

Timely Detection and Mitigation of Stealthy DDoS Attacks via IoT Networks

Keval Doshi, Yasin Yilmaz, and Suleyman Uludag

Abstract—Internet of Things (IoT) networks consist of sensors, actuators, mobile and wearable devices that can connect to the Internet. With billions of such devices already in the market which have significant vulnerabilities, there is a dangerous threat to the Internet services and also some cyber-physical systems that are also connected to the Internet. Specifically, due to their existing vulnerabilities IoT devices are susceptible to being compromised and being part of a new type of stealthy Distributed Denial of Service (DDoS) attack, called Mongolian DDoS, which is characterized by its widely distributed nature and small attack size from each source. This study proposes a novel anomaly-based Intrusion Detection System (IDS) that is capable of timely detecting and mitigating this emerging type of DDoS attacks. The proposed IDS's capability of detecting and mitigating stealthy DDoS attacks with even very low attack size per source is demonstrated through numerical and testbed experiments.

Index Terms—IoT networks, DDoS attacks, anomaly detection, sequential detection, nonparametric methods

I. INTRODUCTION

The emergence of Internet of Things (IoT) has been one of the most significant technological advances of the last decade [1]. With the development of various miniaturized embedded systems and many web services along with cloud computing, it is virtually possible to make any isolated system to communicate with another machine. Moreover, the increased capabilities of new System on Chip (SoC) devices, and a drastic reduction in their sizes have led to an exponential increase in the number of devices that communicate through the Internet. With the number of IoT devices in use today already in a few billions, and with an exponential growth, the amount of data generated and transmitted is also witnessing a proportional increase. This has made the IoT paradigm a prime target for a legion of attackers, hackers, cybercriminals and occasionally governments [1].

Unfortunately, the security of IoT devices is not able to keep up with the hardware development and now more and more vulnerabilities are detected on a regular basis leading to security threats and privacy concerns [1]. For example, such compromised devices can be utilized to perform Distributed Denial of Service (DDoS) attacks. DDoS is a type of cyber-attack in which the perpetrator attacks an online service typically by flooding traffic using a large number of sources. Volumetric attacks as the name suggests are characterized by

enormous amount of traffic, and they normally do not require a large amount of traffic to be generated by the hackers themselves causing it to be the simplest and most common type of DDoS attack [2]. In this paper, we consider this type of DDoS attacks, especially the stealthy ones, which are challenging to detect and mitigate due to their widely-distributed nature and low-rate anomalous traffic from each source which can easily bypass traditional filters (i.e., stealth attacks). In stealthy DDoS attacks, such as the recent Mongolian DDoS attacks [3], although the increase in traffic from each source is small, collectively they are still capable of achieving their goal of disrupting the targeted service due to being widely distributed. Although a number of practical solutions have been deployed against DDoS, many problems still exist [4], especially due to the new genre of DDoS attacks through IoT devices.

A. DDoS via IoT

There has been a sharp increase in the number of IoT devices with an estimated number of 8.4B devices in 2017 which is expected to reach 20B by 2020 [5]. According to a study by Gartner, a high percentage of new businesses and systems will include an IoT component by 2020 [6]. The convenience provided by IoT technologies has led to a wide-scale deployment of a variety of Internet-connected sensors such as thermostats, security cameras, smart lights among many others. Unfortunately, the rapid spread of IoT also brings about a proliferation of security risks. Even though IoT is evolving at an expeditious pace, it is still very much in its inception stage. Hence, at this stage there is a significant risk that hacked IoT devices can be used for nefarious purposes such as being used as a part of a botnet to launch DDoS attacks [4].

Currently, IoT network security faces four major challenges:

- (C1) *Minimally invasive mitigation*: Because of the distributed nature of recent DDoS attacks, it is very difficult to detect the attacking devices. However, the attack should be mitigated with minimal interruption of services to benign users who want to legitimately use the services under attack.
- (C2) *High dimensionality*: Considering the large number of devices in a typical IoT network, and the abundant data generated by those devices, computationally efficient solutions that can achieve effective network monitoring, i.e., joint monitoring of devices, are required.
- (C3) *Unknown attack patterns*: Since there is a wide range of vulnerabilities that attackers can exploit, and new attack techniques are continuously developed by attackers, the predictability of attack patterns is quite low compared

This work was supported in part by the U.S. National Science Foundation under the Grant CNS-1737598 and in part by the SCEE-17-03 Grant.

K. Doshi and Y. Yilmaz are with the Electrical Engineering Department, University of South Florida, Tampa, FL USA (e-mail: keval-doshi@mail.usf.edu, yasiny@usf.edu).

S. Uludag, is with the Department of Computer Science, University of Michigan - Flint, MI USA (e-mail: uludag@umich.edu).

to the traditional Internet security. Hence, conventional signature-based detection techniques, as well as parametric probabilistic models are not feasible.

- (C4) *Timely detection and mitigation*: Due to the highly interconnected IoT ecosystem including the Internet and critical infrastructure such as Smart Grid, and the potential disastrous effects of cyberattacks, timely detection and mitigation of attacks is crucial.

We state some of the real-world examples to scrutinize the damage that can be caused by cyberattacks via IoT.

- 1) The Mirai botnet, that was launched in 2016, caused one of the most prolific series of attacks in the DDoS history [7], [8]. This particular botnet infected numerous IoT devices (primarily older routers and IP cameras), and reached data rates higher than 600Gbps. Through flooding the DNS provider Dyn, the Mirai botnet took down many popular websites such as Etsy, GitHub, Netflix, Shopify, SoundCloud, Spotify, and Twitter. Mirai took advantage of devices running out-of-date firmware, and relied on the fact that most users do not change the default usernames/passwords on their devices. The fact that its source code has been released on a hacker forum (and now it is available online [9]) facilitated its derivations.
- 2) In November 2016, hackers shut down the heating of two buildings in the city of Lappeenranta, Finland. This was another DDoS attack, and in this case, the attack was specifically targeted towards an attribute of smart home. The attackers managed to cause the heating controllers to continually reboot the system in a loop so that the heaters never worked. The attack was significant because the temperatures are well below freezing at that time of the year, and such a scenario can be life threatening.
- 3) In early 2017, Verizon Wireless released a report that included an unnamed university that experienced an attack from more than 5,000 IoT devices, such as vending machines and smart light bulbs.

These examples portray the severity of the situation if such botnets are acquired by sophisticated hackers and employed against critical infrastructures such as Nuclear Plants, Smart Grids, etc.

B. Contributions

In this paper, we propose a practical anomaly-based detection and mitigation technique for IoT-based DDoS attacks, especially the challenging *stealthy DDoS attacks with data rate increase per device as low as 10%*, which is significantly lower than the considered rates in the literature, and can easily bypass most of the existing approaches. Specifically, the proposed technique is based on a statistical anomaly detection algorithm called Online Discrepancy Test (ODIT) that mitigates the attack with minimal interruption of regular service; scales well to large systems; does not rely on presumed baseline and attack patterns; and achieves quick and accurate detection and mitigation thanks to its sequential nature. The major contributions of this paper are as follows:

- A novel detection and mitigation technique for stealthy DDoS attacks is proposed, and its time and space complexity is analyzed;
- Asymptotic optimality of the proposed detector is proven in the minimax sense as the training data size grows;
- Solution to a dynamic scenario in which the number of devices in the network changes is provided;
- A comprehensive performance evaluation is provided using a testbed implementation, the N-BaIoT dataset, and simulations.

We first present the problem formulation in Sec. III, then provide the proposed IDS in Sec. IV, experimental and testbed results in Sec. V, discuss the limitations in Sec. VI, and finally conclude the paper in Sec. VII.

II. RELATED WORKS

DDoS attacks via IoT networks are relatively less addressed compared to other security issues in the IoT environment. However, it is recently attracting considerable interest, e.g., [4], [7], [8], [10]. In [11]–[14], a wide range of vulnerabilities because of which conventional signature-based detection techniques fail, are discussed. In [15], authors propose a solution to UDP flood attack in an IoT environment using 6LoWPAN and IEEE 802.15.4. However, it has high overheads and complex architectures and components which do not suit an IoT environment [16]. An agent-based DDoS mitigation approach is proposed in [17]. The authors propose a two-part algorithm in which the attack detection part has been performed in the border router. [18] proposes a hardware based solution for mitigating volumetric DDoS attacks by using programmable switches. A similar entropy-based solution is proposed in [19], but with the requirement that the packet contents should be detectable. They also do not consider scenarios in which the entropy does not change but the number of packets do. Xiang et al. [20] proposes an information metric approach to quantify the differences between legitimate and attack network traffic by assuming that the legitimate traffic follows a Gaussian distribution, whereas the attack traffic follows a Poisson distribution. Recently, a new category of DDoS attacks known as Link flooding attacks (LFA) [21], which can degrade or even cut off network connectivity, have gained attention. In [22], Xue et al. proposed a system known as LinkScope, that employs both the end-to-end and the hop-by-hop network techniques for detecting LFA attacks. However, they are unable to mitigate the attack by tracing the botnets. In [23], Wang et al. propose a scheme called Woodpecker, which makes LFA more difficult to take effect. They employ a proactive probe approach which is capable of quickly localizing the congested links. Machine learning algorithms are also gaining attention as recent anomaly detection research shows promise [24]. Doshi et al. [25] presents the performance of popular machine learning algorithms such as SVM, k-nearest-neighbors, neural networks etc. in detecting malicious traffic. However, they require training data for malicious traffic (supervised anomaly detection), and extract features which are specific to certain IoT devices without considering other devices that might be present in the network such as laptops or smartphones. In [26], Nomm et al. proposes using feature selection with popular anomaly detection techniques such as one-class SVM to detect IoT botnet attacks. They use three different approaches for

extracting useful features and provide their results over the N-BaIoT dataset. Kurt et al. [27] proposes an online anomaly detection algorithm based on dimensionality reduction, that is capable of detecting anomalies in high dimensional settings. They also present their results using the N-BaIoT dataset. Meidan et al. [28] proposes using deep autoencoders for detecting DDoS attacks at the network level. They achieve small false positive rates by training a deep autoencoder for each individual device in the network, which might not scale well to large networks with many devices. They also use a window-based majority voting scheme to detect attacks, which is not well suited for quick detection. More recently, [29] proposes an ensemble of autoencoders for online intrusion detection, which can even run on an embedded system like a raspberry pi and still detect in real-time.

III. PROBLEM FORMULATION AND BACKGROUND

A. System Model

In the considered architecture (Fig. 1), each IoT device sends its data to the node connected to it. Nodes direct the data traffic to a center, such as a web server, data center or utility center. The architecture is scalable in such a way that a node may represent a smart home consisting of tens of devices or a smart building/neighborhood access point consisting of thousands of devices.

Each device typically has different data communication characteristics. In particular, the data content is typically different (for example, a thermostat would have considerably smaller packet size as compared to a CCTV camera), and the communication protocol used might be different (such as TCP, UDP or HTTP). Also, for the same device the data rates might differ significantly based on the location or type of connection, e.g., a laptop connected via a fiber optic cable might send 1000 packets per second whereas the same device would be sending 10 packets per second on a slower connection. Even the active communication frequency varies considerably, e.g., devices like printers update its status once a minute, whereas CCTV cameras send data every second. In this work, we only assume that they perform a packet-based data communication, and the probability distribution of a device's packet communication is the same in training and testing, which is a common assumption in machine learning. Note that the probability distributions are allowed to be quite complex and different from each other.

B. Threat Model

We consider a volumetric DDoS attack scenario in which data rates (packet/sec.) from a number of devices increase at some point in time (see Fig. 1). Particularly, we consider a threat model in which some IoT devices are compromised and start to send more than usual number of data packets. We do not assume further attack specifications such as knowledge on how devices are compromised (e.g., through a vulnerability in the firmware, spoofing attack, man-in-the-middle attack, use of default password), the attack magnitude (i.e., percentage of increase) and duration, and whether the data content changes




Fig. 1. System model consisting of IoT devices such as thermostat, CCTV, light bulb, smartphone, etc. In the threat model, the bold arrows imply an increased packet rate.

or not. It is not tractable to mitigate such attacks at the center since accurate identification of all attacking devices is not tractable due to the highly distributed nature of the attacks. Moreover, due to the low-rate nature of the attacks, it is very difficult to detect them locally at the nodes. We propose a general Intrusion Detection System (IDS) that can run locally and is capable of detecting and mitigating an attack even when it is not possible to inspect the data content, which is a prerequisite for many IDS algorithms, e.g., [19], [30].

Although the standard volumetric attacks that are studied in the literature include high increase in the data rate of a device, with the increasing number of compromised IoT devices low-rate stealthy DDoS attacks (e.g., with a 20% increase per device) started to become threatening [3]. Due to the proliferation of IoT, cyber-criminals can launch widely-distributed and highly-effective stealthy DDoS attacks, that can bypass conventional filters and IDSs. Hence, in this paper we study DDoS attacks with data increase rates as low as 20% per device. There are existing works which consider low-rate DDoS attacks (e.g., [20]), nevertheless the considered increase rates are still significantly higher than what we consider in this paper.

As a result of such widely-distributed DDoS attacks (e.g., the almost uniform distribution of attack traffic in Mirai [31]), it is not tractable to have a single global solution running at the server end, and thus we propose in the next section a local IDS that runs at each node. Such local solutions also facilitate accurate mitigation.

IV. PROPOSED ANOMALY-BASED IDS

In this section, we present our detection and mitigation strategy for the proposed IDS. As shown in Fig. 2, we detect an attack based on the cooperative test statistic, and once an attack is detected, we monitor each device individually to identify the attacking devices (see Sections IV-B and IV-C). We also analyze the computational complexity and a practical scenario in which the number of devices in the network are dynamic in nature.

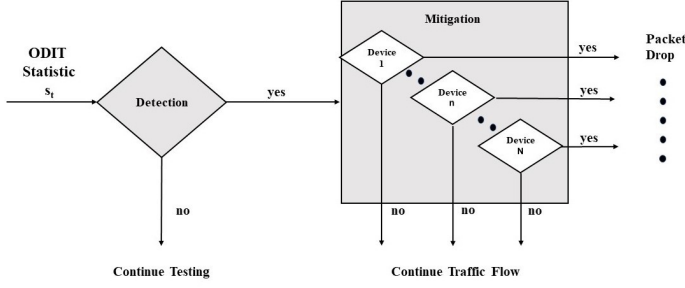


Fig. 2. Proposed Intrusion Detection and Mitigation System




Fig. 3. ODIT procedure with $M_1 = 5$, $M_2 = 10$, $\alpha = 0.2$, $k = 2$. $D_1^n = d[\log L_1^n - \log \tilde{L}_{(\alpha)}^n]$ and $D_2^n = d[\log L_2^n - \log \tilde{L}_{(\alpha)}^n]$ are used as in (1) 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, 0.5 mean, and 0.1 standard deviation.

A. Proposed Detection Strategy

The heterogeneous nature of an IoT network makes parametric anomaly detection approaches for DDoS detection less effective since they assume probabilistic models for nominal and anomalous conditions. In practice, it is difficult to know/estimate the anomalous and even the nominal probability distributions. Hence, parametric anomaly-based IDSs, as well as many conventional signature-based IDSs are not feasible in addressing stealthy DDoS attacks through IoT. Recently, an online and non-parametric detector called the Online Discrepancy Test (ODIT) was proposed for detecting persistent and abrupt anomalies [32]. Thanks to its nonparametric operation, ODIT does not need to know baseline or anomalous distributions beforehand, hence can address the challenge (C3) stated in Section I-A. ODIT is a sequential method which accumulates evidence in time, and makes a decision at each time based on the accumulated evidence so far, instead of making a hard decision based on a single data point. This sequential nature of ODIT is tailored for timely detection, thus it is able to address the challenge (C4). Moreover, ODIT can handle monitoring large number of devices together (see Sections IV-D and V), which addresses the challenge (C2).

In this work, we propose a novel modification for ODIT, and prove that this modified version, as the training data size grows, asymptotically becomes the Cumulative Sum

(CUSUM) test, which is the optimum sequential change detection algorithm in the minimax sense. CUSUM is a parametric test which assumes both the nominal and the anomalous distributions are completely known [33].

We next show the procedure for the proposed ODIT-based IDS for a node n , which observes a d -dimensional normalized data vector $\mathbf{x}_t^n \in [0, 1]^d$, where d is the number of devices, at each time t . Here, in the context of DDoS attack detection, \mathbf{x}_t^n typically denotes the number of packets received from the d devices in the network at time t and normalized by the corresponding maximum number of packets for each device. Normalization of each dimension of \mathbf{x}_t^n into $[0, 1]$ is performed to deal with the typical heterogeneity in the data communication characteristics of IoT devices. Then, in Section IV-B, we show how to achieve cooperation among nodes.

Training: Given an attack-free training dataset $\mathcal{X}_M^n = \{\tilde{\mathbf{x}}_1^n, \dots, \tilde{\mathbf{x}}_M^n\}$ which represents the baseline (i.e., nominal) operation, we begin our training procedure.

- 1) Randomly split \mathcal{X}_M^n into two subsets $\mathcal{X}_{M_1}^n$ and $\mathcal{X}_{M_2}^n$ with M_1 and M_2 points, where $M_1 + M_2 = M$, for computational efficiency, as in the bipartite GEM algorithm [34].
- 2) For each point $\tilde{\mathbf{x}}_i^n$ in $\mathcal{X}_{M_1}^n$ find the k th-nearest-neighbor (k NN) distance \tilde{L}_i^n with respect to the points in $\mathcal{X}_{M_2}^n$.
- 3) For a significance level α , e.g., 0.05, store the $(1 - \alpha)$ th percentile $\tilde{L}_{(\alpha)}^n$ of k NN distances $\{\tilde{L}_1^n, \dots, \tilde{L}_{M_1}^n\}$ to use as a baseline statistic for computing the anomaly evidence of test instances.

The training procedure is illustrated in Fig. 3, where the training set consists of $M = 15$ points, which is then randomly split into two sets of $M_1 = 5$ points (denoted by green) and $M_2 = 10$ points (denoted by blue). In this example, for each point in the first set, we find the second-nearest-neighbor ($k = 2$) distance with respect to the second set. The largest k NN distance ($\alpha = 0.2$) among the points in the first set is used as the baseline statistic $\tilde{L}_{(\alpha)}^n$.

Testing: When a new data \mathbf{x}_t^n is observed at time t ,

- 1) Compute

$$D_t^n = d[\log L_t^n - \log \tilde{L}_{(\alpha)}^n] \quad (1)$$

where L_t^n is the k NN distance of the new data point \mathbf{x}_t^n with respect to the points in $\mathcal{X}_{M_2}^n$, and $\tilde{L}_{(\alpha)}^n$ is obtained in the training.

- 2) Treating the statistic D_t^n as a positive/negative evidence for anomaly accumulate it over time as in CUSUM:

$$s_t^n = \max\{s_{t-1}^n + D_t^n, 0\}, s_0^n = 0, \quad (2)$$

where D_t^n is given in (1).

- 3) Decide to continue taking a new data point \mathbf{x}_{t+1}^n if the accumulated evidence s_t^n is not sufficient for raising an attack alarm, and stop and raise an alarm at the first time $s_t^n \geq h_n$, i.e., at time

$$T_n = \min\{t : s_t^n \geq h_n\}, \quad (3)$$

where $h_n > 0$ is a predetermined threshold.

As compared to anomaly evidence $D_t = L_t - \tilde{L}_{(\alpha)}$ presented in the original ODIT algorithm [32], we use the




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.

form given in Eq. 2. This modification enables an asymptotic optimality proof, which is presented in Theorem 1. The intuition behind this modification is to explicitly show the analogy between the attack evidence D_t and log-likelihood ratio.

Theorem 1. *When the nominal distribution $f_0(\mathbf{x}_t^n)$ is finite and continuous, and the attack distribution $f_1(\mathbf{x}_t^n)$ is a uniform distribution whose support includes \mathbf{x}_t^n , as the training set grows, the anomaly evidence D_t^n converges in probability to the log-likelihood ratio,*

$$D_t^n \xrightarrow{p} \log \frac{f_1(\mathbf{x}_t^n)}{f_0(\mathbf{x}_t^n)} \quad \text{as } M_2 \rightarrow \infty, \quad (4)$$

i.e., the proposed Odit detector converges to CUSUM, which is minimax optimum in minimizing expected detection delay while satisfying a false alarm constraint.

Proof: See the Appendix. ■

Since the proposed detector does not train on anomalous data or assume any model for anomaly, the uniform distribution condition on f_1 for asymptotic optimality is expected.

Parameter Selection: The detection threshold h_n 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_n can be chosen to satisfy a given false alarm rate. The number of neighbors k also affects the trade-off between early detection and small false alarm rate. Smaller k would result in being more sensitive to anomaly, hence supports earlier detection, but at the same time it causes to be more prone to the false alarms due to nominal outliers. Larger k would result in vice versa. The choice for M_1 and M_2 is typically skewed towards M_2 , i.e., $M_2 > M_1$, since M_2 determines the degree of resemblance between k NN distance likelihood under the nominal case, as explained in Theorem 1. The significance level α is an intermediate parameter whose effect can be compensated by the threshold h . As a rule of thumb, a small α value, such as 0.05, should be first selected, and then h_n should be set to satisfy a desired false alarm rate.

Remark 1. *A training set that is free of anomaly can be obtained through either human supervision or through an isolated secure system. We should emphasize here the difference between anomaly and outlier. The training set may contain outliers that are generated under no-attack conditions. Outliers correspond to “tail events” that can occur under nominal settings with low probability. Although outliers are rare under normal operations, they can still exist in the training set, and their natural existence does not harm the regular operation of the proposed detector. On the other hand, anomaly is a change in the system behavior, i.e., in the probability distribution of the generated data. In other words, anomaly can be defined as the existence of “persistent outliers”, as opposed to the sporadic nominal outliers.*

B. Cooperative Operation

Multiple nodes running the proposed IDS given in Eq. (1)–(3) can cooperate for earlier detection and mitigation of attacks by leveraging the hierarchical structure shown in Fig. 1. Following the cooperative CUSUM with independence assumption in [35] we propose a cooperative detector which sums the local statistics computed at the nodes to obtain a global statistic $s_t = \sum_{n=1}^N s_t^n$. That is, at each time t , each node n updates its local statistic s_t^n using (2) and transmits it to the center, which sums them to obtain the global detection statistic s_t . Then, the center decides whether or not there is an attack similarly to (3), i.e., raises an alarm at time $T = \min\{t : s_t \geq h\}$. This cooperative detector can detect attacks earlier than the single-node Odit detector thanks to the spatial diversity, i.e., accumulated attack evidences from multiple nodes. Note that the statistics from nodes without any attacked device typically take small values close to zero, but never become negative according to (2). Thus, they do not negatively contribute to the global statistic s_t , and consequently do not cause extra delay in detection.

The cooperative scheme through summing local statistics in a hierarchical architecture enables the proposed detector to easily scale to arbitrarily large networks. The optimal statistical detection would normally require multivariate analysis for all the devices in the entire network. However, due to the abundance of IoT devices this is not feasible; and more importantly due to the natural hierarchical structure in IoT networks such a large-scale multivariate analysis is unnecessary. Specifically, IoT devices are grouped under nodes such as smart home routers, and devices under different nodes can be typically modeled independently under no-attack conditions. Although an attack will normally correlate them, it is reasonable to relax that constraint since a practical attack strikes devices asynchronously, i.e., each of the attacked devices has different attack onset times. This relaxation is critical for CUSUM to be applicable to any multi-dimensional setting since it is not tractable to estimate the set of attacked devices to perform multivariate analysis [35]. After detection, the mitigation procedure given in the next subsection can be applied at each node.

C. Mitigation Strategy

Timely detection of a DDoS attack is a necessary, but not a sufficient condition for ensuring the security of a system. We need a mitigation strategy that is capable of stopping the DDoS attack by identifying the attacking IoT devices, and then blocking the traffic originating from those devices.

We perform an in-depth analysis to determine which IoT devices are causing the increase. We begin by examining the cooperative statistic s_t calculated in Sec. IV-B and determining which nodes are causing the increase. Once the attacking nodes are identified, we examine every dimension, which represent an IoT device connected to the node, of the distance L_t^n , calculated in Sec. IV-A. $(L_t^n)^2$ corresponds to the squared Euclidean norm of the d -dimensional distance vector $\mathbf{y}_t^n = \mathbf{x}_t^n - \hat{\mathbf{x}}_{(k)}^n$ whose j th entry $y_t^{n,j}$ is the distance of the data from device j at time t to the j th dimension of k th nearest neighbors, i.e., $(L_t^n)^2 = \|\mathbf{y}_t^n\|_2^2$. If $y_t^{n,j}$ comes out large, then it contributes to a large L_t^n towards an alarm, which provides an evidence that the device is under attack. Hence, after an alarm is raised at time T , we

- 1) first determine the time instance τ when the test statistic s_t started to increase since the last time it was zero ($\tau = 6$ in Fig. 4), which can be seen as an estimate for the attack onset time,
- 2) then, compute the average statistic

$$\bar{s}_n = \frac{1}{T - \tau + 1} \sum_{t=\tau}^T s_t^n \quad (5)$$

for each node n , and compare it with a threshold θ_1 to determine the attacking nodes, i.e., node n has attacked devices if $\bar{s}_n \geq \theta_1$.

- 3) then, for each node identified as attacking, compute the average distance

$$\bar{y}_{n,j} = \frac{1}{T - \tau + 1} \sum_{t=\tau}^T y_t^{n,j} \quad (6)$$

for each device j under it, and compare it with a threshold θ_2 to decide as attacking or not, i.e., device j is identified as attacking if $\bar{y}_{n,j} \geq \theta_2$.

As usual the selection of threshold θ_1 and θ_2 controls a balance between the False Positive Rate (FPR) and True Positive Rate (TPR). As shown in Fig. 15, the proposed mitigation technique achieves high TPR for almost all FPR even in the challenging attack scenario investigated in Section V. We should note that the procedure in Eq. (6) requires some memory to store the most recent distance values for all dimensions local to the node performing the procedure.

Combining the detection and mitigation strategies, the proposed IDS technique is summarized in Algorithm 1.

D. Computational Complexity

The following theorem shows that the proposed algorithm can scale well to large systems.

Theorem 2. *The online time complexity and space (i.e., memory or storage) complexity of Algorithm 1 linearly scales*

Algorithm 1: Proposed detection & mitigation algorithm

```

1: Initialize:  $s_0 = 0, t = 0$ 
2: for  $n = 1, \dots, N$  do
3:   Partition training set into  $\mathcal{X}_{M_1}^n$  and  $\mathcal{X}_{M_2}^n$ 
4:   Determine  $\tilde{L}_{(\alpha)}^n$ 
5: end for
6: while  $s_t < h$  do
7:    $t \leftarrow t + 1$ 
8:   Get new data  $\{\mathbf{x}_t^n\}$  and compute  $\{D_t^n\}$  as in (1)
9:    $s_t^n = \max\{s_t^n + D_t^n, 0\}$ 
10:   $s_t = \sum_{n=1}^N s_t^n$ 
11: end while
12: Declare attack at  $T = t$ 
13: for  $n = 1, \dots, N$  do
14:   Compute  $\bar{s}_n$  as in (5)
15:   if  $\bar{s}_n \geq \theta_1$  then
16:     for  $j = 1, \dots, d$  do
17:       Compute  $\bar{y}_{n,j}$  as in (6)
18:       if  $\bar{y}_{n,j} \geq \theta_2$  then
19:         Block traffic from device  $j$ 
20:       end if
21:     end for
22:   end if
23: end for
```

with M_2 , the number of points in the second training set, and d , the number of devices, i.e., $O(M_2 d)$. The offline training time complexity is $O(M_1 M_2 d)$.

Proof: See the Appendix. ■

Remark 2. *There are efficient ways of finding (approximate) k nearest neighbors that scale even better to high-dimensional systems. For instance the method proposed in [36] has a time complexity of $O(\lambda d)$ in online testing where λ is the maximum number of points to examine for finding k nearest neighbors. λ can be chosen much smaller than M_2 at the expense of decreasing the accuracy of k NN approximation. Hence, a balance between approximation quality and computational complexity should be sought while choosing λ . Consequently, using a fast k NN method instead of straightforward computation the real-time operation capability and/or the scalability of the proposed method can be significantly enhanced.*

E. Dynamic Environments

A major challenge for any anomaly-based intrusion detection system is adaptability to dynamic environments. This means that the system should be adaptable to changes in the environment, while still recognizing abnormal activities. In a dynamic IoT network such as a university or a shopping mall, the number of devices may frequently change based on the number of people, time of the day, day of the year etc. With varying number of devices over time the challenge for the proposed IDS is computing the k NN distance under varying number of data dimensions.

A key observation here is that data rates are specific to applications rather than devices. For instance, video streaming has




Fig. 5. Comparison of the estimated $\tilde{L}_{(\alpha)}^n$ statistic to the computed one for different combinations with increasing number of devices.

certain data rates regardless of the streaming device. Hence, considering a list of applications that are used in the network, such as web browsing, music and video streaming, we can deal with the changing number of dimensions. Specifically, we first collect training data for the extreme scenario with maximum number of devices running each application simultaneously. During online testing, at each time we modify the training set by ignoring the unused dimensions for each application, and compute L_t^n .

Since there is a huge number of possible combinations for the number of devices running each application, performing the expensive training procedure (see Theorem 2) for each such combination to compute the baseline statistic $\tilde{L}_{(\alpha)}^n$ is not feasible. Hence, we propose to build a function approximator for $\tilde{L}_{(\alpha)}^n$. We collect data for several different combinations, and compute $\tilde{L}_{(\alpha)}^n$ for each such combination. Using the number of devices running each application as input we train a regression model to estimate the $\tilde{L}_{(\alpha)}^n$ value for a given combination. The results for Gaussian process regression is shown in Fig. 5. A simpler method (e.g., linear regression) or a more sophisticated method (e.g., deep neural network) can be used for the regression model. In Fig. 5, we compare the estimated $\tilde{L}_{(\alpha)}^n$ statistic to the computed one for different combinations with increasing number of devices. We see that the estimated statistic closely matches that of exact ODIT, which is infeasible to compute for all combinations. Furthermore, the baseline statistic $\tilde{L}_{(\alpha)}^n$ depends on the selected significance level α , which is a design parameter. Our simulations show that a small mismatch between the estimated and computed $\tilde{L}_{(\alpha)}^n$ values is not critical for the algorithm's performance.

V. EXPERIMENTAL RESULTS

In this section we evaluate the performance of the proposed IDS using real data, an IoT testbed, and simulations.

A. N-BaIoT Dataset

We firstly consider the N-BaIoT dataset [28]¹, which contains data from various IoT devices under both nominal

¹This dataset is available at the UCI Machine Learning Repository.

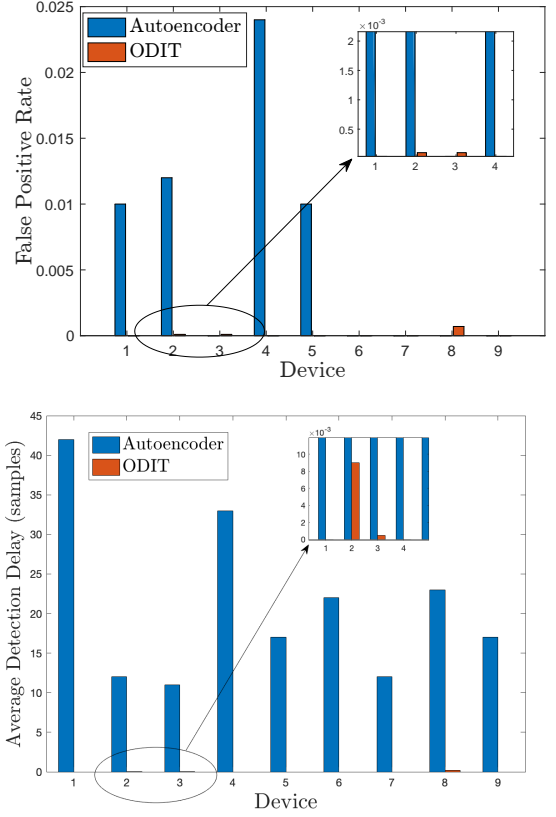


Fig. 6. Comparison between deep autoencoder-based IDS and the proposed ODIT-based IDS in terms of false positive rate (top) and average detection delay (bottom). The x-axis corresponds to the index of the attacked device.

and attack conditions. The network consists of of nine devices, namely a thermostat, a baby monitor, a webcam, two doorbells, and four security cameras connected via WiFi. The dataset description, as well as attacks considered in our experiments are presented in Table I. We do not consider UDP and TCP attacks in our experiments since the provided data does not correlate with typical UDP and TCP attacks.

Our results show that the proposed ODIT-based IDS significantly outperforms the deep autoencoder-based IDS proposed in the N-BaIoT paper [28], and by extension Isolation Forest [37], SVM [38] and LOF [39], which are shown to be outperformed by the autoencoder method. In Fig. 6, we present the false positive rate and the average detection delay when each device is under attack. In the presented experiment results, both our algorithm and the autoencoder-based algorithm successfully detect all the attacks, i.e., achieve 100% true positive rate in terms of attacks. In terms of *observed data instances*, the average detection delay is analogous to the false negative rate as all the misclassified attack data instances would contribute to the average detection delay. The proposed ODIT-based method achieves much more accurate and quicker detection than the autoencoder-based method except for device 8. On further analyzing the data, we see that there are a few outliers in the data for device 8, which causes some false alarms. Note that by increasing the decision threshold h , we are able to reduce the false positive rate at the cost of a slightly higher detection delay. Conversely, a smaller detection delay can be achieved by setting a lower threshold at the cost

TABLE I
OVERVIEW OF THE N-BAIoT DATASET PROPERTIES, BOTNET INFECTIONS, AND CONSIDERED ATTACK TYPES

Device ID	Dataset Properties		Botnet Infections		Considered Bashlite Attacks					Considered Mirai Attacks				
	Device Make and Model	Device Type	Mirai	BASHLITE	Combo	Junk	Scan	TCP	UDP	ACK	Scan	SYN	UDP	UDP Plain
1	Danmini	Doorbell	✓	✓	✓	✓	✓	-	-	✓	✓	✓	-	✓
2	Ennio	Doorbell	x	✓	✓	✓	✓	-	-	-	-	-	-	-
3	Ecobee	Thermostat	✓	✓	✓	✓	✓	-	-	✓	✓	✓	-	✓
4	Philips B120N/10	Baby Monitor	✓	✓	✓	✓	✓	-	-	✓	✓	✓	-	✓
5	Provision PT-737E	Security Camera	✓	✓	✓	✓	✓	-	-	✓	✓	✓	-	✓
6	Provision PT-838	Security Camera	✓	✓	✓	✓	✓	-	-	✓	✓	✓	-	✓
7	SimpleHome XCS7-1002-WHT	Security Camera	✓	✓	✓	✓	✓	-	-	✓	✓	✓	-	✓
8	SimpleHome XCS7-1003-WHT	Security Camera	✓	✓	✓	✓	✓	-	-	✓	✓	✓	-	✓
9	Samsung SNH 1011 N	Webcam	x	✓	✓	✓	✓	-	-	-	-	-	-	-




Fig. 7. IoT testbed consisting of NodeMCUs, smart switches, security camera, Amazon Echo Show, laptop, tablet, and Raspberry Pi.




Fig. 8. Setup for hardware implementation.

of a few more false alarms. Thanks to its sequential nature, the ODIT-based IDS detects the attacks right after it occurs while satisfying very small false alarm rates. Whereas, the autoencoder-based IDS applies a majority voting in a moving window for attack detection, thus its detection delay is at least half the window size. The optimum window sizes reported in [28] for each device are used for comparisons.

B. Testbed Results

We next present a hardware implementation of our proposed attack detection and mitigation algorithm. Even though there are several dataset already available, none of them addresses stealthy DDoS attacks. In existing works, even a 300% increase from the nominal data rate of a device is considered low-rate DDoS, e.g., [20]. However, due to the

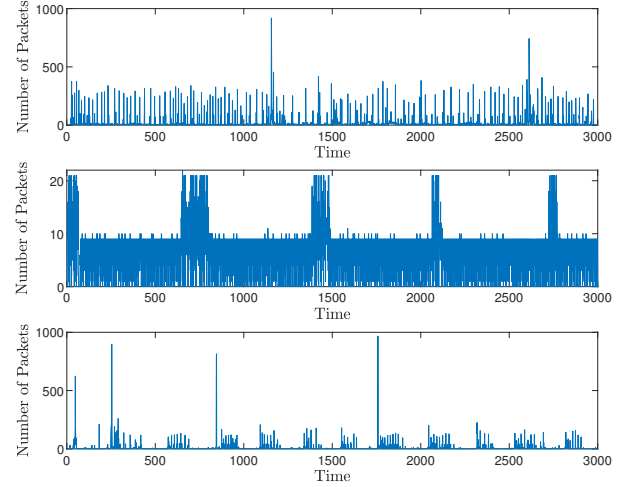


Fig. 9. Time series for number of packets sent by a computer (top), a NodeMCU (middle), and an Amazon Echo Show (bottom).

exponential increase in the number of IoT devices, much lower data rates, e.g., a 30% increase, might be sufficient for performing effective DDoS attacks, as exemplified by the recent Mongolian DDoS attacks [3]. Also, most of the datasets seem to concentrate only on vulnerable devices, but in a typical network, there are also devices that cannot be compromised easily, and account for a significant amount of background data. We first present the testbed setup, and then provide the experimental results using the testbed data.

Testbed Setup: To demonstrate a typical IoT network, we collected the network traffic data from devices that were connected via Wi-Fi to an access point, which is wire connected to a router. For sniffing the network traffic, we performed port mirroring on the router, and recorded the data using Wireshark. The goal here is to design a setup that is as close to a real life scenario as possible for studying stealthy DDoS attacks. We consider a network consisting of 15 popular IoT devices, namely a laptop computer, a tablet, 7 NodeMCUs, 4 smart switches, an Amazon Echo Show device, and a security camera as shown in Fig. 7. The purpose of using a computer and a tablet is to consider devices which cannot be easily compromised yet they account for a significant amount of traffic passing through the router. The NodeMCUs, which may represent various IoT devices on the market, are configured to

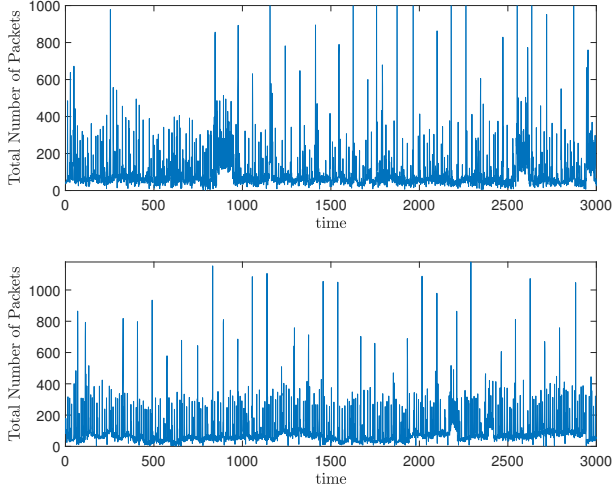


Fig. 10. Time series for total number of packets sent over the entire network without (top) and with the stealthy Http flooding attack (bottom).

TABLE II
TESTBED ATTACK CHARACTERISTICS

Attack Name	Attack Type	Magnitude	Compromised Device
Http Flooding	Application	Low-Rate	Node MCU
ICMP Flooding	Volumetric	Low-Rate	Node MCU
Ping of Death	Protocol	Low-Rate	Laptop
UDP Flooding	Volumetric	High-Rate	Laptop

update their status on a local server. A Raspberry Pi acts as a command and control (C&C) server which is used to start and stop attacks. In Fig. 8, we show the setup for our hardware implementation. A comparison between the number of packets transmitted from a computer, a NodeMCU, and an Amazon Echo Show over a period of time is shown in Fig. 9. It is seen that each device has two major operating states: an active state and an idle state. However, there is a stark difference between the transmission patterns for each device.

Attack Models & Results: We implemented 4 different attacks, as shown in Table II. The data was captured in pcap format by using Wireshark and is publicly available ². In Http flooding, at $t = 80$ sec., we slightly increase the mean rates of updating the server for two of the NodeMCUs. Particularly, there is small increase in the number of *Get* requests from two NodeMCUs. To depict that the attack magnitude does not have to be consistent across all devices, we increase the mean packet rate of NodeMCU 1 by 10% and of NodeMCU 2 by 30%. In Fig. 10, we show a comparison between the total number of packets sent over the entire network with and without the stealthy Http attack. It is seen that due to the low-rate nature of the attack, there is no visible increase in the total number of packets. Hence, filter-based methods which monitor the total number of packets transmitted in the network, would fail to detect such a stealthy attack. In all cases, the input to the proposed IDS is the number of each packet type from each individual device.

In Fig. 11, we see that from $t = 0$ to $t = 80$, the ODIT statistic under the Http attack does not increase considerably,

TABLE III
NOMINAL DATA MODEL FOR DIFFERENT IOT DEVICES.

IoT Devices	Data Model			
	Active State Probability	Mean Packet Rate	Idle State Probability	Mean Packet Rate
Thermostat	0.25	25	0.75	5
Smart Light	0.05	10	0.95	5
Security Camera	1	80	0	0
Smart Printer	0.05	75	0.95	5
Smart TV	0.3	120	0.7	10

but after $t = 80$, it steadily increases. This figure is taken from a real-time demonstration which is available online and explains in detail the testbed setup and attack generation ³. By adjusting the threshold we can have a trade-off between the detection delay and number of false alarms. In Fig. 12, we compare the proposed ODIT-based IDS with the information metric-based algorithm proposed in [20] in terms of average detection delay under all attack cases. The method in [20] uses an information distance metric based on the generalized (Rényi) entropy. It uses a window to compute the information metric on the aggregate traffic at each node, which causes loss in time resolution, and also in early detection ability. Since ODIT monitors each packet type from every individual device, it is able to detect the attack with a much smaller detection delay for the same false alarm rates.

Similarly in the other three attack cases, ODIT quickly and accurately detects the attacks by closely monitoring the data traffic in each type from each device thanks to its multivariate nature. It takes much longer for the information metric IDS to detect the attacks at the same false alarm rate as there is no significant increase in the number of total number of packets. Although in ICMP flooding, the data rates of the two NodeMCUs are again increased 10% and 30%, this time it is easier for the information metric method to detect since the number of ICMP packets in the network is much smaller than the number of Http packets. In the case of ping of death attack, which results in an increase in the number of ICMP packets, the proposed IDS achieves zero detection delay in all trials. Finally, in the UDP flood attack, we considered a higher attack rate by increasing the nominal data rate of the laptop by 100%. In this case, the performance of information metric method improves, but ODIT still outperforms it by detecting the attack under 0.2 second on average for a false alarm rate of 0.01.

C. Simulation Results

We finally present simulation results to evaluate the performance of the proposed IDS in a large network with many nodes, where a stealthy DDoS attack from many compromised IoT devices can actually take down a server.

Simulation Setup: The simulation setup consists of 10 nodes each of which monitors 100 devices (Fig. 1). The IoT devices considered here are those that are most likely to be compromised or devices that are present in every smart home. The devices are assumed to have two states of operation, idle state and active state. The assumed probabilities of the devices

²<https://github.com/kevaldoshi17/DDoSAttackDetection>

³<https://youtu.be/zQexZgB5AMs>




Fig. 11. Live implementation of the proposed IDS in the stealthy Http flooding attack case. At $t = 80$ attack starts by increasing the data rate of device 1 by 10% (third from top) and device 2 by 30% (bottom). In such a stealthy attack, there is no visible change in the total number of packets in the network (second from top). The detection statistic of the proposed IDS steadily increases right after the attack, and alarms when it crosses the threshold (top).

being in idle or active state are given by Table III. We inject the attack traffic of different rates into the simulated dataset, and apply the proposed detection and mitigation scheme to diagnose these attacks. We perform this repeatedly for different kinds of scenarios in which various combinations of devices get attacked to compute the Average Detection Delay vs. False Alarm Rate (i.e., false alarm probability) performance.

In each node, we assume data from each IoT device is independent and identically distributed (iid) following the pattern given in Table III. The probabilities listed are based on heuristics and standard day-to-day usage of the mentioned devices. For example, a smart TV is considered to be used approximately 7 to 8 hours in a day, so its active state probability is given as 0.3. With the shown probabilities, devices may or may not switch state after a session. We consider the following session durations: 5 sec. for thermostat, 10 sec. for smart light, 40 sec. for smart printer, 900 sec. for smart TV, and “always on” for security camera. The mean packet rates are determined by considering the amount of data that is transmitted per second and the average packet size. For each device, the number of packets are generated using the mixture of two Gaussian distributions defined by the active state probability, mean packet rates and a common standard deviation, chosen as 5. To obtain the number of packets, the generated real-valued numbers are rounded to the nearest nonnegative integer. For the purpose of simulations, the training data consists of 40 hours of attack-free data from 100 different devices in each node.

Attack Model: Here we consider a practical scenario in which the IoT devices could be under attack, but the node is assumed to be secure. To parameterize the attack size, we consider 10% of the devices to be compromised. The attacked devices are randomly selected to assume a general

model. During the attack phase, the data rates of the selected devices are increased slightly by 10%. To demonstrate the effectiveness of such a stealthy attack, we plotted in Fig. 13 the total number of packets received by the server when attacked from 100,000 devices with 10% increase in their data rates.

Comparisons: We compare our proposed model with an IDS based on cooperative CUSUM [35], which knows the exact parameters of the nominal model and the anomalous model. CUSUM knows exactly the mean and standard deviation of the Gaussian distribution, as well as the probability of being active for each device. Note that due to rounding to the nearest nonnegative integer value, the real probability distribution of number of packets deviates from the generative bimodal Gaussian. Hence, the proposed ODIT detector even sometimes outperforms CUSUM, which exactly knows the generative Gaussian model. The results for Average Detection Delay vs False Positive Rate are shown in Fig. 14. We see that the cooperative ODIT-based IDS, proposed in Section IV-B, performs better than the clairvoyant CUSUM detector, which exactly knows the generative probabilistic model, for false alarm rates less than 0.1. It significantly outperforms the information metric method proposed in [20], which monitors the aggregate traffic at each node. In Fig. 14, it is seen that the cooperation among nodes facilitates earlier detection by our algorithm (ODIT vs. Cooperative ODIT). Through the proposed computationally efficient cooperation scheme, given in Section IV-B, the ODIT-based IDS is able to handle large networks with thousands of devices. This result can be easily extended to even larger networks with millions of devices. Finally, in Fig. 15, we evaluate the mitigation performance of the proposed method (see Section IV-C). Since the method in [20] monitors the aggregate traffic at nodes, it

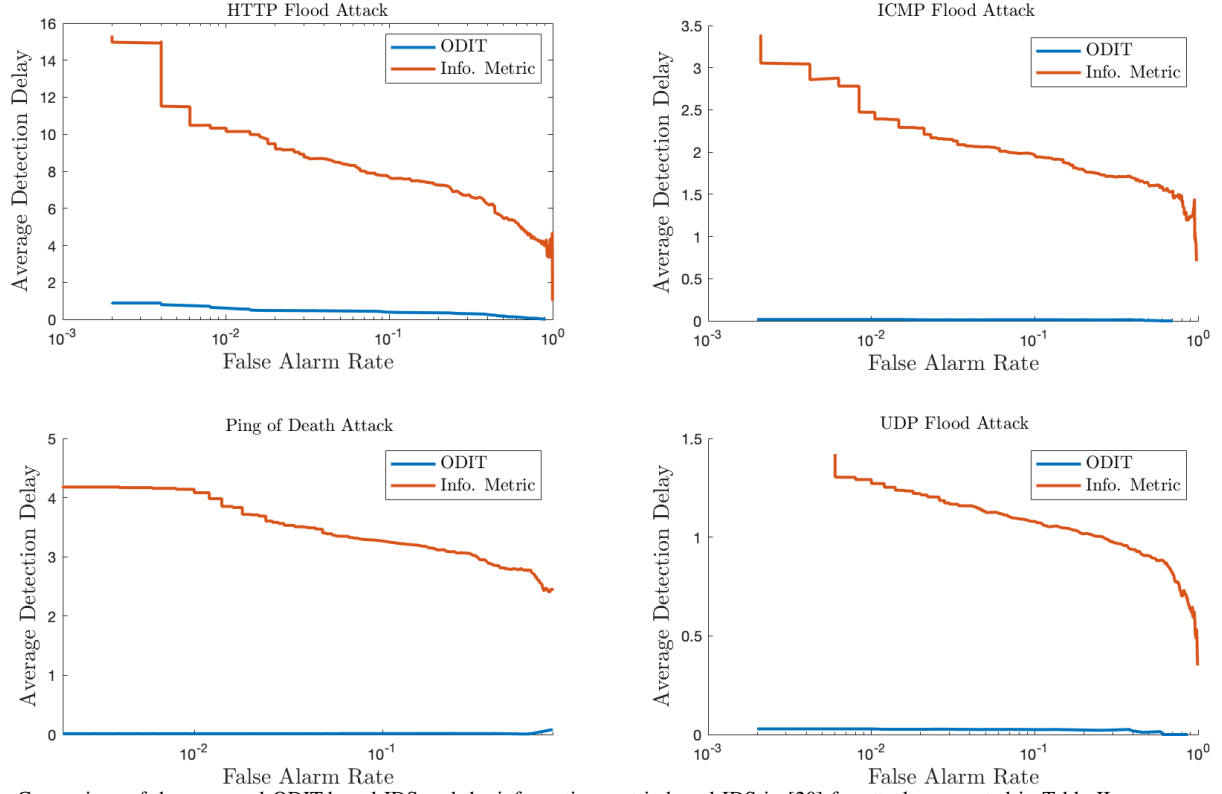


Fig. 12. Comparison of the proposed ODIR-based IDS and the information metric-based IDS in [20] for attacks presented in Table II.




Fig. 13. Impact of stealthy DDoS attack on the server. The attack consists of 100,000 devices with 10% increase in their data rates.

is not straightforward for it to detect the attacking devices. Thus, to evaluate our mitigation performance, we consider the data filtering method, which simply applies a threshold to the observed raw data. The reported Area Under the Curve (AUC) values in Fig. 15 illustrate the successful mitigation performance of the proposed method under a challenging stealthy attack scenario.

VI. LIMITATIONS AND FUTURE WORK

In this work, we proposed a novel intrusion detection system which is capable of quickly and accurately detecting and mitigating a broad set of IoT-empowered attacks, in particular stealthy low-rate DDoS attacks. However, there are still some limitations which need to be addressed to make the system more robust to attacks in the future. First, it is assumed that the nominal behavior of the devices does not change over time, so the IDS needs to be trained only once. However,

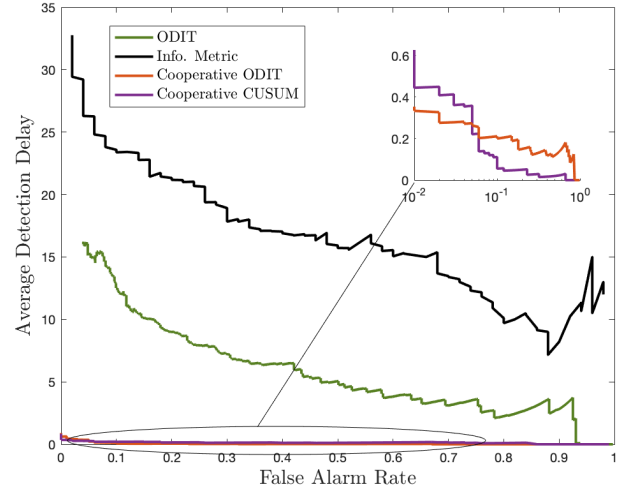


Fig. 14. Average detection delay vs. False positive rate for the proposed cooperative ODIR-based IDS, the IDS based on cooperative CUSUM [35], and information metric-based IDS [20]. The network consists of 10 nodes each of which has 100 devices connected to it. 10% of the devices attack with 10% increase in data rate with respect to their nominal rates.

in a real system implementation the IDS needs to be updated periodically. Secondly, feature extraction plays an important role as number of packets or packet size might not always exactly represent the characteristics of a real network. For future work, we plan to investigate other aspects of dynamic networks such as continual learning under changing nominal network traffic.




Fig. 15. Mitigation performance of ODIR vs. Data filtering method which applies a threshold on the data rates. After detecting an attack, the proposed IDS successfully identifies attacked devices, and blocks their traffic.

VII. CONCLUSION

With the proliferation of IoT devices, and the ease of triggering DoS attacks even by unsophisticated malicious parties, there is an increasing need for developing solutions to DDoS via IoT, especially the recent stealthy DDoS attacks. In this context, we presented a general and emerging threat model for hierarchical IoT networks. We then introduced a novel intrusion detection and mitigation framework that employs an online, scalable and nonparametric anomaly detection algorithm. Through real and simulated data, as well as an IoT testbed we evaluated the performance of proposed detection and mitigation scheme under challenging stealthy DDoS attack scenarios. Applications of the proposed scheme to large and dynamic networks with varying number of devices were also considered.

REFERENCES

- [1] Arsalan Mosenia and Niraj K Jha. A comprehensive study of security of internet-of-things. *IEEE Transactions on Emerging Topics in Computing*, 5(4):586–602, 2017.
- [2] Imperva. Types of ddos attacks. <https://www.imperva.com/learn/application-security/ddos-attacks/>.
- [3] Nexusguard. Q3 2018 threat report: Distributed denial of service (ddos). https://www.nexusguard.com/hubfs/2019%20PTC/Nexusguard_Q3%202018%20Threat%20Report.pdf, 2019.
- [4] Elisa Bertino and Nayeem Islam. Botnets and internet of things security. *Computer*, (2):76–79, 2017.
- [5] Yasin Yilmaz and Suleyman Uludag. Mitigating iot-based cyberattacks on the smart grid. In *Machine Learning and Applications (ICMLA), 2017 16th IEEE International Conference on*, pages 517–522. IEEE, 2017.
- [6] Laurence Goasduff. Gartner says 5.8 billion enterprise and automotive iot endpoints will be in use in 2020. <https://www.gartner.com/en/newsroom/press-releases/2019-08-29-gartner-says-5-8-billion-enterprise-and-automotive-iot>, 2019.
- [7] Constantinos Kolias, Georgios Kambourakis, Angelos Stavrou, and Jeffrey Voas. Ddos in the iot: Mirai and other botnets. *Computer*, 50(7):80–84, 2017.
- [8] Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J Alex Halderman, Luca Invernizzi, Michalis Kallitsis, et al. Understanding the mirai botnet. In *USENIX Security Symposium*, pages 1092–1110, 2017.
- [9] Mirai source code. <https://github.com/jgamblin/Mirai-Source-Code>.
- [10] Keval Doshi, Mahsa Mozaffari, and Yasin Yilmaz. Rapid: Real-time anomaly-based preventive intrusion detection. In *Proceedings of the ACM Workshop on Wireless Security and Machine Learning*, pages 49–54, 2019.
- [11] Alma D Lopez, Asha P Mohan, and Sukumaran Nair. Network traffic behavioral analytics for detection of ddos attacks. *SMU Data Science Review*, 2(1):14, 2019.
- [12] Ahmad Riza'ain Yusof, Nur Izura Udzir, and Ali Selamat. Systematic literature review and taxonomy for ddos attack detection and prediction. *International Journal of Digital Enterprise Technology*, 1(3):292–315, 2019.
- [13] Mark Shtern, Roni Sandel, Marin Litoiu, Chris Bachalo, and Vasileios Theodorou. Towards mitigation of low and slow application ddos attacks. In *2014 IEEE International Conference on Cloud Engineering*, pages 604–609. IEEE, 2014.
- [14] Enrico Cambiaso, Gianluca Papaleo, and Maurizio Aiello. Taxonomy of slow dos attacks to web applications. In *International Conference on Security in Computer Networks and Distributed Systems*, pages 195–204. Springer, 2012.
- [15] Prabhakaran Kasinathan, Claudio Pastrone, Maurizio A Spirito, and Mark Vinkovits. Denial-of-service detection in 6lowpan based internet of things. In *2013 IEEE 9th international conference on wireless and mobile computing, networking and communications (WiMob)*, pages 600–607. IEEE, 2013.
- [16] Vipindev Adat and BB Gupta. A ddos attack mitigation framework for internet of things. In *Communication and Signal Processing (ICCSP), 2017 International Conference on*, pages 2036–2041. IEEE, 2017.
- [17] Krushang Sonar and Hardik Upadhyay. An approach to secure internet of things against ddos. In *Proceedings of International Conference on ICT for Sustainable Development*, pages 367–376. Springer, 2016.
- [18] Menghao Zhang, Guanyu Li, Shicheng Wang, Chang Liu, Ang Chen, Hongxin Hu, Guofei Gu, Qianqian Li, Mingwei Xu, and Jianping Wu. Poseidon: Mitigating volumetric ddos attacks with programmable switches. In *Proceedings of NDSS*, 2020.
- [19] Ping Du and Shunji Abe. Ip packet size entropy-based scheme for detection of dos/ddos attacks. *IEICE transactions on information and systems*, 91(5):1274–1281, 2008.
- [20] Yang Xiang, Ke Li, and Wanlei Zhou. Low-rate ddos attacks detection and traceback by using new information metrics. *IEEE transactions on information forensics and security*, 6(2):426–437, 2011.
- [21] Min Suk Kang, Soo Bum Lee, and Virgil D Gligor. The crossfire attack. In *2013 IEEE symposium on security and privacy*, pages 127–141. IEEE, 2013.
- [22] Lei Xue, Xiaobo Ma, Xiapu Luo, Edmond WW Chan, Tony TN Miu, and Guofei Gu. Linkscope: Toward detecting target link flooding attacks. *IEEE Transactions on Information Forensics and Security*, 13(10):2423–2438, 2018.
- [23] Lei Wang, Qing Li, Yong Jiang, and Jianping Wu. Towards mitigating link flooding attack via incremental sdn deployment. In *2016 IEEE Symposium on Computers and Communication (ISCC)*, pages 397–402. IEEE, 2016.
- [24] Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. *ACM computing surveys (CSUR)*, 41(3):15, 2009.
- [25] Rohan Doshi, Noah Apthorpe, and Nick Feamster. Machine learning ddos detection for consumer internet of things devices. *arXiv preprint arXiv:1804.04159*, 2018.
- [26] Sven Nömm and Hayretin Bahşi. Unsupervised anomaly based botnet detection in iot networks. In *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 1048–1053. IEEE, 2018.
- [27] Mehmet Necip Kurt, Yasin Yilmaz, and Xiaodong Wang. Real-time nonparametric anomaly detection in high-dimensional settings. *arXiv preprint arXiv:1809.05250*, 2018.
- [28] Yair Meidan, Michael Bohadana, Yael Mathov, Yisroel Mirsky, Asaf Shabtai, Dominik Breitenbacher, and Yuval Elovici. N-baiotnetwork-based detection of iot botnet attacks using deep autoencoders. *IEEE Pervasive Computing*, 17(3):12–22, 2018.
- [29] Yisroel Mirsky, Tomer Doitshman, Yuval Elovici, and Asaf Shabtai. Kitsune: an ensemble of autoencoders for online network intrusion detection. In *Proceedings of NDSS*, 2018.
- [30] Lersak Limwiwatkul and Arnon Rungswang. Distributed denial of service detection using tcp/ip header and traffic measurement analysis. In *Communications and Information Technology, 2004. ISIT 2004. IEEE International Symposium on*, volume 1, pages 605–610. IEEE, 2004.
- [31] Imperva. Breaking down mirai: An iot ddos botnet analysis. <https://www.imperva.com/blog/malware-analysis-mirai-ddos-botnet/>.
- [32] Yasin Yilmaz. Online nonparametric anomaly detection based on geometric entropy minimization. In *Information Theory (ISIT), 2017 IEEE International Symposium on*, pages 3010–3014. IEEE, 2017.

- [33] Michèle Basseville, Igor V Nikiforov, et al. *Detection of abrupt changes: theory and application*, volume 104. Prentice Hall Englewood Cliffs, 1993.
- [34] Kumar Sricharan and Alfred O Hero. Efficient anomaly detection using bipartite k-nn graphs. In *Advances in Neural Information Processing Systems*, pages 478–486, 2011.
- [35] Yajun Mei. Efficient scalable schemes for monitoring a large number of data streams. *Biometrika*, 97(2):419–433, 2010.
- [36] Marius Muja and David G Lowe. Scalable nearest neighbor algorithms for high dimensional data. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (11):2227–2240, 2014.
- [37] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *2008 Eighth IEEE International Conference on Data Mining*, pages 413–422. IEEE, 2008.
- [38] Bernhard Scholkopf and Alexander J Smola. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2001.
- [39] Markus M Breunig, Hans-Peter Kriegel, Raymond T Ng, and Jörg Sander. Lof: identifying density-based local outliers. In *ACM sigmod record*, volume 29, pages 93–104. ACM, 2000.



Keval Doshi received the B.Sc. degree in Electronics and Communications Engineering from Gujarat Technological University, India, in 2017. He is currently a Ph.D. student at the Electrical Engineering Department at the University of South Florida, Tampa. His research interests include machine learning, computer vision, and cybersecurity.



Yasin Yilmaz (S'11-M'14-SM'20) received the Ph.D. degree in Electrical Engineering from Columbia University, New York, NY, in 2014. He is currently an Assistant Professor of Electrical Engineering at the University of South Florida, Tampa. He received the Collaborative Research Award from Columbia University in 2015. His research interests include statistical signal processing, machine learning, and their applications to computer vision, cybersecurity, IoT networks, energy systems, transportation systems, and communication systems.



Suleyman Uludag received the Ph.D. degree in Computer Science from DePaul University, Chicago, IL, in 2007. He is currently an Associate Professor of Computer Science at the University of Michigan - Flint. He received the Fulbright U.S. Scholar Program Core Award in 2012 and 2018. His research interests include security, privacy, and optimization of data collection particularly as applied to the Smart Grid and Intelligent Transportation Systems.