

RAPID: Real-time Anomaly-based Preventive Intrusion Detection

Keval Doshi
University of South Florida
Tampa, USA
kevaldoshi@mail.usf.edu

Mahsa Mozaffari
University of South Florida
Tampa, USA
mmozaaffari@mail.usf.edu

Yasin Yilmaz
University of South Florida
Tampa, USA
yasiny@usf.edu

ABSTRACT

Intrusion detection systems (IDSs) today face key limitations with respect to detection and prevention of challenging IoT-empowered attacks. We address these limitations by proposing a novel IDS called RAPID, which is based on an online scalable anomaly detection and localization approach. We show that the anomaly detection algorithm is asymptotically optimal under certain conditions, and comprehensively analyze its computational complexity. Considering a real dataset and an IoT testbed we demonstrate the use of RAPID in two different IoT-empowered cyber-attack scenarios, namely high-rate DDoS attacks and low-rate DDoS attacks. The experiment results show the quick and accurate detection and prevention performance of the proposed IDS.

KEYWORDS

IoT networks, DDoS attacks, anomaly detection, sequential detection, nonparametric method

ACM Reference Format:

Keval Doshi, Mahsa Mozaffari, and Yasin Yilmaz. 2019. RAPID: Real-time Anomaly-based Preventive Intrusion Detection. In *Proceedings of Submission to WiSec '19*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

The past decade has seen an exponential rise in the number of devices connected to the Internet with 8.4B devices in 2017 and expected to hit 20B by 2020 [21]. With an ever growing global Internet audience using various online services, the continuous integrity and availability of the Internet has never been more important. However, there has been a proportional increase in the number of cyber-attacks as well. This rise can be attributed to the vulnerabilities that Internet of Things (IoT) devices present [18, 20]. Recent studies have demonstrated that IoT devices such as cameras and locks can be compromised by an adversary and manipulated to perform attacks [8, 17]. The abundance of low-security IoT devices has enabled new genre of attacks. One such attack was caused by the Mirai botnet, that was launched in 2016, which led to one of the most prolific series of Distributed Denial of Service (DDoS) attacks

in history [2, 11]. This particular malware infected numerous IoT devices (primarily older routers and IP cameras), and reached data rates higher than 600Gbps. IoT botnets, such as Mirai, also enable "stealth" low-rate DDoS attacks, that are quite challenging to detect and mitigate due to the highly distributed nature and the low-rate increase in the local data traffic which looks very similar to the nominal traffic [11]. In this paper, we consider the real-time detection and prevention for two challenging IoT-empowered attack types described above: (i) High-rate DDoS via IoT botnets such as Mirai, (ii) Low-rate DDoS via IoT botnets. To address them we propose an intrusion detection and prevention system called RAPID (Real-time Anomaly-based Preventive Intrusion Detection), which is indeed applicable to a broader set of attacks, such as the MadIoT attack targeting the power grid [20].

1.1 Challenges and Solutions

Next, we summarize the challenges Anomaly-based Intrusion Detection and Prevention System (AIDPS) faces in dealing with the IoT-empowered attacks, and our key ideas to tackle them.

- (C1) *Timely and accurate detection*: Due to the highly interconnected IoT ecosystem including the Internet, timely and accurately detection and mitigation of attacks is crucial. A major criticism against AIDPSs is their high false alarm rates. This criticism is based on the fact that most AIDPS works like statistical outlier detection methods, which are known to be prone to frequent false alarms [3]. Our key idea to address (C1) is to use a novel sequential change detection method. Sequential methods, such as the Cumulative Sum (CUSUM) test, aim to minimize the detection delay and at the same time satisfy a false alarm constraint by looking for persistent outliers.
- (C2) *Curse of dimensionality and heterogeneity*: IoT networks typically consist of many devices of different types with different nominal behaviors. Since anomaly detectors in general learn a statistical description of nominal behavior and detect significant deviations from that, the high-dimensionality and heterogeneity of IoT networks pose a significant challenge for learning the network-wide nominal behavior. RAPID is a nonparametric AIDPS that scales well to high-dimensional and heterogeneous IoT networks. The proposed method, in a computationally efficient way, learns a nominal statistic that asymptotically well approximates the network-wide nominal probability distribution (Section 3.1).
- (C3) *Minimally invasive prevention*: For successful mitigation of cyber-attacks, the attackers should be accurately identified and blocked with a minimal interruption of services to benign users which is difficult due to the distributed and possibly low-rate nature of DDoS attacks. To this end, after

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Submission to WiSec '19, Miami, USA,

© 2019 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

timely and accurate detection of attacks, RAPID applies a statistical test, called t -test, on the detection statistic (Section 3.2). Under challenging scenarios, we show that this method achieves high true positive rate and low false positive rate.

- (C4) *Network integration*: A practical solution should be such that it could be integrated seamlessly with the existing infrastructure. Many existing anomaly detection algorithms have high computational overhead and require specialized hardware, such as GPU, making them infeasible in a practical setting. Due to its computational efficiency, RAPID can easily run on inexpensive devices such as a Raspberry Pi, facilitating its seamless integration with the existing architecture. The computational complexity analysis is given in Section 3.3.

1.2 Related Works

In the recent years, there has been a sharp escalation in the number of cyber-attacks on infrastructure networks, leading to a subsequent rise in both industry and academic research for detecting network related attacks. For example, there has been an increase in the popularity of entropy-based methods to detect low-rate DDoS attacks. Chen et al. [5] introduce two novel information metrics to detect the such attacks: Fourier Power Spectrum Entropy (FPSE) and Wavelet Power Spectrum Entropy (WPSE). They propose that since the energy of a low-rate DDoS attack is concentrated in the low frequency range, the two metrics should be able to detect such attacks. In [23], an information metric based algorithm is proposed by Xiang et al. They present two statistical metrics, generalized entropy and information distance to identify low-rate DDoS attacks. They assume that an attack-free network follows a Gaussian distribution, whereas during an attack it follows a Poisson distribution. They make a decision based on the difference in the information metrics between normal traffic and malicious traffic. Another statistical algorithm is proposed by Wu et al. [22] who design their IDS based on Hurst coefficient in order to detect low-rate DDoS attacks. Using experimental results, they show early stage detection of low-rate DDoS attacks. However, they consider increases in traffic rate as high as 600% which are comparatively much easier to detect. Another entropy-based model is presented by Ping Du et al. [6]. They develop an IDS using packet size distribution based on the change in entropy when traffic is affected by an attack. Their results show that they are able to detect short-term as well as long-term attacks. More recently, a new method was proposed in [14] based on deep autoencoders for detecting IoT botnet attacks. They propose to train a deep autoencoder for every device in the network, and they achieve low false alarm rates, however due to the training of autoencoders being computationally expensive, this method is difficult to scale to thousands of devices.

1.3 Organization

Section 2 presents the system and attack models for the considered scenarios. In Section 3, we introduce the proposed intrusion detection and prevention system, RAPID, along with an analysis of its computational complexity. In Section 4, the performance of RAPID is evaluated in the two challenging attack scenarios using an available dataset and an IoT testbed with respect to the state-of-the-art detectors. Finally, we provide limitations and motivation for future work in Section 5, and conclusion in Section 6.

2 SYSTEM AND ATTACK MODELS

In this section, we define the considered system and attack models that can be used to launch DDoS attack against a web server or a critical infrastructure such as Smart Grid.

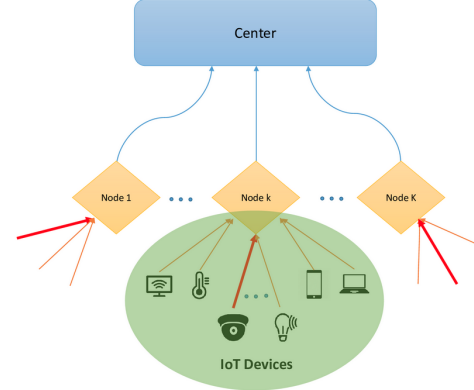


Figure 1: DDoS threat model. Bold arrows imply an increased packet rate.

System Model: We consider a setup in which each IoT device sends its data to a parent node connected to it, such as a router, as shown in Figure 1. Parent nodes direct the data traffic to a center, such as a web server, data center or utility center. Depending on the size of network to be monitored, the hierarchical architecture can be scaled to include multiple levels of such parent nodes. Furthermore, each parent node may represent a variety of networks, such as a smart home with tens of devices, or a university network with thousands of devices. Due to the possibility of widely-distributed compromised devices, it is not convenient to have a central IDS running at the center for quickly and accurately mitigating a major DDoS attack, as illustrated by the Mirai botnet [10]. Hence, in this paper we consider local IDSs for a global solution. Focusing on a single node we propose that each node implements a local IDS to effectively detect and mitigate volumetric DDoS attacks. Each device typically has different data communications characteristics. In particular, the data content is typically different (e.g., a thermostat would have considerably smaller packet sizes as compared to a security camera), and the communications protocol used might be different (such as TCP, UDP or HTTP).

Attack 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, which is called flooding. Particularly, we consider a threat model in which some 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 or not. That is, we also consider the case where it is not possible to inspect the data content, which is a prerequisite for many IDS algorithms, e.g., [7], [12]. Due to the proliferation of IoT, cyber-criminals can launch widely-distributed and highly-effective low-rate DDoS attacks that can bypass conventional filters and IDSs. Hence, in this paper we

also study DDoS attacks with increase in data rates as low as 20%. There are existing works which consider low-rate DDoS attacks, nevertheless the considered increase rates are still significantly higher than what we consider in this paper (e.g., 300% of the baseline in [23]).

3 PROPOSED REAL-TIME ANOMALY-BASED IDS

In this section, we present our detection and prevention strategy, called RAPID, followed by a computational complexity analysis. Before adapting RAPID to specific scenarios in Section 4, we explain it in this section using a generic model in which the IDS runs locally on a node with a number of devices connected to it. The IDS jointly monitors the data received from the devices in real-time, and raises an alarm when there is sufficient statistical evidence. We next discuss the detection strategy, and then the prevention strategy in detail.

3.1 Timely and Accurate Detection

The proposed IDS runs on a local node which observes a d -dimensional instance \mathbf{x}_t at each time t . The d dimensions correspond to the devices connected to the node and the data is considered to be the number of packets sent by the device at each time instance. Based on the computational resources available and the requirement, it is straightforward to extend the observed data (i.e., features) to include other possible observations such as packet size, number of packets from each packet type, etc. To learn the nominal traffic behavior RAPID takes an attack-free training dataset $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$, and normalizes each data dimension, e.g., using the mean and standard deviation, or upper and lower bounds, to handle heterogeneous ranges among dimensions. For each point in \mathcal{X} , it finds the k nearest neighbors (k NN) among the other training points, and specifically the distance to k th nearest neighbor $r_k(\mathbf{x}_i)$, $i = 1, \dots, N$. Training is finished by selecting the $(1 - \alpha)$ th percentile $r_k(\mathbf{x}_{(K)})$ of $\{r_k(\mathbf{x}_1), \dots, r_k(\mathbf{x}_N)\}$, where $K = \text{round}[(1 - \alpha)N]$, to notice statistical deviation at significance level α , e.g., 0.05.

In the test phase, as the observations arrive sequentially, at each time t the proposed IDS computes the k NN distance $r_k(\mathbf{x}_t)$ with respect to the training points in \mathcal{X} , and then the instantaneous attack evidence Δ_t :

$$\Delta_t = d [\log r_k(\mathbf{x}_t) - \log r_k(\mathbf{x}_{(K)})], \quad (1)$$

where d is the number of data dimensions. This specific form of Δ_t enables its asymptotic optimality in the minimax sense, as shown in Theorem 1. Finally, the proposed IDS updates a running decision statistic s_t ,

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

and declares an alarm the first time s_t exceeds a threshold at time T given by,

$$T = \min\{t : s_t \geq h\}. \quad (3)$$

The choice for the decision threshold h sets a balance for the trade-off between the two objectives, smaller detection delay and smaller false alarm rate. Higher h helps decrease the false alarm rate, but it also increases the average detection delay; and vice versa for lower h . Similarly, the choice k , the number of nearest neighbors in computing the distance $r_k(\mathbf{x})$, and the significance level α determine a trade-off between robustness to nominal outliers

(i.e., noise) and sensitivity to attacks. While smaller k and larger α yield higher sensitivity to attacks, i.e., cause a quicker rise in s_t and therefore quicker detection, they are also less robust to statistical outliers, i.e., more prone to false alarms. Note that the significance level α does not have a central role in the proposed IDS as it is not used to decide on attack for each observation, as opposed to the anomaly detection methods based on significance tests (e.g., outlier detection as in [9]). Indeed, it is auxiliary to the decision threshold h , which is the main parameter that controls the trade-off between detection delay and false alarm in the proposed IDS. Hence, α is first set to a typical value such as 0.05, and then h is chosen to satisfy a false alarm rate. The instantaneous attack evidences computed at each time t using (1) are accumulated over time using (2) to make an attack decision.

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

$$\Delta_t \xrightarrow{p} \log \frac{f_1(\mathbf{x}_t)}{f_0(\mathbf{x}_t)} \text{ as } N \rightarrow \infty, \quad (4)$$

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

PROOF. Consider a hypersphere $\mathcal{S}_t \in \mathbb{R}^d$ centered at \mathbf{x}_t with radius $r_k(\mathbf{x}_t)$, the k NN distance of \mathbf{x}_t with respect to the training set \mathcal{X} . The maximum likelihood estimate for the probability of a point being inside \mathcal{S}_t under f_0 is given by k/N . It is known that, as the total number of points grow, this binomial probability estimate converges to the true probability mass in \mathcal{S}_t in the mean square sense [1], i.e., $k/N \xrightarrow{L^2} \int_{\mathcal{S}_t} f_0(\mathbf{x}) d\mathbf{x}$ as $N \rightarrow \infty$. Hence, the probability density estimate $\hat{f}_0(\mathbf{x}_t) = \frac{k/N}{V_d r_k(\mathbf{x}_t)^d}$, where $V_d r_k(\mathbf{x}_t)^d$ is the volume of \mathcal{S}_t , converges to the actual probability density function, $\hat{f}_0(\mathbf{x}_t) \xrightarrow{p} f_0(\mathbf{x}_t)$ as $N \rightarrow \infty$, since \mathcal{S}_t shrinks and $r_k(\mathbf{x}_t) \rightarrow 0$. Similarly, considering a hypersphere $\mathcal{S}_{(K)} \in \mathbb{R}^d$ around $\mathbf{x}_{(K)}$ which includes k points with its radius $r_k(\mathbf{x}_{(K)})$, we see that as $N \rightarrow \infty$, $r_k(\mathbf{x}_{(K)}) \rightarrow 0$ and $\hat{f}_0(\mathbf{x}_{(K)}) = \frac{k/N}{V_d r_k(\mathbf{x}_{(K)})^d} \xrightarrow{p} f_0(\mathbf{x}_{(K)})$. Assuming a uniform distribution $f_1(\mathbf{x}) = f_0(\mathbf{x}_{(K)})$, $\forall \mathbf{x}$, we conclude with $\log \frac{V_d r_k(\mathbf{x}_{(K)})^d}{V_d r_k(\mathbf{x}_t)^d} = d [\log r_k(\mathbf{x}_t) - \log r_k(\mathbf{x}_{(K)})] \xrightarrow{p} \log \frac{f_1(\mathbf{x}_t)}{f_0(\mathbf{x}_t)}$ as $N \rightarrow \infty$. \square

3.2 Minimally Invasive Prevention

In the previous part, we presented a method for quick and accurate detection of attacks causing anomaly in the observations. However, further analysis after detection may be required in many cases for an effective attack prevention strategy that minimizes the interruption to regular services. For instance, detecting a DDoS attack is not sufficient for ensuring the availability of system to legitimate users. An effective localization strategy that determines the attacking devices and blocks the data traffic originating from them is needed.

To this end, after an attack is detected, we perform an in-depth analysis by examining the recent k NN distances of every data dimension. Firstly, the attack onset time is estimated as the last time the detection statistic s_t was zero, i.e., $\tau = \max\{t < T : s_t = 0\}$, where T is the detection time. Then, for each dimension j , we use a t-test to decide whether the average of M recent k NN distances $\{r_k(x_{\tau+1}^j), \dots, r_k(x_{\tau+M}^j)\}$ is significantly greater than μ_j , average of nominal distances $\{r_k(x_1^j), \dots, r_k(x_N^j)\}$. Since typically the number of training points N is large, we assume the variance of nominal distances is negligible, and for each dimension j compute the t statistic as

$$t_j = \frac{\bar{r}_j - \mu_j}{\frac{\sigma_j}{\sqrt{M}}}, \quad (5)$$

where \bar{r}_j , σ_j and M are the sample mean (i.e., average), sample standard deviation and sample size of $\{r_k(x_{\tau+1}^j), \dots, r_k(x_{\tau+M}^j)\}$. Finally, we decide that the data dimension j is anomalous if $t_j \geq q_M$, where q_M is a certain percentile, such as 95th, of Student's t -distribution with $M - 1$ degrees of freedom. After determining the anomalous dimensions, we block the traffic from device j if $t_j \geq q_M$. The percentile q_M is selected according to a desired significance level, and controls the balance between sensitivity to true anomalies and robustness to nominal outliers. Given the percentile level and M , it is easily found from a lookup table. The sample size M determines a trade-off between the accuracy of t-test and the reaction delay for the prevention strategy. Bigger M increases the accuracy at the expense of longer reaction delay, and vice versa for smaller M values. Note that the reaction time is lower bounded by the detection time T , and it is greater than T for $M > T - \tau$. The proposed algorithm for RAPID is summarized in Algorithm 1.

Algorithm 1: Proposed RAPID algorithm

```

1: Input:  $\mathcal{X}, k, \alpha, h, q_M$ 
2: Train: Find  $r_k(\mathbf{x}_{(K)})$  and  $\{\mu_j\}$  (see (1) and (5))
3: Initialize:  $s = 0, t = 0$ 
4: while  $s < h$  do
5:    $t \leftarrow t + 1$ 
6:   Get new data  $x_t$  and compute  $\Delta_t$  using (1)
7:    $s = \max\{s + \Delta_t, 0\}$ 
8: end while
9: Declare attack at  $T = t$ 
10: for  $j = 1, \dots, d$  do
11:   Compute  $t_j$  using (5)
12:   if  $t_j \geq q_M$  then
13:     Block data from device  $j$ .
14:   end if
15: end for
```

3.3 Computational Complexity

In this section, we analyze the computational complexity of our proposed algorithm. The training phase of RAPID requires the computation of k NN distances between each pair of data instances in the training set. Therefore, the time complexity of training phase is $O(N^2d)$. N data instances of the training set are required to be stored for the testing phase. Hence, the space complexity of the training phase is $O(Nd)$. Similarly, in the testing phase, computation

of the k NN distances of a single data instance to all data points of the training set takes $O(Nd)$ time. The computation of k NN distances for high-dimensional systems with abundance of training data could be the bottleneck in implementing the RAPID algorithm, which we address next.

k NN approximation: Straightforward computation of the k NN distance for a test point requires the computation of its Euclidean distance to each data point in the training set. We employ a scalable k NN distance approximation algorithm called *priority search k -means tree* which was proposed in [16]. This algorithm performs hierarchical clustering by constructing a k -means tree, and approximates the k NN distance by performing a priority search in the k -means tree. Consequently, the computational complexity of the training reduces to $O(N \log N^{\frac{L+KJ_{max}}{\log K}} d)$ where K is the number of clusters and $L \ll N$ is the maximum number of data points to examine.

Experiment: In an experiment, we used this k NN approximation method in our algorithm to improve the computational efficiency. The dimensionality of data is $d = 50$, the training data size is $N = 200,000$, and the anomaly is defined as an increase in the mean of the observations by 3 standard deviations in 10% of the dimensions. We set the branching factor for building the priority search k -means tree as 100, and the maximum number of points to examine during search for the k nearest neighbors as 1000. The average computation time for both RAPID based on the exact and approximate k NN distances is summarized in the Table 1, which presents the time spent for the computation of (1) and (2) per observation. We consider two different cases: (i) data arrives every 1 sec, (ii) data arrives every 0.01 sec. We see that depending on the sampling period, either exact k NN or approximate k NN could be more advantageous. For a sampling period that is smaller than the computation overhead, (see the bottom figure Figure 2), approximate k NN computations are preferred over the exact k NN computations. Whereas, for a large sampling period which the exact method can keep pace with, the delay is mainly due to the sample delay, thus exact k NN computations are favorable in this case, as shown in the top figure in Figure 2.

Table 1: Average computation overhead of (1) and (2) per sample

Average execution time (sec.)	
Exact k NN	Approximate k NN
0.0472	0.0051

4 PERFORMANCE EVALUATION

In this section, we evaluate the performance of the proposed IDS under two scenarios, namely high-rate DDoS and low-rate DDoS.

4.1 High-Rate DDoS

To evaluate the performance of RAPID in IoT-empowered high-rate DDoS attacks we consider the N-BaIoT dataset, which is collected from a real IoT botnet [14] and available at the UCI Machine Learning Repository. This dataset contains network traffic statistics from various IoT devices under both nominal and attack conditions. The network consists of 9 devices, namely a thermostat, a baby monitor, a webcam, two doorbells, and four security cameras connected via




Figure 2: Comparison between RAPID performances based on exact and approximate k NN distances in terms of seconds for $T_{\text{sampling}} = 1$ sec. (top) and $T_{\text{sampling}} = 0.01$ sec. (bottom). WiFi. We compare RAPID with the deep autoencoder method proposed also in [14] in terms of false alarm rate and detection delay. To be able to compare RAPID with the autoencoder method we applied RAPID to each device separately using the 115-dimensional data provided for each device. In Fig. 3, we present the the detection delay for each case. The x-axis corresponds to the attacked device and y-axis corresponds to the average detection delay when that particular device is attacked. The false alarm rates for the autoencoder based method for each attacked device are $[0.01, 0.012, 0, 0.024, 0.01, 0, 0, 0]$ whereas for RAPID, they are $[0, 0.0001, 0.0001, 0, 0, 0, 0.0007, 0]$. We see that in terms of false alarm rate and average detection delay, RAPID considerably outperforms the autoencoder method, which was shown to work better than several other state-of-the-art methods, namely Isolation Forests [13], SVMs [19] and LOFs [4]. Also, in Figure 3 we see that RAPID has a comparatively small detection delay in all devices. The sequential nature of RAPID is the main reason for the big difference between the performances. The autoencoder method uses a window-based decision rule for real-time operation, which requires the majority of the instances in the window to be labelled as anomalous.

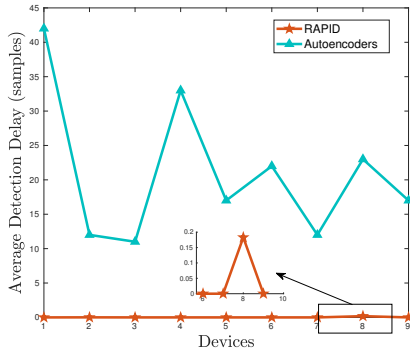


Figure 3: Comparison between the deep autoencoder method [14] and RAPID in terms of quickness of detection.

4.2 Low-rate DDoS

We first provide a motivational study for IoT-empowered low-rate DDoS attacks. As discussed in [15], a DDoS attack with a throughput of more than 10,000 packets/sec can be considered as an effective attack that can take down a small scale server. Based on this

Increase in nominal rate (%)	No. of Security Cameras	No. of Personal Computers	No. of NodeMCUs	Total No. of Devices
25	32	16	16	64
50	16	8	8	32
75	12	6	6	24
100	8	4	4	16

Table 2: Average number of devices required to successfully attack a small scale online server based on increase in their nominal packet rates.

information, we can extrapolate a distributed attack scenario in which there is a slight increase in the mean data rate of some IoT devices. To put things in perspective, we consider three vulnerable IoT devices, namely security cameras, personal computers and NodeMCUs, which can be found in a variety of IoT devices. After monitoring the nominal behavior of these devices in our IoT testbed (see Figure 4), we found that in its active state, a security camera sends approximately 800-1000 packets/sec, a personal computer sends 1800-2600 packets/sec, whereas a NodeMCU might send 30-90 packets/sec depending on its application. Based on these numbers, in Table 2, we present the approximate number of compromised devices required to successfully attack a web server depending on their mean increase. We assume that an attack does not interfere with the normal operation of a device. For example, if the increase in nominal rate is 25%, it would mean a security camera would be attacking with 200-250 packets/sec on top of its normal operation data rate. Hence, in the 25%-increase case, which we call a low-rate DDoS scenario, a total of 64 devices would be sufficient to successfully attack a small scale server. Therefore, by compromising sufficient number of devices attackers may perform "stealth" low-rate DDoS attacks which are much more challenging to detect.

For a complete solution to a widely-distributed DDoS attack, we propose that the IDS runs at local nodes, such as routers, by monitoring the data traffic from the connected IoT devices and blocking them whenever necessary. We consider one such node in a testbed setup as shown in Figure 4, which observes at each time t the number of packets $\mathbf{x}_t = [x_t^1, \dots, x_t^{15}]$ received from the connected 15 devices. Specifically, the testbed consists of 15 popular IoT devices such as an Amazon Echo Show, a Raspberry Pi, a security camera, smart switches and a few NodeMCUs. We implemented a low-rate HTTP flooding attack in which there was a slight (25%) increase in the number of *GET* and *POST* requests from the two attacked devices. The data was captured in pcap format by using Wireshark and is publicly available along with a detailed description of the attack scenario and a real-time demo video of the proposed IDS ¹.

In Fig. 5, we compare the performance of RAPID with a state-of-the-art IDS for low-rate DDoS [23], which is based on an information metric. Since RAPID jointly monitors data rates from all devices, it easily detects the attack with a small detection delay and false alarm rate. Although the information metric based IDS also detects the attack with a reasonable delay, it needs a considerably longer time due to its window-based operation. In Fig. 6, we compute the mitigation statistic as shown in (5) for each device. We see

¹The video of the testbed setup is available at <https://www.youtube.com/watch?v=sQPDZ98mUGw>




Figure 4: Testbed setup consisting of IoT devices such as smart plugs, Amazon Echo Show, laptops and computers.

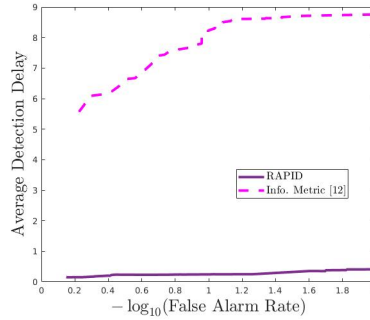


Figure 5: Comparison between RAPID and Information Metrics for a low rate HTTP attack.




Figure 6: Mitigation Statistic computed for the low-rate HTTP attack. Device 2 and Device 6 correspond to the attacked devices.

that there is an obvious increase in the statistic for device 2 and 6, which are indeed the attacked devices.

5 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 a "stealth" low-rate DDoS attack. However, there are still some limitations which need to be addressed to make the system more robust to attacks in the future. First, the current system assumes the network to be static, i.e. the number of devices connected to a node do not change, which might not always be the case. A dynamic scenario needs to be considered in which a new device might also join the network. Secondly, 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, in a real system implementation the IDS needs to be updated periodically.

6 CONCLUSION

With the proliferation of IoT devices, and the ease of initiating cyber-attacks through these devices, there is an increasing need for developing effective solutions for IoT-empowered cyber-attacks. We considered two different scenarios for these attacks: high-rate DDoS attacks and low-rate DDoS attacks. In this context, we have presented a novel intrusion detection and prevention system (IDPS) called RAPID that employs a scalable, online, and nonparametric anomaly detection and localization algorithm. We have illustrated the quick and accurate detection and mitigation performance of the proposed IDS through a real dataset and simulations.

REFERENCES

- [1] Alan Agresti. 2018. *An introduction to categorical data analysis*. Wiley.
- [2] Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J Alex Halderman, Luca Invernizzi, Michalis Kallitsis, et al. 2017. Understanding the mirai botnet. In *USENIX Security Symposium*. 1092–1110.
- [3] Monya Baker. 2016. Statisticians issue warning over misuse of P values. *Nature News* 531, 7593 (2016), 151.
- [4] Markus M Breunig, Hans-Peter Kriegel, Raymond T Ng, and Jörg Sander. 2000. LOF: identifying density-based local outliers. In *ACM sigmod record*, Vol. 29. ACM, 93–104.
- [5] Zhaomin Chen, Chai Kiat Yeo, Bu Sung Lee, and Chiew Tong Lau. 2018. Power spectrum entropy based detection and mitigation of low-rate DoS attacks. *Computer Networks* 136 (2018), 80–94.
- [6] Ping Du and Shunji Abe. 2008. IP packet size entropy-based scheme for detection of DoS/DDoS attacks. *IEICE transactions on information and systems* 91, 5 (2008), 1274–1281.
- [7] Ping Du and Shunji Abe. 2008. IP packet size entropy-based scheme for detection of DoS/DDoS attacks. *IEICE transactions on information and systems* 91, 5 (2008), 1274–1281.
- [8] Earlene Fernandes, Jaeyeon Jung, and Atul Prakash. 2016. Security analysis of emerging smart home applications. In *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, 636–654.
- [9] Alfred O Hero. 2007. Geometric entropy minimization (GEM) for anomaly detection and localization. In *Advances in Neural Information Processing Systems*. 585–592.
- [10] Ben Herzberg, Dima Bekerman, and Igal Zeifman. 2016. Breaking down mirai: An IoT DDoS botnet analysis. *Incapsula Blog, Bots and DDoS, Security* (2016).
- [11] Constantinos Kolias, Georgios Kambourakis, Angelos Stavrou, and Jeffrey Voas. 2017. DDoS in the IoT: Mirai and other botnets. *Computer* 50, 7 (2017), 80–84.
- [12] Lersak Limwiwatkul and Arnon Rungswang. 2004. Distributed denial of service detection using TCP/IP header and traffic measurement analysis. In *Communications and Information Technology, 2004. ISCIT 2004. IEEE International Symposium on*, Vol. 1. IEEE, 605–610.
- [13] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. 2008. Isolation forest. In *2008 Eighth IEEE International Conference on Data Mining*. IEEE, 413–422.
- [14] Yair Meidan, Michael Bohadana, Yael Mathov, Yisroel Mirsky, Asaf Shabtai, Dominik Breitenbacher, and Yuval Elovici. 2018. N-BaLoT: Network-Based Detection of IoT Botnet Attacks Using Deep Autoencoders. *IEEE Pervasive Computing* 17, 3 (2018), 12–22.
- [15] David Moore, Colleen Shannon, Douglas J Brown, Geoffrey M Voelker, and Stefan Savage. 2006. Inferring internet denial-of-service activity. *ACM Transactions on Computer Systems (TOCS)* 24, 2 (2006), 115–139.
- [16] Marius Muja and David G Lowe. 2014. Scalable nearest neighbor algorithms for high dimensional data. *IEEE Transactions on Pattern Analysis & Machine Intelligence* 11 (2014), 2227–2240.
- [17] Muhammad Naveed, Xiao-yong Zhou, Soteris Demetriou, XiaoFeng Wang, and Carl A Gunter. 2014. Inside Job: Understanding and Mitigating the Threat of External Device Mis-Binding on Android. In *NDSS*.
- [18] Eyal Ronen, Adi Shamir, Achi-Or Weingarten, and Colin O’A’Flynn. 2017. IoT goes nuclear: Creating a ZigBee chain reaction. In *Security and Privacy (SP), 2017 IEEE Symposium on*. IEEE, 195–212.
- [19] Bernhard Scholkopf and Alexander J Smola. 2001. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press.
- [20] Saleh Soltan, Prateek Mittal, and H Vincent Poor. 2018. BlackIoT: IoT Botnet of high wattage devices can disrupt the power grid. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*. 15–32.
- [21] Rob Van Der Meulen. 2015. Gartner says 6.4 billion connected things will be in use in 2016, up 30 percent from 2015. *Press Release*. Online unter: <http://www.gartner.com/newsroom/id/3165317> (Abgerufen am 20.10. 2016) (2015).
- [22] Zhi-jun Wu, Liyuan Zhang, and Meng Yue. 2016. Low-rate DoS attacks detection based on network multifractal. *IEEE Transactions on Dependable and Secure Computing* 1 (2016), 1–1.
- [23] Yang Xiang, Ke Li, and Wanlei Zhou. 2011. Low-rate DDoS attacks detection and traceback by using new information metrics. *IEEE transactions on information forensics and security* 6, 2 (2011), 426–437.