, Phasor Measurement Units Security, Big Data Management, Anomaly Detection.**

I. Introduction

Traditionally, power grid operators had limited information about dynamically varying system states in the grid. Many major faults in the grid are usually preceded by ephemeral warning signs (e.g., voltage sags) that Supervisory Control And Data Acquisition (SCADA) measurements (with data resolution of several seconds) could not capture as shown in [11]. To alleviate this problem, PMUs are deployed to capture fine grained high resolution time series data. These PMUs form the crucial endpoint device for the PMU Infrastructure, one of the key cornerstones of the modern smart grid design. Furthermore, with the increasing market penetration of Distributed Energy Resources (DERs) (e.g. solar panels), two-way electricity flows, and novel loads (such as electric vehicles), the grid requires real time grid monitoring, making the integrity of PMU data streams of strategic importance.

The PMUs record time-synchronized measurements of voltage, current, phase angle and frequency (collectively known as synchrophasor data) and sends it to an aggregator called Phasor Data Concentrator (PDC). The PDC, in turn, relays such data to a control center, allowing grid operators to localize and infer the type, time and location of a fault or disturbance as well as support critical control-actuation operations such as state estimation, maintain optimal power flow, based on the measured PMU data streams. The architecture of a typical PMU-PDC infrastructure is shown in Fig. 1.

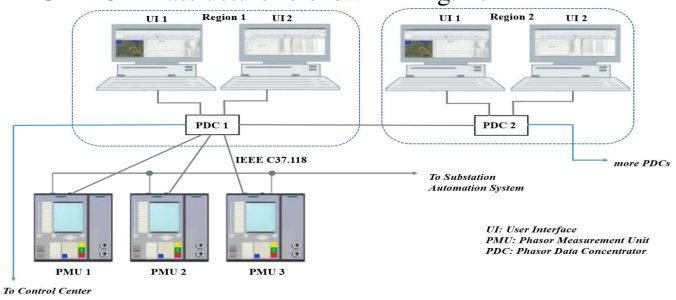

Figure 1: Architecture of a PMU Infrastructure.

However, in recent years, power distribution systems have faced cyber-attacks, threatening their security, reliability of operations. The report of US National Research Council highlights potential multi-state blackouts as a result of coordinated False Data Injection (FDI) attacks on power systems [2]. Such an attack on the Ukrainian power grid resulted in the loss of service for approximately 225,000 customers in three different territories which lasted for several hours [3]. Stuxnet worm has directly affected more than 100,000 industrial components [4]. However, the widely accepted IEEE C37.118-2 protocol for synchrophasor communication is highly vulnerable to cyber-attacks [6], [8]. In fact, most synchrophasor data transmission happen on non-reliable and insecure IP networks. Heavy encryption is not possible due to the latency critical nature of PMU data applications, thus increasing the chances of FDI attacks. This motivates the need for anomaly based intrusion detection in PMUs. While some existing research [7], [14] offer solutions, they have the following limitations: [7] focus on transmission layer PMUs, where data is very stable, thus making anomaly detection easy. The [5] considers the problem of only voltage data falsification, which is stable and hence easy to detect, ignoring current data falsification.

In this paper, we first discuss multiple attack strategies for data falsification attacks in PMUs. Then, we propose a process variable selection that reduces the dimensionality

of the anomaly detection problem. Then, we use a ratio of harmonic means to arithmetic means of the active power derived from the synchrophasor data sent from PMUs as a data-driven 'invariant' for anomaly detection. Specifically, we find the appropriate spatial and temporal considerations of the PMU network, such that an 'invariant' is highly stable under no attacks but shows unique changes under various kinds of data falsification attacks. Then, we propose a two-tier threshold based detection criterion involving stateless and stateful residuals of the anomaly detection metric, that better improve the false alarm versus detection sensitivity trade-off. The two-tier detector uses the sum of long term residuals from the median absolute deviation of the ratio based metric observed over the training phase. Finally, we validate our work by using real PMU datasets collected from Lawrence Berkley National Lab across 12 days.

The main benefits of our approach are to provide a practical framework for compromised PMU identification that (i) real time, light weight, semi- supervised, (ii) enables quick identification, and (iii) simultaneously works for a variety of data falsification attack types.

The rest of this paper is organized as follow. Section II, discusses related work. Section III discusses PMU dataset description, system and threat models, Section IV presents the proposed detection framework, Section V and VI offers experimental results and conclusions, respectively.

II. Related Work

In [7], a mechanism based on continuous monitoring of phase-wise equivalent transmission line impedance was proposed, for detecting data falsification on the voltage data from transmission system PMUs. However, they require two PMUs deployed at both ends of the transmission line and one of them needs to be honest. More importantly, we found that the PMU data streams at transmission level were inherently stable making anomaly detection a less challenging problem.

In [5] a Support Vector Machine (SVM) was used for detection, against a mirroring spoof attack strategy on the voltage data at distribution level PMUs. However, only falsification of voltage stream was considered which is relatively stable and makes anomaly detection less challenging.

The [9] proposed a decision tree based anomaly detection scheme to differentiate between normal tripping and malicious tripping by training on specific attack samples. However, it is not feasible to generate 100% of all the possible legitimate line tripping cases for training in [9].

In [14] a smart Time Synchronization Attack (TSA) based on GPS spoofing was shown to be equivalent to modifying the phase angle measurement from PMUs. However, they have not discussed any defense mechanism.

In [13] a density-based local outlier factor (LOF) analysis was used to detect the anomalies among the data, to describe spatio-temporal outliers among all the synchrophasor measurements from the grid. However, this method might not be able to detect attacks in real time and in their proposed method the authors have only considered an attack on voltage magnitude.

A critical analysis of all previous works on the detection of PMU data falsification revealed that current data falsification for PMU streams was not investigated. Furthermore, we found that unlike transmission level PMUs, the distribution level PMU's current synchrophasor data shows high dynamic variations in benign conditions, making anomaly detection challenging. Finally, all previous defenses are stream specific in the sense that they only work for either voltage or phase falsification. Since each PMU contains 4 streams and has 3 phases, a stream specific defense will require 12 different defense models that need complex cross-coordination.

III. System and Threat Models

A. PMU System Architecture

Here we first describe the PMU infrastructure network architecture. Most PMUs measure time-stamped voltages and current magnitudes and their phase angles denoted by $V_t(j)$, $I_t(j)$, $\theta_t^V(j)$, $\theta_t^I(j)$ respectively, where t is the time stamp and j is the j-th phase. These PMUs are deployed at strategic points of the transmission and distribution layers of the smart grid. Each PMU sends its data to a regional decentralized data aggregator known as PDC. The corresponding PDC in turn relays the aggregated data from multiple PMUs to a Local Controller Center (LCC). Various local controller centers communicate with each other forming a wide network for synchronizing local and global PMU data. In this paper, we are specifically interested in a *decentralized anomaly detection that runs on a PDC or a LCC and facilitates early attack detection from a bunch of PMUs that are geographically proximate in terms of the PMU network.*

Dataset Description: We use a dataset collected from the Power Standards Lab (PSL) at LBNL in Berkeley, CA, which developed high-precision μ-PMUs for showing how steps in our framework related to a real PMU system. The LBNL dataset contains three μ-PMUs that are deployed at multiple utility and LBNL campus locations on 12 kV distribution grid. The μ-PMU devices are named as: Grizzly, A6, and Bank514 in the dataset. Each μ-PMU device produces 12 streams of 120 Hz high-precision values with timestamps accurate to 100 ns (the limit of GPS). The 12 streams of data include both magnitude and phase angle for both voltage and current for all three phases on a true distribution network [10].

B. Threat Model

This section describes three features characterizing the threat model (e.g.,attack types, falsification margins, and falsification distributions) that can be employed by organized adversaries.

Threat Model Scope: PMU being a comparatively new research area, real malicious data samples from PMUs are hard to find. Therefore we generated the malicious samples by applying the three aspects of adversarial strategy over the real data. We have ensured that the falsification strategies used, do not favor or suit our proposed defense mechanism.

In simple electrical terms, the term load is equivalent to the current magnitude in each phase. Typically, in any phase,

there could be two possibilities of load change. Either there could be an increase or decrease in current, both creating an imbalance in the power grid. An increase in the phase current will cause the phase voltage to drop. If the current increases too much, then the phase is shed or the load is switched to other phases. Imbalance can also occur if the current drops in any phase, making the system inefficient in terms of utilization. *This creates a motivation to falsify current measurements.*

Attack Types: Attacks can be categorized in different types based on how data is changed across multiple PMUs. Organized adversaries can falsify data from single or multiple compromised PMU(s) simultaneously. Based on the objective and intent of the adversary, any of the four streams (*Voltage Magnitude, Voltage Angle, Current Magnitude, Current Angle*) of each phase can be falsified.

We assume the adversary falsify the 'current magnitude'. Let $I_t^i(act)$ be the actual current magnitude of i-th PMU at time t, while I_t^i be it's reported value. Under no attacks, the actual and reported value $I_t^i = I_t^i(act)$, while under attacks the reported value I_t^i can be biased by the following ways:

Deductive: In this case I_t^i from the i-th compromised PMU at time t is changed to $I_t^i(act) - I_{\delta_t}$, where $I_{\delta_{min}} \leq I_{\delta_t} \leq I_{\delta_{max}}$, for $I_{\delta_{min}} > 0$ is the false bias. Deductive attacks disrupt the efficiency of the grid by reducing the power utilization.

Additive: An additive attack can be launched by a rival utility to make the control center believe in a sudden increase in load which might lead to load shedding in that particular phase. Therefore, for additive falsification, the modified attack sample is $I_t^i = I_t^i(act) + I_{\delta_t}$ from a compromised PMU.

Alternating Attack: The adversary alternates between additive and deductive falsification for equal time duration over the time domain with the same average bias value of I_{δ_t}. In such a case, the effect of additive and deductive falsification will cancel each other's effect over a particular time period making it hard over most device specific statistical anomaly detectors to detect such attacks.

FDI Margin: We consider $I_{\delta_{avg}}$ as the *average margin of false data* for each compromised PMU. The strategic value of $I_{\delta_{avg}}$ is selected by an adversary as some value that ensures some minimum damage to the system. We keep this as an uncontrolled variable to test detection sensitivity since there could be various applications of PMU data. We consider that the attack is uniformly distributed $I_{\delta_t} \in [I_{\delta_{min}}, I_{\delta_{max}}]$ that does not change the resultant shape of the load distribution drastically, making it a smarter and less obvious attack.

Attack Strategies: We consider three types of attack strategies: *(a) Step attack:* In this case the adversary modifies all samples by $I_{\delta_{avg}}$ in the attack period. *(b) Ramp attack:* Here adversary gradually increases the I_{δ_t} in each time slots to reach $I_{\delta_{max}}$ and then again gradually decreases I_{δ_t} [7]. *(d) Mirroring:* Here attacker captures I_t^i for some period and then replaces the actual current measurements with the mirror image of captured I_t^i.

IV. Proposed Framework

The proposed framework is divided into four steps: (1) Propose a derived process variable (active power from synchrophasor measurements) that will form the basis for the anomaly detection process; (2) Design the invariant metric by optimizing spatial and temporal granularities of the process variable; (3) Design of a stateless and a stateful detection thresholds that identify the normal region of invariants under no attacks from the training set, such that false alarms are not drastically sacrificed for detection sensitivity improvement; (4) Determine the detection criteria parameters, based on learning from the training and cross validation steps, and apply it on the testing set, such that the predictive accuracy of distinguishing between legitimate changes versus malicious attacks is improved.

A. Choosing Process Variable for Anomaly Detection

Given the high velocity of the data, quick lightweight analytical tools are required for big data summarization to ensure the security and integrity of the dataset. However, due to 12 streams of data per PMU, the variety of data is extremely large. With multiple data streams per PMU, the anomaly monitoring of all these streams separately increases the computational cost and latency in anomaly detection analytics.

Hence, we propose the active power calculated from synchrophasor data streams per PMU, as the process variable over which the data driven invariant is designed. The active power ($P(j)$) per phase from PMU measurements are calculated using the following standard power equations:

$$P(j) = V(j)I(j)\cos\theta(j) \tag{1}$$

where $j \in \{1, 2, 3\}$ denote the phases and $V(j), I(j), \theta(j)$ are voltage magnitude, current magnitude, and angle difference between voltage and current phases respectively, for the j-th phase. This reduces the complexity of the monitoring each stream separately unlike existing works.

Another advantage is that any deliberate falsification of the voltage or current (both in terms of magnitude and phase) will impact the active power, and hence we can potentially detect an attack on any of the data streams from PMUs. Therefore, for our anomaly detection, we propose to use the phase wise monitoring of the active power $P(j)$ as a starting point. To clean the raw dataset [10] we have also applied 95% Winsorization before proceeding with our model.

B. Achieving an Invariant for Anomaly Detection Metric

For real time anomaly detection in CPS, it has been established that a metric which is invariant under normal operating conditions (without any attack) is ideal for attack detection. However, unlike tightly controlled industrial CPS applications, the distribution level synchrophasor data is affected by randomness and renewable power outputs and consumption patterns, causing traditional statistical invariants to have high randomness. As shown in Fig. 2a the arithmetic mean of the time series is not stationary. Prior works such as [12] propose the use of derived smoothing statistics of the arithmetic mean

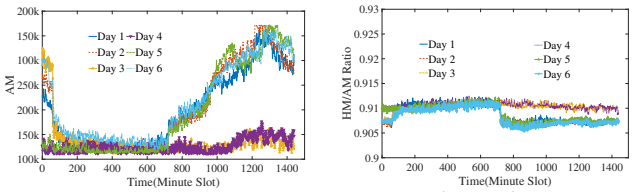

(a) AM. (b) HM/AM.

Figure 2: Illustrations of AM and HM/AM: Weekdays: Day 1,2,3,6; weekends: 4,5.

(such as ARMA, EWMA, CUSUM control charts) for time series anomaly detection. However, Fig. 2a shows that time series of PMUs active power fluctuates greatly over time windows, making it difficult to distinguish legitimate changes from a malicious one. Any moving average or smoothing technique either loses sensitivity for a small margin of attacks (since the moving average does not reflect the changes beyond already existing deviations or has large false alarms.

Let $\boldsymbol{P_t} = [P_t^1, ..., P_t^N]$ denote the active power from N PMUs at time slot t. We have taken second wise average of active power for our analysis, thus $t = 1$ second. Recently, in [1], we have shown that the ratio of harmonic mean and arithmetic mean of positively correlated variables exhibit invariance in their time series even when the individual means show non-stationarity. Additionally, [1] showed that the data perturbations in any variable cause the ratio to lose its invariance. However, this stability is guaranteed for appropriately correlated variables only. Hence, our primary goal is to investigate how to apply this on active power from PMUs. To this aim, we need to find the appropriate spatial and temporal granularity that maximizes the correlation between active powers on a given phase across different PMUs, which ensures invariance in the following metric:

Harmonic to Arithmetic Mean Ratio: Let the harmonic mean (HM_t) and arithmetic mean (AM_t) of P_t at time slot t be defined as:

$$HM_t = N(\sum_{i=1}^{N} P_t^i)^{-1} \quad \text{and} \quad AM_t = \frac{1}{N}\sum_{i=1}^{N} P_t^i. \tag{2}$$

We calculate HM_t and AM_t for slot t over a time window T of length n slots. Then we calculate the average HM_t to AM_t ratio, $Q^r(T)$, at the end of each window as follows:

$$Q^r(T) = \frac{\sum_{t=1}^{n} HM_t}{\sum_{t=1}^{n} AM_t} \tag{3}$$

where $0 \leq Q^r(T) \leq 1$, as $HM_t \leq AM_t$.

*1) **Optimizing the Spatial Granularity**:* Intuitively, a group of PMUs connected to the same feeder or serving proximate geographical areas should exhibit some interdependence in the synchrophasor data streams. We use the pairwise Pearson correlation coefficient to identify clusters that show some level of positive correlation. The higher the desired level of invariance, the higher is the required level of positive correlation. We calculate hourly Pearson's correlation among all pairs of PMUs in the training set to find groups

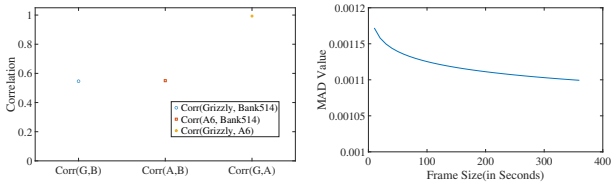

(a) Correlation among the PMUs. (b) MAD over time windows.

Figure 3: PMU clustering and MAD over time window.

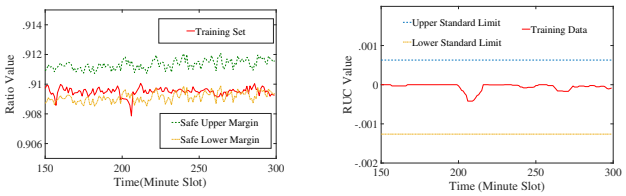

(a) Stateless Residual. (b) Stateful Residual.

Figure 4: Stateless and Stateful Residuals for $\epsilon = 0.85$.

having a maximum correlation. In the LBNL dataset, the mean of hourly correlations between Grizzly, A6 is 0.98; between Grizzly and Bank514 is 0.54; between A6 and Bank514 is 0.55 as shown in Fig. 3a. It is evident from the mean correlations that Grizzly and A6 are connected to the same feeder and thus can be considered in a single cluster. The average correlation identifies PMUs to be clustered under one instance of the anomaly detection technique.

*2) **Optimizing Temporal Granularity**:* Now we focus on choosing the appropriate time granularity over which the ratio metric is calculated. The time granularity should be such that the invariance in the ratio metric is maximized (i.e., minimize the measure of dispersion in the ratio statistic). Therefore, we solve the following search problem:

$$T = \underset{T^*}{\text{argmin}} MAD(Q^r(T^*)) \tag{4}$$

where $MAD(Q^r(T^*))$ is the median absolute deviation (MAD) of the resulting ratio time series with candidate time granularity $T^* \leq 360$ seconds. We choose T^* that minimizes the MAD of the ratio time series (shown in Fig. 3b).

C. Stateless and Stateful Residual based Threshold Design

Intuitively, The anomaly detection needs to identify a proximate spatial region around the ratio time series that specifies the behavior of the invariant under no attacks. Usually, a threshold is calculated by tracking the difference between the actual time series value and its smoothed value over time. However, a simple threshold based approach, cannot decrease both false alarms and missed detections simultaneously [12]. Hence, we put forward a two-tier approach with stateless and stateful residuals.

*1) **Stateless Residuals:*** The stateless residual is an instantaneous residual per time window T. Our method computes the mean μ_r and median absolute deviation, m_Q, from the probability distribution of ratio values $Q^r(T)$ for each PMU cluster (shown in Fig. 4a).

Unlike our previous work [1], we propose the use of Median Absolute Deviation (MAD) as a scale parameter for designing the stateless residual rather than the standard deviation (SD), because MAD is more robust to outliers. Thus, MAD can automatically adjust the resultant safe margin under errors and outliers in the training. The MAD is robust than SD since it is based on a squared error from the mean, so a finite number of outliers can influence SD easily compared to MAD, thus reducing sensitivity to small attack strengths.

Stateless residual is parameterized as $\kappa = \epsilon m_Q$ where $\epsilon \in (0, 4]$, such that $\kappa \in (0, 4m_Q]$ and m_Q is the MAD. Intuitively, larger κ values produce wider safe margins, thus reducing false alarms but increasing misdetection and vice-versa. Hence, a trade-off is necessary for selecting a threshold that will automatically generalize into lowering false alarms while not sacrificing the detection sensitivity, which is taken care of by the stateful residual as shown in Fig. 4b.

Our framework calculates a parameterized 'stateless residual' with two values; $\Gamma_l(T)$, and $\Gamma_h(T)$ around the observed instantaneous ratio values $Q^r(T)$, on every time window on the training dataset, such that:

$$\Gamma_h(T) = Q^r(T) + \epsilon m_Q. \tag{5}$$

$$\Gamma_l(T) = Q^r(T) - \epsilon m_Q. \tag{6}$$

To first derive, an instantaneous stateless residual $\nabla(T)$ which is the 'signed residual distance' between the observed ratio and the stateless residuals as:

$$\nabla(T) = \begin{cases} Q^r(T) - \Gamma_h(T), & \text{if } Q^r(T) > \Gamma_h(T); \\ Q^r(T) - \Gamma_l(T), & \text{if } Q^r(T) < \Gamma_l(T); \\ 0, & \text{otherwise.} \end{cases} \tag{7}$$

The value of $\nabla(T)$ could be positive (or negative) depending on whether the ratio sample observed is above (or below) the upper (or lower) safe margin $\Gamma_h(T)$ (or $\Gamma_l(T)$). Thus, $\nabla(T)$ is zero when the ratio observed is within $[\Gamma_h(T)\ ,\ \Gamma_l(T)]$.

*2) **Stateful Residuals:*** Our framework now maintains the sum of residuals between the ratio value and the $\Gamma_h(T)$ and $\Gamma_l(T)$ over a sliding frame of past K time windows. We denote this sum as $RUC(T)$. To calculate this metric. Now, the framework calculates $RUC(T)$ over a sliding frame of past K time windows as:

$$RUC(T) = \sum_{j=T-K}^{T} \nabla(j). \tag{8}$$

D. Optimizing Standard Limits of $RUC(T)$

We need to calculate an upper and a lower threshold from the RUC values that prevent underfitting and overfitting and improves detection performance in the test set. The procedure for calculating the upper and lower thresholds is similar. Algorithm 1, shows the method for τ_{max}.

For this, we define a cost C, and penalty P, as the loss functions. The cost and penalty function represents the loss due to missed detection and false alarms respectively. One key consideration in time series attack detection is to minimize false alarms, since the actual probability of being under attack is much lesser. Therefore, seemingly low false alarm rates, do not necessarily indicate a good usable attack detector. Therefore, we need to give more importance to the false alarms. Hence, the loss due to false alarm (penalty P) gets more weight, compared to the loss due to missed detection (cost C) as is evident in Algorithm 1. In the end, we choose a threshold τ_{max} (and τ_{min}) which minimizes the absolute difference between total cost and penalty values for the positive RUC samples (and negative RUC samples).

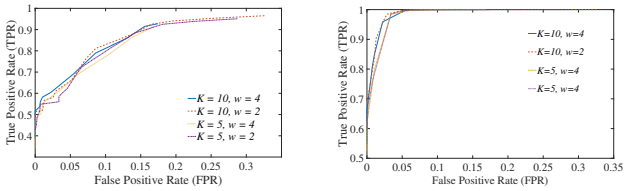

(a) ROC(CV) Additive Attack. (b) ROC(CV) Deductive Attack.

Figure 5: Parameter Selection from Cross Validation.

Algorithm 1 Calculate τ_{max}

Input: list of τ: $[\tau]$
Result: τ_{max}
for T, $[\tau]$ **do**
  if $RUC(T) > 0$ **then**
    if $RUC(T) < \tau$ **then**
      $C_{max} : \frac{|\tau - RUC(T)|}{w}$
    else
      $P_{max} : w|RUC(T) - \tau|$
    end
  end
end
$\tau_{max} = \text{argmin}_\tau\, |sum(C_{max}) - sum(P_{max})|$

The frame size K and weight w of C and P can be determined optimally, by using a small cross validation set with a few attack samples and test what values of K and w are best. We plot the Receiver Operating Characteristic (ROC) curve for the cross validation set (See Figs. 5a and 5b) for various values of K and w, and choose that combination that gives the steepest ROC curve.

E. Detection Criterion in Test Set

The main idea behind attack detection is that RUC in the test set ($RUC(T^C)$) should not deviate from the standard limit obtained from the training set. We first calculate the stateless residuals for each time window of the testing set T^C such that $\Gamma_h(T^C) = Q^r(T^h) + \kappa_{opt}$ and $\Gamma_l(T^C) = Q^r(T^h) - \kappa_{opt}$, where κ_{opt} is the margin that resulted in the optimal standard limit. The historical value of the ratio on that time window $Q^r(T^h)$, where T^c is the current time window and T^h is the corresponding time window in the training set, $\Gamma_{high}(T^c)$ and $\Gamma_{low}(T^c)$ are the safe margins at T^c of the test set.

From $\Gamma_h(T^C)$ and $\Gamma_l(T^C)$, we calculate the $RUC(T^C)$ using Eqn. 9. Then we check whether $RUC(T^C)$ violates the standard limit range identified during training set.

$$RUC(T^c) : \begin{cases} \in [\tau_{min}, \tau_{max}], \text{No Anomaly}; \\ \notin [\tau_{min}, \tau_{max}], \text{Anomaly}. \end{cases} \tag{9}$$

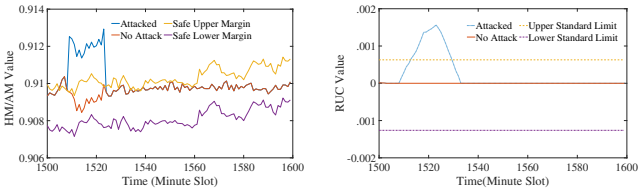

(a) Additive Attack: Tier 1. (b) Additive Attack: Tier 2.

Figure 6: Anomaly Detection for Additive Attack.

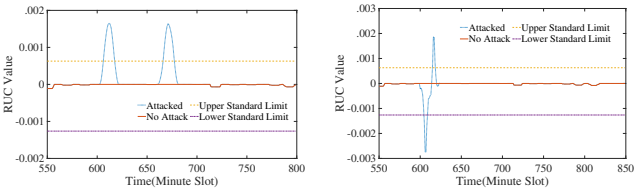

(a) Additive Ramp Attack: Tier 2. (b) Mirroring Attack: Tier 2.

Figure 7: Anomaly Detection for Ramp and Mirroring Attack.

V. Experimental Results

Using the LBNL PMU dataset (see Sec. III), we conducted extensive experiments for different falsification margins and attack strategies. For our experimental results, the first seven days are the training set and the next two days of data is used for cross validation and the remaining data are testing set. We divide this section into two parts: (1) Snapshot Results that show how our method works under several attack strategies and types (2) Performance Evaluation that shows the sensitivity versus the false alarm across varying attack margins.

Snapshot Results: We randomly selected a period from the test set and introduced an attack on the current magnitude from A6 PMU with δ_{avg} 1 p.u (which is approximately 0.16 amps). Tier 1 detection scheme to infer the presence of an attack is shown in Fig. 6a and subsequently tier 2 is applied to confirm the presence of the attack as shown in Fig. 6b. The detection of the ramp and mirroring attacks are shown in Fig. 7a and Fig. 7b. For both of these types, we have randomly selected a period of 15 minutes and introduced the respective attacks.

Performance Evaluation: For performance evaluation, we generate the ROC curve that characterizes the trade-off between the probability of attack detection vs. the probability of false alarm. we vary the δ_{avg} from 1 p.u to 2.5 p.u ($\approx$ 0.4 amps) using a step strategy to show the ROC for the additive attacks in Fig. 8a. A comparative analysis of ROCs of additive, deductive, and alternating attacks for an attack margin of 1 p.u. is shown in Fig. 8b. A report on accuracy (A), false positive (FP), and false negative (FN) for different attack margins in case of deductive and alternating attacks is given in Table I.

Table I: Experimental Results.

Attack On	Attack Type	Margin(p.u.)	A(%)	FP(%)	FN(%)
Curr. Mag.	Deductive	1	99	1	1
Curr. Mag.	Deductive	1.5	100	1	0
Curr. Mag.	Alternating	1	91	1	9
Curr. Mag.	Alternating	1.5	99	1	1

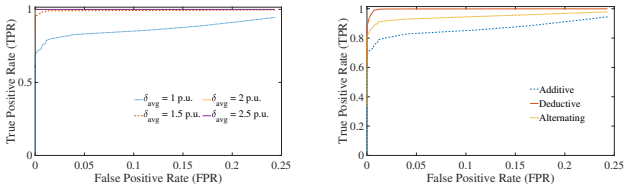

(a) ROC for Additive Attack. (b) ROCs for $\delta_{avg} = 1$ p.u.

Figure 8: Performance Analysis using ROCs.

VI. Conclusions

In this work, we presented a real time anomaly based attack detection for current magnitude falsification in PMU data streams. We showed that harmonic to arithmetic mean ratios can be used an effective invariant that is stable without attacks but show changes during attacks. We showed that even if the attacker has knowledge about the underlying time series we are still able to identify anomaly with a low false alarm rate in real time. Also, unlike many existing bad data detection methodologies, it does not require the topology of the grid network. In future, we will extend the idea to capturing voltage and phase angle falsification, and validate effectiveness against bigger PMU networks.

Acknowledgment

This work is supported by NSF grants SaTC-2030611, OAC-2017289, CNS-1818942, and SaTC-2030624.

References

[1] S. Bhattacharjee, S. K. Das, "Detection and Forensics against Stealthy Data Falsification in Smart Metering Infrastructure," *to appear, 2020.*

[2] L.M. Branscomb, R.D. Klausner, "Making the nation safer: the role of science and technology in countering terrorism" *National Research Council*, 2002.

[3] Defense Use Case, "Analysis of the cyber attack on the Ukrainian power grid," *Electricity Information Sharing and Analysis Center (E-ISAC), vol. 388*, 2016.

[4] T.M. Chen, "Stuxnet, the real start of cyber warfare?[Editor's Note]" *IEEE Network, vol. 24, no. 6*, pp. 2-3, 2010.

[5] J. Jiang, X. Zhao, S. Wallace, E. Cotilla-Sanchez, and R. Bass., "Mining PMU Data Streams to Improve Electric Power System Resilience," *ACM BDCAT*, pp. 95-102, 2017.

[6] T. Morris, S. Pan, J. Lewis, J. Moorhead, B. Reaves, N. Younan, R. King,M. Freund, and V. Madani., "Cybersecurity testing of substation phasor measurement units and phasor data concentrators" *7th Annual ACM CSIIRW*, pp. 12-14, 2011.

[7] S. Pal, B. Sikdar, and J.H. Chow., "Classification and detection of PMU data manipulation attacks using transmission line parameters," *IEEE Transactions on Smart Grid, vol. 9, no. 5*, pp. 5057-5066, 2017.

[8] D.P. Shepard, T.E. Humphreys, and A.A. Fansler., "Evaluation of the vulnerability of phasor measurement units to GPS spoofing attacks," *International Journal of Critical Infrastructure Protection, vol. 5, no. 3-4*, pp. 146-153, 2012.

[9] V.K. Singh and M. Govindarasu., "Decision Tree Based Anomaly Detection for Remedial Action Scheme in Smart Grid using PMU Data," *IEEE Power & Energy Society General Meeting*, pp. 1-5, 2018.

[10] E. Stewart, A. Liao, and C. Roberts., "Open μpmu: A real world reference distribution micro-phasor measurement unit data set for research and application development," *LBNL Tech. Rep. 1006408*, 2016.

[11] NDT Team., "Synchrophasor Monitoring for Distribution Systems: Technical Foundations and Applications," *North American Synchrophasor Initiative, Tech. Rep*, 2018.

[12] D.I. Urbina, J.A. Giraldo, A.A. Cardenas, N.O. Tippenhauer, J. Valente,M. Faisal, J. Ruths, R. Candell, and H. Sandberg., "Limiting the impact of stealthy attacks on industrial control systems," *ACM CCS*, pp. 1092-1105, 2016.

[13] M. Wu and L. Xie., "Online detection of false data injection attacks to synchrophasor measurements: A data-driven approach," *HICSS*, 2017.
[14] Z. Zhang, S. Gong, A.D. Dimitrovski, and H. Li., "Time synchronization attack in smart grid: Impact and analysis," *IEEE Transactions on Smart Grid, vol. 4, no. 1*, pp. 87-98, 2013.