8-1-2018

Power System Spoof Detection with a Hybrid Hardware/Software Benchmarking Framework

Keaton Dieter
Oregon State University

Ben McCamish
Oregon State University

Eduardo Cotilla-Sanchez
Oregon State University

Robert B. Bass
Portland State University, robert.bass@pdx.edu

Scott Wallace
Washington State University

See next page for additional authors

Citation Details

K. Dieter, B. McCamish, E. Cotilla-Sanchez, R. B. Bass, S. Wallace and X. Zhao, "Power System Spoof Detection with a Hybrid Hardware/Software Benchmarking Framework," 2018 IEEE Power & Energy Society General Meeting (PESGM), Portland, OR, 2018, pp. 1-5. doi: 10.1109/PESGM.2018.8585743

Power System Spoof Detection with a Hybrid Hardware/Software Benchmarking Framework

Keaton Dieter, Ben McCamish
and Eduardo Cotilla-Sanchez
Oregon State University
Corvallis, Oregon, USA

Robert B. Bass
Portland State University
Portland, Oregon, USA

Scott Wallace
and Xinghui Zhao
Washington State University
Vancouver, Washington, USA

***Abstract*—The integration of monitoring and control networks at different voltage levels and across utility boundaries has made it harder to maintain and asses the resilience of power systems due to increasing cyber attacks. On the software side, a variety of research efforts pursue cyber protection algorithms, such as spoof detection techniques. On the hardware and firmware side, research has demonstrated the feasibility of adversarial attacks by providing an entry point at the device level. This work proposes and evaluates two detection performance metrics for a variety of cyber spoofing attacks introduced in a realistic Phasor Measurement Unit (PMU) network for a hybrid transmission and distribution power system. This research finds that both proposed metrics show promise in aiding a spoof detection algorithm in consistently detecting spoofs in power system measurements.**

I. Introduction

The proliferation of smart grid devices at the distribution level creates new attack surfaces that can be used by adversaries planning cyber intrusions. The *Advanced Metering Infrastructure* (AMI) is one place that additional attack surfaces are created, and the number of AMI devices is growing rapidly. In 2014, nearly 60 million AMI devices were present on the U.S. grid, a number that is expected to grow to nearly 1.1 billion by 2022 [1], [2].

Along with the growth in AMI adoption there has been proliferation in the number of methods used to compromise them, including compromising single meters through reverse engineering or imitating RF communications. More seriously, once a single meter is compromised the attacker could utilize the access to the network to spread *malware* to other meters on the network, a process which has been demonstrated by researchers [3], [4].

AMI has the ability to both track loads in real time and remotely connect or disconnect the loads they monitor [5]. A compromised device then, not only gives adversaries the ability to control the measurements sent to operators, but also to maliciously manipulate system loads [6]. Such load manipulations could be used to physically damage the system by oscillating large amount of load [7], [8], which could seriously affect the stability of large areas of the grid [9]. Additionally, the ability to control measurements could not only allow adversaries to cover up the measurements that show operators the load manipulations, but could allow adversaries to utilize economic attacks by adjusting usage information [3].

Therefore, it is important for system operators to be able to detect an intrusion, specifically when the data reported from the grid has been spoofed and is the result of an adversary changing real measurements to artificially-generated data. Previous work has been done detecting spoofed data in measurements from the transmission level of the power grid [10]. However, the threats posed by compromised AMI affect the distribution level. It is anticipated that threats at this level will be harder to detect due to the higher varieties of loads and smaller amounts of electrical coupling between locations of different areas of the distribution network. Our work seeks to improve the work in [10], and show that the methods in that research can extend to the distribution level. To do this, we leverage our inter-institution phasor measurement unit (PMU) network to test the performance of two spoof detection metrics. These metrics indicate when it is likely that a spoof may be occurring. We then show the performance of these metrics on both spoofed and non-spoofed versions of two events, data from normal operational data, and one from a load manipulation we created at the OSU campus.

II. Distribution level PMU network

Our research PMU network consists of seven PMUs located at the Oregon State (OSU), Portland State (PSU) and Washington State, Vancouver (WSU) University campuses. The PMUs provide monitoring at the utilization level (120/208 V), with the exception of two PMUs at OSU, which monitor a 20 kV distribution substation and auto-transformer providing 480 V service into OSU's power engineering laboratory. All PMUs monitor three phase service.

The PMU network is of particular value to research in spoof detection because two of the monitored locations have controllable loads that can be used to demonstrate the generation of real system events that can be spoofed for the purposes of testing our algorithms. The first of these locations is OSU's Wallace Energy Systems & Renewables Facility (WESRF) power laboratory, where a 750 kVA, 480 V three-phase load exists that is controlled by a PC workstation. The second location is at PSU where a similar, but smaller, laboratory connection exists. At both of these locations, either of the controllers may be used to mount a physical attack on the load. Then the attacker could mask the physical attack by using a man-in-the-middle injection point to spoof the data coming

from the PMU monitoring the location. In this paper, a proof of concept attack is staged at the OSU campus. However, it is clear from the variety of controls, machines, and loads under the control of the researchers that many more power system events could be generated in the future.

Of the seven PMUs in the network, five of the PMUsare concentrated and very close to the OSU network. The other two, at PSU and WSU allow for visibility from a larger distance away from the OSU campus.

The PMUs report data at 60 samples per second to a Phasor Data Concentrator (PDC) located at the OSU campus. The PDC stores data in CSV, SynchroWAVe, and PDAT formats.All of these file-types are archived for permanent storage so that we may reliably access them from a variety of software tools.

III. Metrics for spoof detection

This section describes the metrics that we use to aid spoof detection. First, we describe the inter-PMU correlations. We then describe two additional metrics designed to detect when a PMU ceases to follow the same trends as the others, and therefore is likely being spoofed. These metrics are then used by an MLA, along with other information, to identify when data from a PMU is being spoofed.

A. Inter-PMU correlations

Previous work determined that inter-PMU correlations were useful for detecting spoofs on transmission level PMU data [10]. Inter-PMU correlations are calculated as the Pearson correlation coefficients for one measurement type for any pair of PMUs on the network. Since it was anticipated that the inter-PMU correlations would be weaker on the distribution data, we focus on the strongest correlations identified in [10]. Namely, inter-PMU correlations in frequency (f), positive sequence voltage magnitude ($|V_1|$), and positive sequence voltage angle (ϕ_1). We tested these correlations for the time windows identified in [10] on distribution data and chose the ten second window because it provided the cleanest signals on our data. Additionally, we used the algorithm in [11] to "unwrap" the phase angles to undo rollovers in the data because these transitions affected the correlations in a way that did not represent the underlying data's continuous nature. For the rest of this paper, an inter-PMU correlation is represented by $X_{t,i,j}$ where t is the correlation type (f, $|V_1|$, ϕ_1), calculated between PMU i and PMU j. X_t denotes the set of all correlations of a given type.

B. Correlation out-of-bounds

First, we created the correlation out-of-bounds metric (OOB). In [10], a similar metric, the MCOOB is used. However, we wanted a more defined version of this that was a binary value for each instant in time, since this would translate better to being the input to an MLA. The OOB determines if the inter-PMU correlation is outside of the envelope determined by the median of the set of all of the inter-PMU correlations for a given type. The width of the envelope is defined by α which we set to $\alpha = 0.3$ for our results, determined experimentally. The purpose of the OOB is to capture when the value of one inter-PMU correlation drifts from the values seen by the rest of the PMUs. Equation 1 shows the mathematical formulation of the OOB for correlation between PMU i and PMU j of type t.

$$OOB_{t,i,j} = [(1-\alpha) \times \tilde{X}_t < X_{t,i,j} < (1+\alpha) \times \tilde{X}] \quad (1)$$

C. Correlation out-of-swing

We observed that the correlations tended to de-correlate/re-correlate at the same time, but with different rates. To account for this, we created the out-of-swing (OOS). The different rates caused the OOB to improperly capture this trend. OOS determines the direction of change for the inter-PMU correlations, or the unit direction, at time k $D_{t,i,j}(k)$, shown in Equation 2. This takes the value of -1 if decreasing, 1 if increasing, and 0 if not changing. Next, we determine if the correlations of type t from a given PMU i are changing in the same direction as all the correlations of the same type. To calculate the direction of change, for a single PMU, we sum $D_{t,i,j}$ for a single PMU, by holding i constant and sweeping all available j values, and then find the unit direction by dividing by the absolute value of the same sum. Then, for the whole group, we do the same type of operation, by summing $D_{t,l,m}$ for all available PMUs l and m, and dividing by the absolute value. We then perform a logical check to determine if the direction of change for the single PMU is different than the whole group. If true, $OOS_{t,i}$ is 1, otherwise 0, as shown in (3). We evaluate the difference between $X_{t,i,j}(k)$ and $X_{t,i,j}(k-\beta)$ instead of simply k and $k-1$, to help filter small changes from step to step that do not reflect the overall change of the data. We found in practice that $\beta = 120$, a time difference of about two seconds worked well.

$$D_{t,i,j}(k) = \frac{X_{t,i,j}(k) - X_{t,i,j}(k-\beta)}{|X_{t,i,j}(k) - X_{t,i,j}(k-\beta)|} \quad (2)$$

$$OOS_{t,i}(k) = [\frac{\sum_j D_{t,i,j}(k)}{|\sum_j D_{t,i,j}(k)|} \neq \frac{\sum_{l,m} D_{t,l,m}(k)}{|\sum_{l,m} D_{t,l,m}(k)|}] \quad (3)$$

IV. Spoofing techniques

An attacker generating a spoofed signal S at time t with the aim of avoiding detection would likely aim to satisfy a couple of criterion: *continuity* - the change in the signal value at the onset of the spoof is small; *precedent* - the signal values generated through the course of the spoof are within normal operating ranges and have distributions similar to non-spoofed data; *time-coherence* - there are small variations between similar to that of non-spoofed data; *signal-coherence* - the spoofed data values must be consistent with each other because they are normally dependent on the system state and each other; and, *geographic-coherence* - the spoofed data must remain consistent with what is being observed at nearby locations in the grid.

In this paper, three spoofing techniques are examined that meet the above desired criteria to various degrees. Each spoof

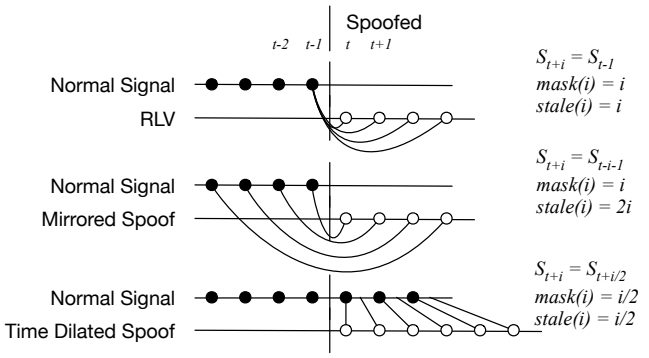

Fig. 1. Visual representation of three spoofing strategies.

is illustrated in Figure 1, where circular nodes indicate real sampled data from a PMU (solid nodes) or spoofed data (unfilled nodes). The approaches are described briefly below and illustrated in Fig 1.

Replace with Last Value (RLV): The simplest spoofing strategy is to replace new measurements with the last value obtained before the spoof began. The RLV strategy shows that no "real" PMU measurements are required after the spoof begins and is able to mask data for the entire duration it is enabled. Since the last signal value is repeated, the spoofed data gets increasingly stale (out of date) as the spoof continues.

Mirroring: We can partially address the weakness of the RLV strategy by using historic data from the moments just prior to the spoof. By playing signals back in reverse, different spoof characteristics are achieved. As with the RLV spoof, all data after the spoof starts is masked. Mirroring produces a much more natural distribution of signal samples, which have the same incremental changes as recently observed values, with the trade off that the mirroring spoof must reach further back to append data so it becomes stale twice as fast, and looses geographic coherence more quickly than RLV.

Time-dilation: Both the RLV and Mirroring strategies only require data obtained prior to the spoof onset at time t. An alternative approach is to create a spoof by resampling time and playing a true (historic) signal back more slowly. The example in Figure 1, shows time-dilation playing back data twice as slowly as normal. 2x-time-dilation masks $i/2$ samples and grows stale at a rate of $i/2$.

The time-dilation approach improves upon RLV by generating non-stationary signals that reasonably follow the anticipated distribution. It improves upon the mirroring spoof because it maintains a higher level geographic consistence since the ratio of $mask(i)/stale(i)$ is higher for time-dilation than mirroring.

V. Tests

To evaluate our spoof detection methods, two different PMU data sets were collected. The first occurs when no known events take place. The second contains a proof-of-concept generator attack performed using the generator at OSU, which is under the control of the research team. Using these events, the spoofing methods described in Section IV were used to replace all of the target PMU's reported measurements for a specified period of time, which varied based on the length needed by each spoof to cover the desired period.

A. Baseline spoofs with normal data

Baseline spoofs with normal data allow us to test the performance of the spoof detection metrics at covering up what we believe to be normal data. This provides us something to compare to the event spoofs, and lets us know that the metrics are not finding abnormalities in the data that are inherent to an event rather than spoofing. For this data set, we selected five minutes of data from our system. The selection was performed arbitrarily. However, we verified that no known events occurred during the time period. Five minutes was chosen as this was the amount of time we anticipated collecting during the event spoofs. Additionally, since we knew we would be spoofing the WESRF PMU in the event spoofs, we decided to spoof it in the baseline spoofs.

B. Event spoofs with abnormal data

The test system at OSU was used to test the spoof metrics on event data. This system is a 300 HP motor, driven by a 250 kW drive.The drive directly is coupled to a wound rotor synchronous machine that can be instantaneously connected or disconnected from a large water rheostat load.Attacks that oscillate large loads have been shown to be valid and potentially harmful [12], [8], so we made our test event resemble the beginning of this type of attack. To do this we connected the generator to the rheostat load at half of its rated output speed for one minute, then disconnected the rheostat load, causing a sudden decrease of power from the lab feeder. The spoofs were then constructed in order to completely hide the sudden drop in power.

VI. Results

This section discusses results for the two sets of data described in Section V using the various methods described in Section IV. We calculated the Pearson correlation-coefficients for all types of correlations and PMUs over a ten second window, then calculated the OOS and OOB, as described in Section III. We calculated the fraction of time the OOS and OOB were active separately for the time leading up to the beginning of the spoof, and during the spoof. We did these calculations this way to compare against various periods of time because the time dilation spoofing method required a different length of time to properly cover the event. We used two groupings for the statistical analysis of the OOS and OOB values: the first, the spoofed group (S) that contains all the PMUs that had data spoofed, in this case only WESRF; the second, the non-spoofed group (NS) that contained the rest of the PMUs that did not have data modified.

In order to evaluate the performance of the OOS and OOB metrics, it is important to note our goals. For these metrics to function well as the inputs to an MLA, they must identify only the PMU that is being spoofed, and only mark it as such during

the spoof. In order to evaluate this, we created the following conditions for both the OOB and the OOS metrics:

1) Before the spoof, the NS group must see minimal OOB and OOS fractions.
2) Before the spoof, the S group must see minimal OOB and OOS fractions.
3) During the spoof, the NS group must see very similar OOB and OOS fractions compared to before the spoof.
4) During the spoof, the S group must see significantly larger OOB and OOS fractions.

A. Baseline spoofs with normal data

The results for the OOS fractions are illustrated in Figure 2. We evaluate the results according to the four criteria. Criteria one is mostly satisfied, though to varying degrees for each type of correlation, with the $|V_1|$ begin the worse, on the upper end indicating that there may be spoofs happening when none are happening almost 30% of the time. Requirement two is also mostly satisfied, with much smaller fractions than the NS group, probably due to the fact that only one PMU is included. The third and fourth criterion can be observed in Table I. We see that the third is again mostly satisfied, as the increases here are on average only as large as 1.72. The fourth is satisfied, as we see the average increase ranging from 5 to 55 times the pre-spoof values.

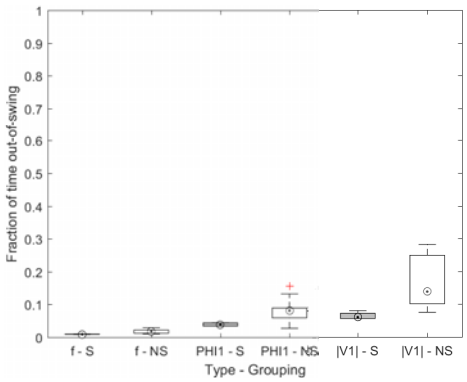

Fig. 2. Pre-spoof baseline OOS fraction distributions for the spoofed PMU (WESRF, denoted by S) and the non-spoofed group (all others, denoted by NS)

Next, we examined the OOB fractions, and evaluated each of the requirements for good performance. Requirements one and two can be observed in Figure 3. From this we see that the first and second criteria are satisfied for the f and ϕ_1 correlation types but probably not for the $|V_1|$ correlation, as we see values here upwards of 0.50, meaning that 50% of the time this will lead a MLA to falsely label data as spoofed. Table I shows the results that help us make a determination on the third and fourth requirements for good spoof detection metrics. We see that there is good performance with ϕ_1, as before the spoof the fraction was 0 and after the spoof it was large for the S group and there was no increase for the NS group. The frequency correlation also yielded acceptable performance for these requirements, though while we see approximately 6 times increase on average for the S group, we see values for increase greater than 2 for the NS group, which is not ideal. The $|V_1|$ correlation type fails requirement four as there is no increase for the S group, but passes requirement three since there is also no increase for the NS group.

TABLE I
BASELINE FACTOR OF INCREASE IN OOS/OOB FRACTIONS FROM PRE-SPOOF TO DURING-SPOOF.

Metric	Type	Spoofed PMU		Non-spoofed PMUs	
		Mean	Median	Mean	Median
OOS	f	55.39	49.83	1.72	0.67
	ϕ_1	14.02	11.14	0.71	0.59
	$\|V_1\|$	5.45	4.98	1.12	1.32
OOB	f	6.90	6.88	1.81	2.74
	ϕ_1	Large	Large	None	None
	$\|V_1\|$	1.32	1.30	1.28	1.34

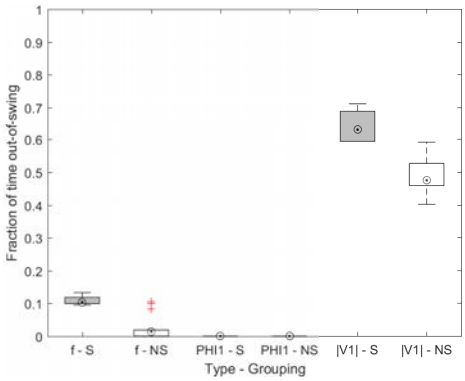

Fig. 3. Pre-spoof baseline OOB fraction distributions for both groups (S, NS)

B. Event spoofs with abnormal data

To begin with, the OOS fractions were analyzed. From the pre-spoof Figure 4 we see how well the first two requirements of a spoof detection metric hold up. Due to a low amount of pre-spoof activity for both groups both requirements are satisfied, though, the $|V_1|$, has higher levels of pre-spoof activity, so it may not be fully met here. Table II shows that requirement three is very well satisfied, with most of the average increases being very close to one for the NS group, and the fourth criteria is also satisfied, though at an average increase of only around three. Again, $|V_1|$ may not fully meet this criteria.

TABLE II
GENERATOR EVENT FACTOR OF INCREASE IN OOS/OOB FRACTIONS FROM PRE-SPOOF TO DURING-SPOOF.

Metric	Type	Spoofed PMU		Non-spoofed PMUs	
		Mean	Median	Mean	Median
OOS	f	26.43	24.14	1.03	1.22
	ϕ_1	6.32	6.04	1.29	1.43
	$\|V_1\|$	2.87	2.80	1.02	1.04
OOB	f	2.82	2.14	2.68	3.43
	ϕ_1	Large	Large	Small	Small
	$\|V_1\|$	1.18	1.16	1.17	1.17

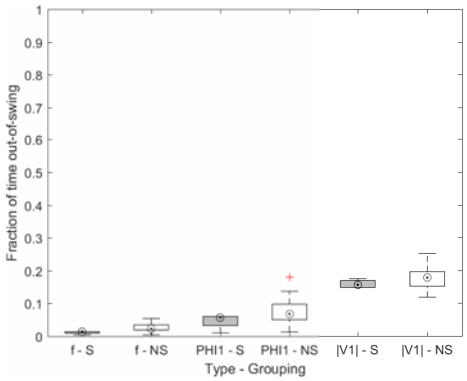

Fig. 4. Pre-spoof generator event OOS fraction distributions both groups (S, NS).

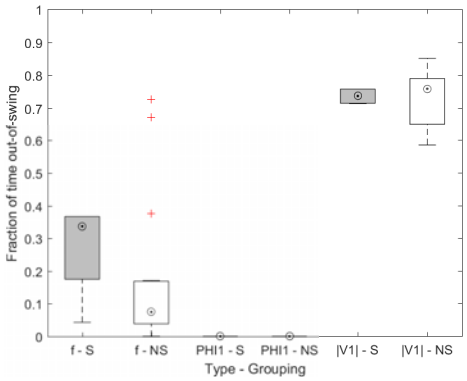

Fig. 5. Pre-spoof generator event OOB fraction distributions for both groups (S, NS)

Rounding out the analysis of the spoofs over the generator event, we analyzed the OOB fractions for three spoofing methods. We see from Figure 5 that with such high pre-spoof values, $|V_1|$ does not satisfy requirement one or two, and f only partially satisfies them, showing values nearing 0.40; ϕ_1, however definitely satisfies both requirements. Finally, from Table II, only ϕ_1 and $|V_1|$ have values at or around one for the NS, and as such are the only ones who satisfy requirement three. Additionally, $|V_1|$ shows no significant increase for the S group during the spoof, and so fails requirement four, and f only shows a moderate increase of around two, and so only marginally satisfies it; once again, however, ϕ_1 definitely satisfies requirement four.

C. Summary

From the results of both the baseline and event spoofs, a couple of things are clear. First, the voltage angle and the frequency correlations are the most useful. The voltage angle coupled with the OOB produced the most binary results (off pre-spoof, on during spoof). Additionally, the frequency correlation produced the best results when coupled with the OOS. These satisfying the requirements that we defined for good spoof metrics. Also, the OOB showed that it is really only reliably effective when paired with the voltage angle, though it proved it could be very effective then. The OOS performed better, satisfying more of the requirements for most of the correlation types.

VII. Conclusion

This work focused on evaluating the performance of two different spoof detection metrics.To test these metrics, we spoofed data using three types of spoofs: mirroring, replace-last-value, and time-dilation. We established a baseline for the data by replacing segments of normal operating data at one PMU. Next, leveraging the WESRF laboratory, it was possible to create a hardware event using a large motor in the WESRF laboratory. The data were collected from this event and spoofed in order to cover the event. This work only examines one event created this way, but future work will be able to leverage this hardware/software framework to do much more in-depth validation of software used to detect spoofing. The OOB and OOS metrics were calculated using the inter-PMU correlations in frequency, and positive sequence voltage magnitude and angle. We determined that the OOS had the most consistent performance of the two metrics, showing promising performance with all three inter-PMU correlation types. However, the OOB may be particularly useful when paired with voltage magnitude inter-PMU correlation as well. This work will be expanded upon by testing with more events, and types of events, created in the laboratory and by expanding the abilities of the metrics by incorporating them into a MLA.

References

[1] "Advanced metering infrastructure installations in the U.S.A," U.S. Energy Information Administration, Tech. Rep.
[2] "Number of smart meter installations, worldwide," Navigant Research, Tech. Rep.
[3] A. Geib, "How privacy-conscious consumers are fooling, hacking smart meters," Natural News, July 2012.
[4] P. McDaniel and S. McLaughlin, "Security and privacy challenges in the smart grid," *IEEE Security Privacy*, vol. 7, no. 3, pp. 75–77, May 2009.
[5] E. Shakshuki, K. Shuaib, Z. Trabelsi, M. Abed-Hafez, A. Gaouda, and M. Alahmad, "Resiliency of smart power meters to common security attacks," *Procedia Computer Science*, vol. 52, pp. 145 – 152, 2015.
[6] "Smart meters and smart meter systems: A metering industry perspective," Edison Electric Institute and Association of Edison Illuminating Companies, Tech. Rep., 2011.
[7] A. Natarayan, "The emerging smart grid: Opportunities for increased system reliability and potential security risks," Ph.D. dissertation, Carnegie Mellon University, 2012.
[8] C. Lassetter, E. Cotilla-Sanchez, and J. Kim, "Load oscillating smart meter attack," in *Signal and Information Processing (GlobalSIP), 2016 IEEE Global Conference on*. IEEE, 2016, pp. 821–825.
[9] K. Higgin, "Smart meter hack shuts off the lights," InformationWeek, October 2014.
[10] J. Landford, R. Meier, R. Barella, X. Zhao, E. Cotilla-Sanchez, R. Bass, and S. Wallace, "Fast sequence component analysis for attack detection in synchrophasor networks," 2017, submitted to Sustainable Energy, Grids and Networks.
[11] V. Venkatasubramanian, "Real-time strategies for unwrapping of synchrophasor phase angles," *IEEE Transactions on Power Systems*, vol. 31, no. 6, pp. 5033–5041, Nov 2016.
[12] S. Liu, S. Mashayekh, D. Kundur, T. Zourntos, and K. L. Butler-Purry, "A smart grid vulnerability analysis framework for coordinated variable structure switching attacks," in *Power and Energy Society General Meeting, 2012 IEEE*. IEEE, 2012, pp. 1–6.