Mitigating IoT-based Cyberattacks on the Smart Grid

Yasin Yilmaz[1] and Suleyman Uludag[2]

[1]Department of Electrical Engineering, University of South Florida, Tampa, FL
[2]Department of Computer Science, University of Michigan - Flint, MI

***Abstract*—The impact of cybersecurity attacks on the Smart Grid may cause cyber as well as physical damages, as clearly shown in the recent attacks on the power grid in Ukraine where consumers were left without power. A set of recent successful Distributed Denial-of-Service (DDoS) attacks on the Internet, facilitated by the proliferation of the Internet-of-Things powered botnets, shows that it is just a matter of time before the Smart Grid, as one of the most attractive critical infrastructure systems, becomes the target and likely victim of similar attacks, potentially leaving catastrophic disruption of power service to millions of people. It is in this context that we propose a scalable mitigation approach, referred to as Minimally Invasive Attack Mitigation via Detection Isolation and Localization (MIAMI-DIL), under a hierarchical data collection infrastructure. We provide a proof-of-concept by means of simulations which show the efficacy and scalability of the proposed approach.**

I. Introduction

With the enhanced automation, computing, communications and control characteristics of the Smart Grid, a crucial need becomes apparent to address the plethora of security and privacy related challenges. The essential nature of the Smart Grid cybersecurity spans availability, integrity, and confidentiality of computing, communications, and/or control devices from intentional or accidental harm and damage. The severity of cybersecurity consequences in the Smart Grid is generally exasperated due to the complexity, sheer volume of the devices and stakeholders, and highly time sensitive operational constraints.

With the proliferation of Internet-of-Things (IoT) devices, 8.4B in 2017 and expected to hit 20B by 2020 [1], across the industries in general and in the power grid in specific, a new genre of attack vectors and malicious capabilities are emerging. One prolific example is Mirai malware, facilitating a huge botnet of IoT devices and significantly improving attack capabilities of even unsophisticated malicious actors by enabling compromised IoT devices, such as surveillance cameras, DVRs, home routers, etc. Mirai botnet initiated attacks peaked to almost 1 Tbps bandwidth from an *army of mice* (IoT devices), showing the effectiveness and potential devastating damage they can cause by means of small, resource constrained, and hard-to-patch devices. What is even more dreading is the fact that Mirai botnet is highly customizable and optimized, and it is up for sale as a service in the dark Web [2]. It is just a matter of time before these capabilities are employed against critical infrastructures, like the Smart Grid. Examples of vulnerable IoT devices in the Smart Grid include smart meters, smart light bulbs, smart thermostats, connected vehicles, electric vehicles, smart street lights, smart home appliances, etc. Finally, in the hands of more sophisticated parties, such as state actors, they may become even more lethal.

Scourging effects of DoS attacks on the Smart Grid, in our expanded definition, have been manifested recently in the real-world attacks in Ukraine, both December 2015 and 2016, where the city of Ivano-Frankivsk with 100K people were cut from power for 6 hours as a result of a cyberattack and an attack (attributed to Mirai) disrupting heating distribution to two housing blocks in Lappeenranta, Finland in November 2016.

While there has been some cybersecurity related work on Smart Grid related areas, to the best of our knowledge, there is no other study in the literature to mitigate the aforementioned new genre of IoT-initiated DoS/DDoS attacks against the Smart Grid, which seems to be a prime target by many categories of adversaries at the first opportunity they get, such as nation states, curious/motivated eavesdropper, terrorists/cyberterrorists, organized crime, disgruntled employees, etc.

In this paper, we present a set of vulnerabilities leading to an expanded definition of DoS attacks in the Smart Grid, especially those initiated via IoT devices. We then provide a framework, called MIAMI-DIL (Minimally Invasive Attack Mitigation via Detection Isolation and Localization). The algorithmic underpinning of MIAMI-DIL's anomaly detection is based on a computationally efficient, online, and non-parametric approach with a distributed statistical inference methodology that scales well to high-dimensional data.

The rest of the paper is organized as follows: Related work is presented in Section II. The system model is explained in Section III. Our anomaly-based intrusion detection system formulations with analytical details are provided in Section IV. A proof-of-concept simulation results are summarized in Section V followed by concluding remarks in Section VI.

II. Related Work

Cyber attacks targeting the availability dimension, usually referred by the umbrella term of Denial of Service (DoS) attacks, are not new. They have been studied for the Internet for a while with many proposed defense mechanisms, e.g., [3]. Yet, providing efficient and effective solutions and mitigation techniques for the Internet DoS attacks are still challenging and elusive.

When it comes to the Smart Grid, the potential damage of such attacks are more profound and due to the peculiar features of the infrastructure the DoS attacks pose an even

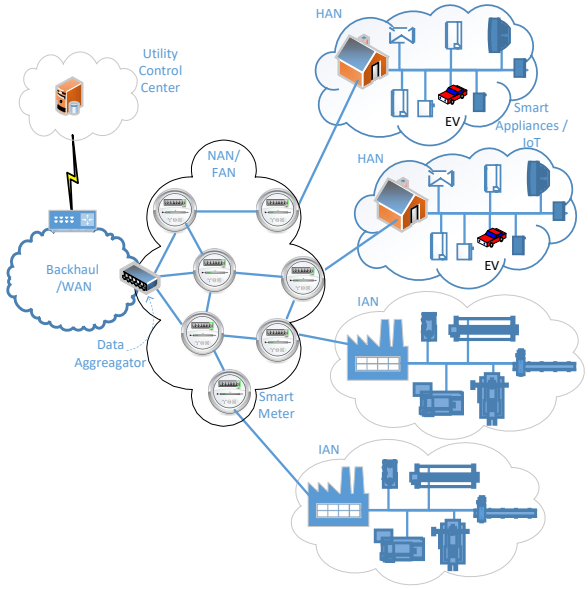

Fig. 1. The Smart Grid Model hierarchy with HAN (Home Area Network, IAN (Industrial Area Network), NAN/FAN (Neighborhood/Field AN, WAN (Wide AN), and the utility control center.

harder challenge. A specification-based intrusion detection system (IDS) is proposed in [4] tailored for the application layer protocol of ANSI C12.22 to catch violations of the specified security policy. However, there are other protocols used in industry and even the same protocol might be deployed with proprietary implementations to make such an approach infeasible for all cases [5]. Another packet-level inspection for intrusion detection is proposed in [6] on encrypted traffic, albeit with the same application layer protocol. A fourth order Markov Chain is used to model the event logs of data aggregators in [7], which can only scale to a small number of aggregators and cannot be deployed on the smart meters. A hierarchical distributed IDS for the Smart Grid is proposed in [8], designed for specific wireless mesh network technology assumptions. An anomaly-based IDS is presented in [9] that can only be deployed at headend and the data aggregators due to its computational complexity. Timely detection of false data injection attacks against voltage phase estimation in the Smart Grid is considered in [10] from a parametric point of view.

With the above literature review and the pertinent features of the Smart Grid infrastructure as discussed in Section III, a centralized intrusion detection would not work due to the heterogeneity of the constituent and independent networks [4], [11]. A fully distributed and computationally efficient technique is needed to facilitate deployment at many system devices. Finally, to make the deployment feasible to as many systems as possible, it should not be tied to a specific protocol or data type. We address all of these features in our approach as explained in the following sections.

III. System Model

We consider the Smart Grid as a hierarchical system where a set of electrical devices[1] is connected to the Smart Grid by means of a smart meter in a Home Area Network (HAN) or Industrial Area Network (IAN), as shown in Figure 1. Neighborhood Area Network (NAN) or Field Area Network (FAN) is used to refer to the logical association of these smart meters. Data aggregators collect, summarize, and report the data from HAN/IAN through the WAN to the utility's headend or the control center.

The smart meters may report a variety of different data back to the utility, from pricing to consumption data to power quality monitoring [12]. Although we consider the consumption data in Section V, the framework proposed in Section IV can be used for different data types. We do not assume a specific protocol.

We believe that DoS attacks in the Smart Grid, as being one of the most critical infrastructures, deserve a more fine-grained and broader definition of DoS attacks targeting availability to involve the following dimensions, as partially stated in [13], [14]. Firstly, DoS may refer to several problems such as denial of control, denial of electric service, and denial of pricing service. Secondly, such problems may be caused by means of compromising data integrity (e.g., misleading state estimation and situational awareness through false data injection), as well as classical DoS methods such as flooding and jamming.

IV. Anomaly-Based IDS and Attack Mitigation

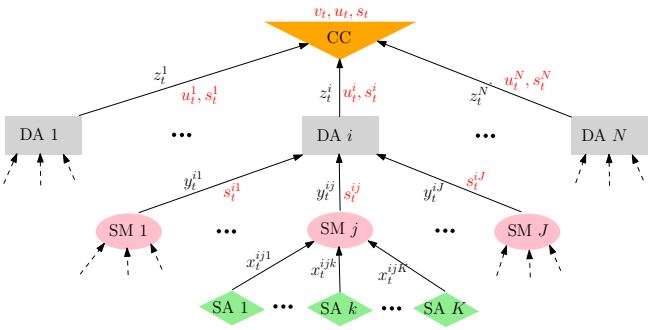

Fig. 2. Proposed hierarchical IDS. Data are shown in black, e.g., x_t^{ijk}, and statistics are shown in red, e.g., s_t^{ij}.

Considering the hierarchical topology of the the Smart Grid and security threats at each level of such hierarchy (see Sec. III) we propose a hierarchical and distributed IDS that consists of several subsystems. Specifically, each smart meter j in each NAN i monitors the streaming data $\{x_t^{ijk} : \forall k, t\}$ from each smart home appliance (i.e., IoT device) k in its HAN, and computes a statistic s_t^{ij} at each time $t = 1, 2, \ldots$, as shown in Fig. 2. Similarly, each data aggregator i monitors the streaming data $\{y_t^{ij} : \forall j, t\}$ from each smart meter j in its NAN, and computes a statistic u_t^i. Furthermore, it gathers the statistics $\{s_t^{ij} : \forall j, t\}$, from its smart meters, and combines them in s_t^i. Finally, the control center monitors the data $\{z_t^i : \forall i, t\}$ from each data aggregator i, using which it computes a statistic v_t,

[1]Smart appliances (e.g., smart bulbs, smart thermostat, electrical vehicle, and other relevant IoT devices) at home, connected machinery in an industrial setting, business equipment in commercial environment, etc.

and also combines the statistics $\{s_t^i\}$ and $\{u_t^i\}$ in s_t and u_t, respectively (see Fig. 2).

For generality, we do not specify the types of the data $\{x_t^{ijk}\}$, $\{y_t^{ij}\}$ and $\{z_t^i\}$. To handle heterogeneity of data types we only assume they are numerical data which can be normalized, e.g., to lie in $[0,1]$ using upper and lower bounds, $\boldsymbol{x}_t^{ij} = [x_t^{ij1} \cdots x_t^{ijK}] \in [0,1]^K$. Some example data types communicated in the Smart Grid are energy consumption, voltage phase, current, active and reactive power, and power factor [12]. The statistics s_t^{ij}, u_t^i and v_t are computed to detect anomalies in the data $\{x_t^{ijk}\}$, $\{y_t^{ij}\}$ and $\{z_t^i\}$, respectively. Anomalies might be caused by various types of threats aiming at a DDoS attack, such as false data injection, man-in-the-middle, spoofing, and jamming [15]. We next show how to compute the statistics s_t^{ij}, u_t^i and v_t, as well as s_t^i, s_t and u_t.

A. Online Nonparametric Anomaly Detection

Anomaly detection in the considered Smart Grid setting is quite challenging due to the following reasons:

(C1) The *attack patterns are typically unknown* since there is a wide range of vulnerabilities for attackers, especially considering the lack of security measures in the IoT devices (such as smart appliances). Hence, parametric anomaly detection-based IDSs that assume probabilistic models for anomalies, as well as conventional signature-based IDSs are not feasible in this emerging security threat.

(C2) The *problem is inherently high-dimensional* given the large number of IoT devices in a typical HAN (i.e., the dimension of $\boldsymbol{x}_t^{ij}$) and the number of smart meters in a NAN (i.e., the dimension of $\boldsymbol{y}_t^i = [y_t^{i1} \cdots y_t^{iJ}]$). Thus, computationally efficient algorithms that can scale well to high dimensionality are required.

(C3) *Timely and accurate detection is critical* given the broad societal impacts of a successful attack to the Smart Grid, and also the high demand-response time resolution in the Smart Grid (e.g., real-time pricing).

Anomaly-based IDS has the capability of detecting unknown attacks under certain conditions. It typically needs to know a statistical description of the nominal (i.e., no attack) behavior, denoted as the baseline, and classifies each outlying instance that significantly deviates from the baseline as an anomaly. This conventional interpretation of anomaly detection is also called *outlier detection*. Ideally, with the nominal probability distribution f_0 completely known, an instance x is deemed an outlier if its likelihood under the nominal distribution is smaller than a predefined threshold, i.e., $f_0(x) < \alpha$. Equivalently, x is declared an outlier if it is outside the most compact set of data points under the nominal distribution, called the minimum volume set Ω_α given by

$$\Omega_\alpha = \arg\min_{\mathcal{A}} \int_{\mathcal{A}} \mathrm{d}y \quad \text{subject to} \quad \int_{\mathcal{A}} f_0(y)\mathrm{d}y \geq 1-\alpha, \quad (1)$$

where a data point is deemed nominal in the region $\mathcal{A}$, and α is the significance level, i.e., constraint on the false alarm probability. In high-dimensional problems like the one considered in this paper, even if f_0 is known, it is very computationally expensive (if not impossible) to determine Ω_α. Hence, in the literature, there are various methods for learning minimum volume sets [16]. One of them, called Geometric Entropy Minimization (GEM), is shown to be very effective with high-dimensional datasets [17] while asymptotically achieving the performance of minimum volume set [18].

B. Online Discrepancy Test (ODIT)

Recently a GEM-based *online and nonparametric* anomaly detector, called *Online Discrepancy Test (ODIT)*, was proposed in [19] to timely detect persistent anomalies. ODIT combines the simplicity of the GEM approach with the timely and accurate detection capabilities of the Cumulative Sum (CUSUM) algorithm [20] to enable online anomaly detection in high-dimensional problems. Hence, in this paper, we use ODIT to develop an effective and efficient IDS that addresses the challenges (C1)–(C3).

We next show the ODIT procedure for smart meter j under data aggregator i, which observes the data vector $\boldsymbol{x}_t^{ij}$ at each time t. ODIT assumes a training dataset $\mathcal{X}_N = \{\boldsymbol{x}_1^{ij}, \ldots, \boldsymbol{x}_N^{ij}\}$ that is free of anomaly, and randomly separates it into two subsets $\mathcal{X}^{N_1}$ and $\mathcal{X}^{N_2}$ for computational efficiency, as in the bipartite GEM algorithm [17]. Then, for each point in $\mathcal{X}^{N_1}$ it finds the k nearest neighbors from $\mathcal{X}^{N_2}$, and forms an M-point k-nearest-neighbor (M-kNN) Euclidean graph $G = (\mathcal{X}_M^{N_1}, E)$ by selecting the M points $\mathcal{X}_M^{N_1}$ in $\mathcal{X}^{N_1}$ with the smallest total edge length and their k closest neighbors in $\mathcal{X}^{N_2}$, where $E = \{e_{m(n)}\}$ is the set of edges with $e_{m(n)}$ denoting the edge between point m in $\mathcal{X}^{N_1}$ and its nth nearest neighbor in $\mathcal{X}^{N_2}$. The total edge length of a point m in $\mathcal{X}^{N_1}$ is given by

$$L_m = \sum_{n=k-s+1}^{k} |e_{m(n)}|^\gamma, \quad (2)$$

where $|e_{m(n)}|$ is the Euclidean distance between point m and its nth nearest neighbor in $\mathcal{X}^{N_2}$, $1 \leq s \leq k$ is a fixed number introduced for convenience, and $\gamma > 0$ is the weight. It is known [17] that $\mathcal{X}_M^{N_1}$ converges to the minimum volume set Ω_α as

$$\lim_{M,N_1 \to \infty} M/N_1 \to 1-\alpha.$$

An example M-kNN graph is shown in Fig. 3, where "Training set 1" and "Training set 2" denote $\mathcal{X}^{N_1}$ and $\mathcal{X}^{N_2}$, respectively, and the edges are shown with solid lines. In this M-kNN graph, 4 out of 5 points in $\mathcal{X}^{N_1}$ are connected to their 2 nearest neighbors in $\mathcal{X}^{N_2}$ with $L_{(M)}$ showing the longest edge in the graph. Two test points and their longest edges (L_1 and L_2) are also shown.

In *outlier detection* by bipartite GEM [17], each test point $\boldsymbol{x}_t^{ij}$ is classified as an outlier if its total edge length L_t is greater than that of the Mth point, which has the largest total edge length, in $\mathcal{X}_M^{N_1}$, i.e., $L_t > L_{(M)}$. On the other hand, in ODIT,

$$D_t = L_t - L_{(M)} \quad (3)$$

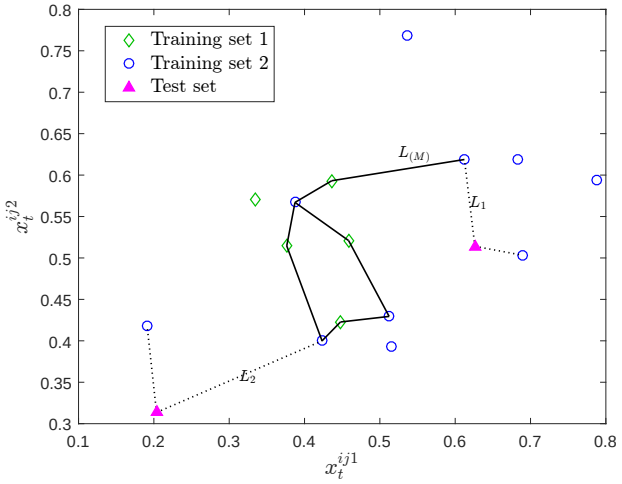

Fig. 3. ODIT procedure with $N_1 = 5$, $N_2 = 10$, $M = 4$, $k = 2$, $s = 1$, $\gamma = 1$. $L_1 - L_{(M)}$ and $L_2 - L_{(M)}$ are used as in (3) for online anomaly detection (see also Fig. 4). Test points are from the same nominal distribution as training points, which is a two-dimensional Gaussian with independent components with 0.5 mean and 0.1 standard deviation.

is treated as some positive/negative evidence for anomaly, which approximates the log-likelihood ratio $\ell_t = \log \frac{p(\boldsymbol{x}_t^{ij}|\mathsf{H}_1)}{p(\boldsymbol{x}_t^{ij}|\mathsf{H}_0)}$ between the alternative hypothesis H_1 claiming $\boldsymbol{x}_t^{ij}$ is anomalous and the null hypothesis H_0 claiming $\boldsymbol{x}_t^{ij}$ is nominal [19]. Assuming the data $\boldsymbol{x}_t^{ij}$ is independent over time, $\sum_{t=1}^{T} D_t$ gives the aggregate anomaly evidence until time T, simimlar to the running log-likelihood $\sum_{t=1}^{T} \ell_t$, which is the sufficient statistic for optimum detection. Leveraging the analogy between the anomaly evidence D_t and the log-likelihood ℓ_t (see [19] for a technical discussion) ODIT mimics the CUSUM procedure for online and nonparametric anomaly detection.

In particular, when the aggregate anomaly evidence $\sum_{t=1}^{T} D_t$ crosses a lower bound, say at time T_1, ODIT decides that there is no change and restarts the test by considering $\sum_{t=T_1+1}^{T} D_t$. The test continues until $\sum_{t=T_n+1}^{T_d} D_t$ crosses an upper bound the first time, say after n restarts. In this case, ODIT stops the test and decides for a change (e.g., an anomaly). Actually, in the ODIT procedure, the lower bound is set to zero not to waste time to decide for a no change decision because it is known that initially there is no change. Hence, it is possible to recursively update the ODIT statistic as

$$s_t^{ij} = \max\{s_{t-1}^{ij} + D_t, 0\}, \;\; s_0^{ij} = 0. \tag{4}$$

The stopping time of ODIT is given by

$$T_d = \min\{t : s_t^{ij} \geq h\}, \tag{5}$$

where $h > 0$ is a predetermined threshold.

The detection threshold h manifests a trade-off between minimizing the detection delay and minimizing the false alarm rate, as can be seen in Fig. 4. Particularly, smaller threshold facilitates early detection, but also increases the probability of false alarm. In practice, h can be chosen to satisfy a given false alarm rate.

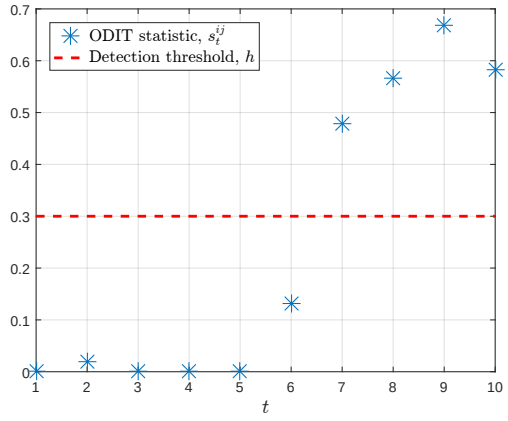

Fig. 4. ODIT statistic and decision procedure using the setup in Fig. 3 and anomalous test points from uniform distribution over $[0, 1]$. Anomaly starts at $t = 6$, and detected at $t = 7$ with the shown threshold.

C. Proposed IDS and Attack Mitigation

We presented the ODIT anomaly detector for a single smart meter in Section IV-B. Leveraging the spatial diversity that is inherent to the hierarchical structure shown in Fig. 2, we propose a system-wide IDS in which each smart meter j does not decide alone based on its data $\{\boldsymbol{x}_t^{ij}\}$, but instead cooperates with other smart meters by passing its test statistic s_t^{ij} to its parent node, data aggregator i. Gathering $\{s_t^{ij}\}$ data aggregator i fuses them into $s_t^i = \sum_{j=1}^{J} s_t^{ij}$, and passes it to the control center, together with the statistic u_t^i of data $\{y_t^{ij}\}$ it receives from its smart meters. It computes u_t^i in the same way as s_t^{ij}, as shown in (5). Note that s_t^{ij} denotes the evidence for anomaly at HAN j, and summing the independent evidences s_t^i gives the total evidence among HANs. Finally, control center, receiving the statistics $\{s_t^i, u_t^i\}$ obtains $s_t = \sum_{i=1}^{N} s_t^i$ and $u_t = \sum_{i=1}^{N} u_t^i$, and computes, through (5), the statistic v_t of data it receives from data aggregators. The statistics s_t, u_t, and v_t measure the anomaly evidence at different levels of hierarchy, namely smart appliances, smart meters, and data aggregators, respectively; hence they potentially exhibit heterogeneity. Using each of them control center runs three separate ODIT procedures

$$\begin{aligned} T_s &= \min\{t : s_t \geq h_s\}, \\ T_u &= \min\{t : u_t \geq h_u\}, \;\; T_v = \min\{t : v_t \geq h_v\}, \end{aligned} \tag{6}$$

and decides for an anomaly the first time one of them stops, i.e.,

$$T_d = \min\{T_s, T_u, T_v\}. \tag{7}$$

After detecting an anomaly in the system, which potentially corresponds to an attack, control center takes action to mitigate its effects in a minimally invasive fashion (i.e., with minimal service interruption). Specifically, it first *isolates* the information flow from suspected nodes, and then *localizes* the actual attack places via further investigation. We call this framework *MIAMI-DIL (Minimally Invasive Attack Mitigation via Detection Isolation and Localization)*, which is summarized below.

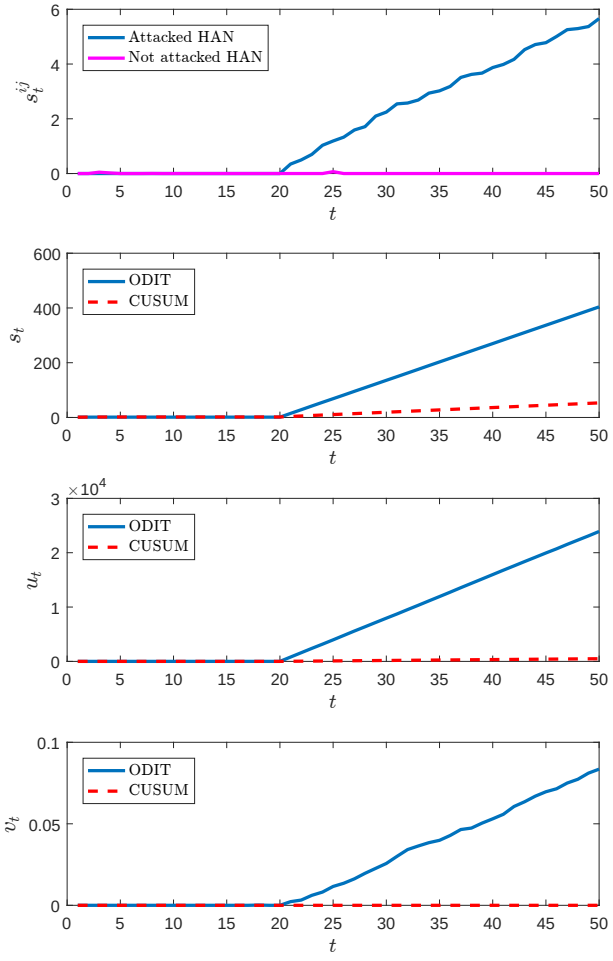

Fig. 5. ODIT and CUSUM statistics when 10% of the HANs are attacked.

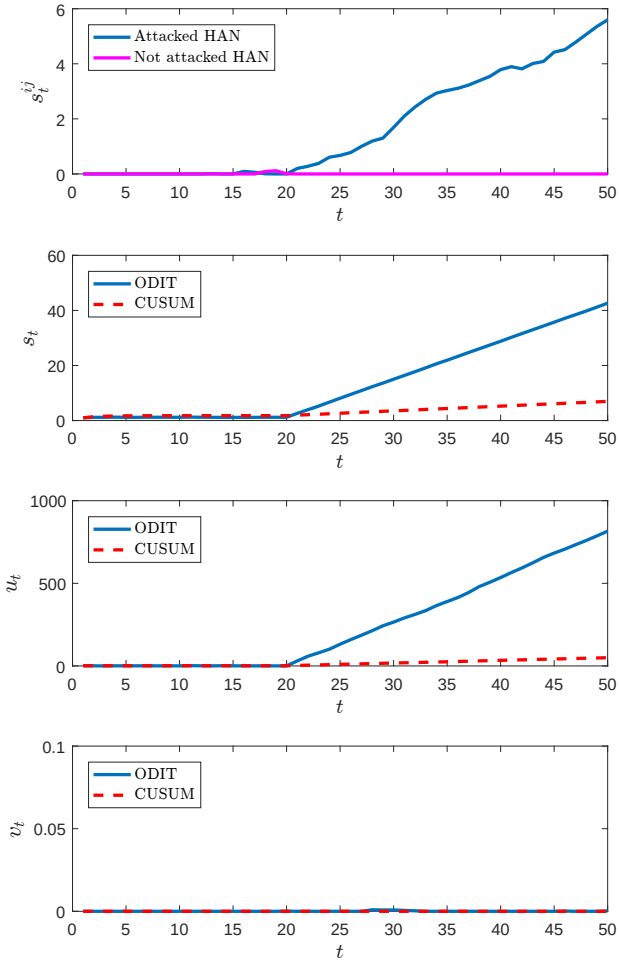

Fig. 6. ODIT and CUSUM statistics when 1% of the HANs are attacked.

Detection: Control center gathers data and statistics from the network, and sequentially detects possible anomalies in the system using the IDS given by (7).

Isolation: After detection, control center identifies the data aggregators which positively contributed to s_t and u_t. If v_t shows no signs of anomaly (i.e., no attack suspected at the data aggregators), then only those data aggregators with highly positive s_t^i and u_t^i are temporarily isolated until further investigation, that is, data from them is disregarded, and historical averages are used instead. If on the other hand v_t is highly positive, then the data flow from all data aggregators are suspended until the identification of attacked ones through further investigation.

Localization: After isolation, detailed investigation is performed for the suspected nodes to localize the actual attack places. This is done by comparing the local data from a suspected node (i.e., smart appliance, smart meter, data aggregator) with its training data that is known to be nominal. After localizing actual attack places, regular service continues with the cleared nodes; however, data from the identified attack places is ignored until they are fixed.

V. Performance Evaluation

In this section we numerically evaluate the performance of the proposed IDS. We consider a smart grid that consists of a control center, $N = 100$ data aggregators, $J = 1000$ smart meters under each data aggregator, and $K = 10$ smart appliances under each smart meter, yielding in total $1,000,000$ smart appliances and $100,000$ smart meters systemwide. Thus, the dimensionality of the system is on the order of one million.

Attack Model: Note that the security level typically decreases as we go down the hierarchy shown in Fig. 2, hence we consider a practical scenario where smart appliances (i.e., IoT devices) and smart meters could be under attack, but data aggregators are secure. To parameterize the attack size we define a as the percentage of home-area networks under attack, so $a\%$ of the $100,000$ (i.e., $1000a$) HANs are attacked. In each HAN, the smart meter is attacked with probability 0.1, and each smart appliance is attacked with probability 0.5, so on average $100a$ meters and $5000a$ appliances are attacked. We consider a false data injection attack in which the attack data is uniformly distributed in $[0, 1]$. In all simulations attack starts after $t = 20$.

Data Model: In each HAN, we assume data from appliances are independent and identically distributed (iid) with $x_t^{ijk} \sim$

$\mathcal{N}(0.5, 0.1)$. We consider the future energy consumption data, which IoT devices can predict from their scheduled usage and historical data. Smart meter can use this data for energy hedging to gain robustness to the volatility of real-time prices. In case of a successful attack, falsified data not only misleads pricing and hedging, but may also cause a demand-supply unbalance in the system, which might destabilize the grid and cause catastrophic outcomes such as a wide-area blackout. In accordance with this example, each smart meter sends the total energy consumption in the HAN, i.e., $y_t^{ij} = \sum_{k=1}^{K} x_t^{ijk}$, and similarly each data aggregator reports $z_t^i = \sum_{j=1}^{J} y_t^{ij}$.

We analyze both a large-scale and a small-scale attack in which 10% and 1% of the HANs are under attack, i.e., $a = 10$ and $a = 1$, respectively. As shown in Fig. 5 and Fig. 6, the ODIT statistic at the attacked HANs steadily increase right after the attack, whereas there is no evidence contributing to alarm at the non-attacked HANs. The advantage of cooperation among HANs is clearly observed in the second figures from top in Fig. 5 and Fig. 6, more emphasized in the large-scale attack as expected. Even with the 1% attack size, in every trial we observed that both ODIT and CUSUM successfully detect the attack (with no false alarm), albeit with different delays. Compared to CUSUM, the ODIT statistics rise much more quickly in both attack scenarios, which brings about much smaller detection delay under the same false alarm constraint. This result may look counterintuitive at first since CUSUM is optimum in terms of minimizing the minimax expected detection delay. However, due to its parametric nature, even a small mismatch between the actual and assumed parameter values degrade the performance, especially when it is not easy to differentiate between the nominal and anomalous distributions, as in our case. It is seen that the dominant statistic is u_t, which measures the smart meter data, since both the attacks to appliances and meters are available to it. On the other hand, the attacks do not appear much in v_t (not at all in Fig. 6) since data aggregators are not attacked, and the attacks in the lower levels are suppressed in the data aggregator data z_t^i due to summing the 1000-dimensional meter data.

VI. Conclusion

With the proliferation of IoT devices to the Smart Grid, and the ease of triggering DoS attacks even from unsophisticated malicious parties, there is an increasing need for developing solutions. In this context, we have presented a general threat model for the data collection subsystems of the Smart Grid. We have then developed a novel intrusion detection framework, called Minimally Invasive Attack Mitigation via Detection Isolation and Localization (MIAMI-DIL) that employs a scalable, online, and nonparametric intrusion detection system (IDS). Another important distinguishing feature of MIAMI-DIL is that it is protocol-agnostic and free from any data type assumptions. We have numerically shown that in a challenging high-dimensional scenario, the proposed IDS is capable of timely and accurately detecting cyber-attacks, even more quickly than the optimally designed CUSUM detector thanks to its nonparametric nature and CUSUM's performance degradation even with small mismatches between the actual and assumed parameter values.

References

[1] Rob van der Meulen, "Gartner Press Release February 2017," feb 2017. [Online]. Available: http://www.gartner.com/newsroom/id/3598917

[2] James Scott and Drew Spaniel, "Rise of the Machines: The Dyn Attack Was Just a Practice Run," Institute for Critical Infrastructure Technology (ICIT), p. 60, dec 2016.

[3] S. T. Zargar, J. Joshi, and D. Tipper, "A Survey of Defense Mechanisms Against Distributed Denial of Service (DDoS) Flooding Attacks," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 4, pp. 2046–2069, 2013. [Online]. Available: http://ieeexplore.ieee.org/document/6489876/

[4] R. Berthier and W. H. Sanders, "Specification-Based Intrusion Detection for Advanced Metering Infrastructures," in *2011 IEEE 17th Pacific Rim International Symposium on Dependable Computing*. IEEE, dec 2011, pp. 184–193.

[5] M. Q. Ali and E. Al-shaer, "Configuration-based IDS for Advanced Metering Infrastructure," *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pp. 451–462, 2013.

[6] R. Berthier, D. I. Urbina, A. A. Cardenas, M. Guerrero, U. Herberg, J. G. Jetcheva, D. Mashima, J. H. Huh, and R. B. Bobba, "On the practicality of detecting anomalies with encrypted traffic in AMI," in *2014 IEEE International Conference on Smart Grid Communications (SmartGridComm)*. IEEE, nov 2014, pp. 890–895.

[7] M. Q. Ali and E. Al-Shaer, "Randomization-Based Intrusion Detection System for Advanced Metering Infrastructure," *ACM Transactions on Information and System Security*, vol. 18, no. 2, pp. 7:1—-7:30, dec 2015.

[8] Y. Zhang, L. Wang, W. Sun, R. C. G. Ii, and M. Alam, "Distributed Intrusion Detection System in a Multi-Layer Network Architecture of Smart Grids," *IEEE Transactions on Smart Grid*, vol. 2, no. 4, pp. 796–808, dec 2011.

[9] F. A. A. Alseiari and Z. Aung, "Real-time anomaly-based distributed intrusion detection systems for advanced Metering Infrastructure utilizing stream data mining," in *2015 International Conference on Smart Grid and Clean Energy Technologies (ICSGCE)*. IEEE, oct 2015, pp. 148–153.

[10] S. Li, Y. Yilmaz, and X. Wang, "Quickest detection of false data injection attack in wide-area smart grids," *IEEE Transactions on Smart Grid*, vol. 6, no. 6, pp. 2725–2735, 2015.

[11] V. Gulisano, M. Almgren, and M. Papatriantafilou, *METIS: A Two-Tier Intrusion Detection System for Advanced Metering Infrastructures*. Cham: Springer International Publishing, 2015, pp. 51–68.

[12] R. R. Mohassel, A. Fung, F. Mohammadi, and K. Raahemifar, "A survey on advanced metering infrastructure," *International Journal of Electrical Power & Energy Systems*, vol. 63, pp. 473–484, 2014.

[13] Y. Yan, Y. Qian, H. Sharif, and D. Tipper, "A Survey on Cyber Security for Smart Grid Communications," *Communications Surveys Tutorials, IEEE*, vol. 14, no. 4, pp. 998–1010, 2012.

[14] W. Wang and Z. Lu, "Cyber security in the Smart Grid: Survey and challenges," *Computer Networks*, vol. 57, no. 5, pp. 1344–1371, 2013.

[15] Y. Lopes, N. C. Fernandes, T. B. de Castro, V. dos Santos Farias, J. D. Noce, J. P. Marques, and D. C. Muchaluat-Saade, "Vulnerabilities and threats in smart grid communication networks," *Security Solutions and Applied Cryptography in Smart Grid Communications*, p. 1, 2016.

[16] C. Scott and R. Nowak, "Learning minimum volume sets," *Journal of Machine Learning Research*, vol. 7, pp. 665–704, 2006.

[17] K. Srichanran and A. O. Hero, "Efficient anomaly detection using bipartite k-nn graphs," in *Advances in Neural Information Processing Systems (NIPS)*, 2011, pp. 478–486.

[18] A. O. Hero, "Geometric entropy minimization (gem) for anomaly detection and localization," in *Advances in Neural Information Processing Systems (NIPS)*, 2006, pp. 585–592.

[19] Y. Yilmaz, "Online nonparametric anomaly detection based on geometric entropy minimization," in *IEEE International Symposium on Information Theory (ISIT)*, 2017.

[20] M. Basseville and I. V. Nikiforov, *Detection of abrupt changes : theory and application*. Prentice Hall, 1993.