

Attribution of Economic Denial of Sustainability Attacks in Public Clouds

Mohammad Karami^{1*}, An Wang² and Songqing Chen²

¹Google, *work performed while at George Mason University

²Department of Computer Science, George Mason University

Abstract

The cloud pricing model leaves cloud consumers vulnerable to Economic Denial of Sustainability (EDoS) attacks. In this type of attacks, an adversary first identifies web resources with high levels of cloud resource consumption, and then uses a botnet of compromised hosts to make fraudulent requests to these costly web resources. The attacker's goal is to disrupt the economical sustainability of the victim by inflicting cost through fraudulent consumption of billable cloud resources.

In this paper, we propose two different Markov-based models to profile the behavior of legitimate users in terms of their resource consumption and the resource request patterns to detect malicious sources engaged in fraudulent use of cloud resources. Our experimental evaluation results demonstrate the effectiveness of the proposed attribution methodology for identifying malicious sources participating in EDoS attacks.

Received on 28 December 2017; accepted on 28 December 2017; published on 28 December 2017

Keywords: Economic Denial of Sustainability, EDoS Detection, Markov Chain, Hidden semi Markov Model

Copyright © 2017 Mohammad Karami *et al.*, licensed to EAI. This is an open access article distributed under the terms of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/3.0/>), which permits unlimited use, distribution and reproduction in any medium so long as the original work is properly cited.

doi:10.4108/eai.4-1-2018.153525

1. Introduction

As a new paradigm, cloud computing is reshaping the entire information technology industry. Cloud service providers enable their consumers to access shared computing resources in a flexible way without the need for upfront investment on infrastructure, platform, and software. Although the adoption of cloud computing has experienced significant growth in recent years, some concerns regarding the unique features of cloud computing environments have hindered its broader adoption. Security and privacy concerns in particular are frequently ranked as one of the top reasons why some organizations are reluctant to adopt cloud computing [5, 26, 27].

The understanding and mitigation of security and privacy risks of the public cloud computing model has been an active area of research in recent years. The research efforts, however, have been primarily focused on protecting the confidentiality and integrity of sensitive data processed in public cloud environments as well as ensuring the continuous availability of cloud services for their intended users [23]. Very little attention has been paid to security threats targeting the

cost model of consumers running their services on the public cloud [11].

Services running on public clouds are vulnerable to fraudulent resource consumption attacks aiming at increasing the financial burden of the victim service. This is enabled by exploiting the utility-based pricing model of the cloud where consumers are charged for the actual consumption of computing resources such as CPU cycles, RAM, bandwidth, and storage [14].

An adversary can conveniently rent a botnet [24] consisting of thousands of bot machines to incur artificial cost to a victim service. The target of the attack will have to pay for the cost of fraudulent resource consumption resulted from requests made by bot clients. By keeping the rate of fraudulent requests made by individual bots low to mimic the behavior of legitimate users, and intelligently focusing on requests that are most costly in terms of resource consumption, an attacker can sustain the attack over an extended period of time and maximize the effectiveness of the attack.

In practice, any device with an Internet connection is capable of launching an EDoS attack. The attacker

can simply instrument the device to send HTTP GET requests to the victim service at the highest rate possible. This is basically the method used in application layer DDoS attacks where the attacker's goal is to render a targeted service unavailable to its intended users by overwhelming victim's resources. However, this will very quickly result in a significant deviation from the request rate of normal users and this artifact can be used for detecting and blocking the offending source [6, 15, 30].

In this paper, we focus on an adversarial scenario in which the attacker's goal is to increase the financial burden of the victim. This attack is also referred to as Fraudulent Resource Consumption (FRC) attack by some researchers in the literature [11, 14]. In a recent empirical study, Wang et al. [29] show how practical EDoS attacks can be launched by abusing popular third-party services provided by companies such as Google, Facebook, etc.

In this paper, we assume that the attacker is intelligent in the sense that she makes requests that are resource-intensive resulting in higher costs for the victim. To be effective, an EDoS attack needs to be stealthy and remain undetected for an extended period of time (e.g., weeks or months). To this end, not only that malicious requests must not cause any noticeable degradation of service quality, but also the quantity of requests made by malicious sources should not be significantly different from those of legitimate users. Although high-rate DDoS attacks with the intention of overwhelming resources of a victim hosted on a public cloud can increase the resource consumption cost for the victim, in our threat model we assume that targets are properly protected against such attacks and we instead only focus on addressing low-rate and stealth EDoS attacks.

As malicious clients participating in a stealth EDoS attack make requests in a similar rate as legitimate users, this type of attacks can be challenging to detect and mitigate. In this paper, we present a method for detecting stealth EDoS attackers by directly assigning a cost to each user request in proportion to the resources consumed to serve that request.

The proposed methodology is based on statistical anomaly detection. First, we process web server logs to identify the sequence of requests made by each individual user over a predefined period of time. Next, according to the amount of resources consumed to serve each request or the time intervals between service requests, a relative cost value is assigned to each request. The result is a dataset consisting of a sequence of request values for each of the legitimate users in the processed web access logs. The sequence of request values for each user is considered as a random or stochastic process and this data is used to construct two different Markov-based models to capture the behavior

of users in terms of the cost they incur to a service over time. We use sequence of request values collected for normal users as training data to estimate the model parameters. Once the parameters are estimated, at the detection phase, the abnormality of a newly observed sequence of request values is tested against the trained model to identify malicious sources participating in an EDoS attack.

We use real-world web access logs of about a month from an academic website to experimentally evaluate the effectiveness of the proposed method. Experimental results are presented for the two Markov-based detection methods that we propose, a simple Markov chain model, and a more complex Hidden semi-Markov Model (HsMM). The experimental results show that our proposed detection methods are very effective in differentiating normal users and malicious users participating in EDoS attacks. While most of previously proposed methods require a malicious source to make significantly more requests than legitimate users to be effective, our proposed attribution methodology can successfully detect malicious sources that try to remain undetected by making only a few resource-intensive or time-periodic requests.

The remainder of this paper is structured as follows. We begin with a discussion on the exploitation of the cloud pricing model that motivates this work. Related work is discussed in section 3. Section 4 presents a brief background on Markov chains and HsMM as well as our proposed Markov-based methods for identifying malicious sources participating in EDoS attacks. The details of experiments designed to validate the proposed methodology and their results are presented in section 5. Finally, discussion and conclusion remarks are presented in chapters 6, and 7, respectively.

2. Exploitation of the Utility-based Pricing Model

The cloud computing technology provides many attractive benefits such as avoiding the need for upfront spendings on computing infrastructure, improved manageability, security, and elasticity to businesses of various sizes. While the flexibility of the "pay-as-you-go" pricing model adopted by cloud service providers can be beneficial to cloud consumers, it leaves them vulnerable to financial risks imposed by EDoS attacks [11, 14].

To launch an EDoS attack, all an attacker needs to do is to simply send seemingly legitimate requests to a victim service to make it consume cloud resources for which the victim will have to pay for the cost. If the attacker is able to enforce significant fraudulent resource consumption over an extended period of time, the economical sustainability of the victim service could be threatened.

In an EDoS attack, the attack target can be a website or web applications hosted on a third party public cloud and we assume that attack targets predominantly serve public content accessible to all Internet users.

Unlike Distributed Denial of Service (DDoS) attacks, an EDoS attack is not meant to cause availability issues or noticeable degradation of service quality for the users of a target service. To be effective, an EDoS attack needs to be stealthy and remain undetected over an extended period of time (e.g., weeks or months). To remain undetected, a wise attacker will want to keep the rate of fraudulent requests low to blend them into the noise of legitimate requests, while trying to focus on requests resulting in high levels of cloud resource consumption to achieve the objective of the attack.

As documented in recent studies, DDoS-for-hire services can be readily located and rented on underground black markets [4, 16, 17]. These abusive services are often supported by botnets consisting of tens of thousands of compromised hosts and offer both network layer and application layer attacks [28]. With the availability of DDoS-for-hire services, an attacker does not need to be capable of building a supporting attack infrastructure.

The potential impact of an EDoS attack can be best quantified by examining a hypothetical attack on a service hosted on a real public cloud service provider. In the sequel, we consider a hypothetical attack on a victim service hosted on Amazon’s Elastic Compute Cloud (EC2) platform. Although cloud consumers are billed for various cloud resources including computing, network, and storage resources, for simplicity, this work only focuses on data transferred from the cloud environment to the Internet to serve received requests. Table 1 shows the cost of outgoing data transfer for Amazon’s EC2 platform [3].

Table 1. Amazon EC2 outgoing data transfer pricing as of February 2016.

Traffic volume	Cost
First 1 GB / month	\$0.00 Per GB
Up to 10 TB / month	\$0.09 Per GB
Next 40 TB / month	\$0.085 Per GB
Next 100 TB / month	\$0.07 Per GB
Next 350 TB / month	\$0.05 Per GB

According to the HTTP Archive [2], which regularly measures the Alexa top 10,000 websites [1], the average page size was 2,225 KB for the homepage of the top 10,000 websites visited in January 2016. However, many websites host a number of much larger web resources such as videos or large compressed files that an attacker can focus on to maximize the cost of resource consumption for a victim operating on a public cloud. For the purpose of our hypothetical EDoS attack, we assume the average size of web resources requested

by malicious bots participating in the attack to be 100 MB.

At the rate of only 100 requests per month which is too low to raise any red flags, a single bot would consume about 10 GB of outgoing bandwidth and the monthly bill will increase by 90¢. Sending requests with the same characteristics as the single bot scenario from 1000 bots will approximately cost the victim \$900 per month. The inflicted cost grows linearly by increasing the request rate, requesting larger files, or employing more malicious bots. For instance, by locating and requesting files that are 1 GB in size, the fraudulent resource consumption cost of the previous EDoS attack scenario will escalate to about \$9000 per month. As seen from the hypothetical attack scenarios, the resource consumption cost accumulated over time can impose an important financial burden to public cloud consumers, especially small businesses. As individual bots show no trace of excessive request rates, most of existing detection schemes that look for a large number of requests in a short period of time [15, 22] will not succeed at detecting the described hypothetical attack.

It is worth noting that leasing a botnet to carry out an EDoS attack will be a cost factor that an attacker would need to take into consideration. However, due to the fact that only a very small fraction of resources available to a compromised host are actually required to make a few requests at a very low rate, the cost of accessing a botnet can be significantly reduced for an attacker by renting non-dedicated botnets shared with other cybercriminals using the bots for various purposes. According to Huang et al. [10], using the pay-per-install marketplace, an attacker can gain access to 1000 compromised machines for as low as \$10.

3. Related Work

So far there are only a few studies in the literature directly concerning the issue of EDoS attacks.

Khor and Nakao [18] propose a mitigation mechanism based on cryptographic puzzles to dissuade clients from submitting fraudulent requests. The basic idea of their proposed scheme called self-verifying Proof of Work (sPoW) is to require clients to present a proof of work before a protected service commits its resources to serve clients’ requests. When a client first requests a resource, it receives a “crypto-puzzle” from sPoW that mediates all communications between clients and the protected service. The puzzle contains encrypted information necessary to reach the intended service such as the IP address and port number as well as a partial encryption key with k bits concealed. The client will have to spend its resources to discover the encryption key by brute forcing the k concealed bits so that it can decrypt the information necessary to contact the requested service.

However, sPoW or any other solution based on the “crypto-puzzle” approach [21] are more relevant when malicious sources are sending requests at a high rate to a target service. In an intelligent and stealth EDoS attack, malicious clients can afford to solve the puzzles to submit only a few well-crafted, resource-intensive requests and succeed at adding financial burden to a victim service protected by sPoW.

Sqalli et al. propose a mitigation scheme called EDoS-Shield to address the issue of EDoS attacks in cloud environments [25]. The main idea of EDoS-Shield is to detect whether an incoming request is initiated by a legitimate user or by an automated source. EDoS-Shield depends on CAPTCHA tests to verify the source of requests. The proposed architecture consists of virtual firewalls (VF) and verifier nodes (V-Nodes) that are deployed as virtual machines in the cloud. The V-Nodes are responsible for verification of request sources, and VF nodes are implemented to decide if incoming packets should be forwarded or dropped based on the verification results received from the V-nodes. One weakness of the EDoS-Shield mitigation scheme has to do with the cost of additional cloud resources required for deploying the verifier nodes and the virtual firewalls. But, more importantly, this approach requires all users to be verified and research studies suggest that CAPTCHA tests could be annoying for some users and even a certain portion of legitimate users may not be able to solve them [7]. In addition, some existing CAPTCHA tests have been shown to be vulnerable to automated attacks [8], and recently inexpensive CAPTCHA solving services backed by crowd sourced human labor can be used to effectively defeat the protection purpose of CAPTCHA tests [20].

In [12] the authors use a number of statistical self-similarity metrics including Zipf’s law, and Spearman’s Footrule distance to detect the occurrence of FRC attacks. The proposed detection mechanism only looks at the aggregate pattern of user requests and does not deal with identification of individual malicious sources participating in an attack. In contrast, our proposed method is concerned with identification of malicious sources exhibiting a similar behavior as legitimate users in terms of request rates, but focusing on resource-intensive requests to maximize the cost for the victim service.

Idziorek and Tannian propose a method that attempts to model the behavior of individual users based on the number of requests per session generated by each user over a fixed period of time [11]. A pause of 900 or more seconds between consecutive requests from the same user is used as the criterion to group user requests into web sessions. The premise is that malicious users generating sessions with a random number of requests would be sufficiently different from the profile of normal users, so that an entropy-based

detection method could be used to identify malicious sources. This method is based on the assumption of malicious users making more requests/web sessions than legitimate users. However, as mentioned earlier, an intelligent attacker does not necessarily need to make malicious sources to send more requests than legitimate users to succeed. By focusing on web resources that are expensive in terms of resource consumption, malicious sources with similar request rates as legitimate users can be still effective.

In [13], the authors propose a methodology for identifying malicious sources trying to inflate the utility bill of a victim by making fraudulent requests. The proposed methodology combines four different usage metrics including the number of sessions, the number of requests, and the average number of requests per session. For the last usage metric, the overall request frequency distribution of documents hosted on a website is computed, and the requests made by individual users are compared against this distribution. To evaluate a user, a probability score is calculated for each of the four metrics and then an overall average probability is computed. The more deviation observed from the normal usage, the higher would be the probability score and the more likely the user would be a malicious client. Again, this model is heavily influenced by the the quantity of requests made by individual users, and it will not be effective for detecting malicious users making a small number of high cost requests. As we will show in section 5, our proposed method is able to detect both malicious sources making an anomalous number of random requests, as well as more subtle malicious sources with a request rate similar to that of legitimate users but focusing on requests that are more costly for the victim.

4. The Proposed Markov-based Models for Detecting Sources Participating in an EDoS Attack

In this study, our goal is to build an anomaly detection system to identify malicious sources participating in EDoS attacks. In this section we introduce our proposed detection methodology and describe our formulation of detecting malicious sources participating in an EDoS attack using two different Markov-based models.

Markov-based models are usually used to model randomly changing systems. It is also assumed that future states depend only on the current state. Such statistical process resembles the process of browsing web pages. Furthermore, since the resources could be categorized into different groups based on their size, it is natural to represent the Markov-based models with a finite state machine.

Most web requests are for HTML documents that are meant to be rendered, and displayed by a user browser. These requests are typically followed shortly by several

subsequent HTTP GET requests to fetch objects such as images, scripts, and CSS files embedded in the main requested document. The requests can also be for downloading objects such as binary files over HTTP. Web servers can be configured to log the details of all user requests including the IP address of the requesting host, the requested document, the type of request (GET, POST, etc.), and the size of data transferred to serve the request. Although all the HTTP request types cause resource consumption on the server side, to simplify our experimentations, we only focus on HTTP GET requests in our work.

Proportional to the amount of data transferred to serve a request, a relative cost value can be calculated and assigned to each request. Based on the data size of various requests, one can decide on a small number of buckets to represent different cost values to be associated with user requests. Furthermore, the time intervals between multiple requests in a sequence capture the patterns that describe user behaviors. These values could also be bucketized and represented by discontinuous values. We will see examples of both in section 5 where we use discrete values from 1 to 5 for requests in our dataset.

Using collected web server logs, requests made by each individual user during a specific period of time can be identified and mapped to the discrete request values. The result would be a sequence of the values for each user. We assume that individual users (both legitimate and malicious) can be uniquely identified by their IP addresses. Using browser fingerprinting techniques [9] can be a potential solution for cases where some users can not be reliably identified by their IP addresses.

The sequence of request values from individual users during a specific period of time can be considered as a discrete-time stochastic process and a Markov-based model can be used to describe the behavior of users in terms of the cost they incur to a service or the frequency they request a service over time. Intuitively, it is trickier to profile users based on their request frequency because EDos attacks do not demand high traffic rate that are distinguishable. A much simpler approach to distinguish between legitimate and malicious users would be to calculate the sum of request values per user over a predefined period of time and apply a threshold value to identify users exceeding the threshold as malicious. However, as we will show in section 5, such a naive approach will result in high positive rates where legitimate users are incorrectly identified as malicious.

We use requests made by legitimate users to estimate the parameters of the Markov model and then use the trained model to compute the likelihood of new request value sequences generated by users. The request value sequences generated by malicious users would be different from legitimate users and this will result in much smaller likelihood values than those of

legitimate users. As we will show in section 5, using the right threshold likelihood value, legitimate users and malicious users can be effectively distinguished.

In this paper, we propose and evaluate the detection performance for two different Markov-based models. The first one is a simple Markov chain model in which the observed request values are considered as the states of the Markov chain. We also evaluate a HsMM which has more complexity and computational cost but can outperform the simple Markov chain model for detecting low rate attacks focusing on high cost or high frequency requests. In our HsMM, the request values are the observable outputs and the hidden states represent different levels of resource consumption by users. The models that we propose and evaluate are both discrete-time and the discrete points in time correspond to user requests as recorded in the web access logs. In the subsequent subsections, we give a brief background on the theory of Markov chains and HsMM and briefly discuss the process of learning model parameters using training data.

4.1. Markov Chain Model

A Markov chain models the state of a system with a random variable taking states from a finite state space as the time passes. Given the current state of a Markov chain denoted as s_t , the state of the chain at time $t + 1$ will only depend on s_t . In a stationary Markov chain, the state transition probabilities are assumed to be constant and independent of time. Consider a Markov chain model with M states denoted as $S = \{s_1, s_2, \dots, s_M\}$. A Markov chain model can then be specified by its parameters as $\lambda = (\{\pi_m\}, \{a_{mn}\})$ where:

- $\pi_m \equiv \Pr[s_1 = m]$ is the initial state probability distribution. s_t denotes the state taken by the model at time t and $m \in S$. The sum of initial state probabilities adds up to 1 ($\sum_m \pi_m = 1$).
- $a_{mn} \equiv \Pr[s_t = n | s_{t-1} = m]$ is the state transition probability for $m, n \in S$, satisfying $\sum_n a_{mn} = 1$.

Given the model parameters, the probability of a particular sequence of states s_1, s_2, \dots, s_T to be taken by the model is computed as follows:

$$P(s_1, s_2, \dots, s_T) = \pi_{s_1} \sum_{t=2}^T a_{s_{t-1}s_t}$$

The initial state probability distribution and the state transition probability matrix can be readily learned from historical observations of the system states. These two model parameters can be learned using the formulas below [19]:

$$a_{mn} = \frac{N_{mn}}{N_m}$$

$$\pi_m = \frac{N_m}{N}$$

where:

- N_{mn} is the number of observed direct transitions from state m to state n .
- N_m is the number of observations where the Markov chain is in state m .
- N is the total number of observations.

In our context, we process web access logs to compute the sequence of request values for normal users and then use this data to learn the parameters of a Markov chain model representing the resource consumption behavior of normal users. At the detection phase, newly observed sequences of request values are analyzed to compute the likelihood of those sequences being supported by the trained Markov chain model. The resource consumption behavior of users participating in EDoS attacks would be different from that of normal users and the requests from these users are therefore expected to receive low likelihood of support when analyzed by the trained model.

4.2. Hidden Semi-Markov Model

A hidden Markov model (HMM) is a Markov model in which the states of the system being modeled are not directly observable (hidden). HsMM extends the traditional HMM by allowing states to have variable durations [31]. The duration of a state represents the number of observations made while in that state. Consider a HsMM with M states denoted as $S = \{s_1, s_2, \dots, s_M\}$. A HsMM can be specified by its parameters as $\lambda = (\{\pi_m\}, \{a_{mn}\}, \{b_m(k)\}, \{p_m(d)\})$ where:

- $\pi_m \equiv \Pr[s_1 = m]$ is the initial state probability distribution. s_t denotes the state taken by the model at time t and $m \in S$. The sum of initial state probabilities adds up to 1 ($\sum_m \pi_m = 1$).
- $a_{mn} \equiv \Pr[s_t = n | s_{t-1} = m]$ is the state transition probability for $m, n \in S$, satisfying $\sum_n a_{mn} = 1$.
- $b_m(k) \equiv \Pr[o_t = k | s_t = m]$, for $m \in S, k \in \{1, \dots, K\}$ is the state output distribution. The observable output at t is denoted by o_t and k is the index into the observable output set with cardinality K . The output distribution satisfies $\sum_k b_m(k) = 1$.
- $p_m(d) \equiv \Pr[\tau_t = d | s_t = m]$ is the state residual time distribution, for $m \in S, d \in \{1, \dots, D\}$. D represents the maximum interval between any consecutive state transitions and the residual time distribution satisfies $\sum_d p_m(d) = 1$.

Then, if at time t , the pair process (s_t, τ_t) takes on the value (m, d) , where $d \geq 1$, the semi-markov chain will remain in state m until time $t + d - 1$ and will transit to the next state at time $t + d$. The states themselves are not directly observable. The observables are a sequence of observations $O = (o_1, \dots, o_T)$. The notation o_a^b represents the observation sequence from time a to time b and conditional independence of observed outputs is assumed so that $b_m(o_a^b) = \prod_{t=a}^b b_m(o_t)$. The model parameters are initially estimated and are then updated as new observations o_t are collected. This process is known as parameter reestimation and it can be done by following the forward and backward algorithm proposed by Yu and Kobayashi [32]. The forward and backward variables are defined as follows:

$$\alpha_t(m, d) \equiv \Pr[o_1^t, (s_t, \tau_t) = (m, d) | \lambda]$$

$$\beta_t(m, d) \equiv \Pr[o_{t+1}^T, (s_t, \tau_t) = (m, d) | \lambda]$$

which can be recursively computed by forward and backward algorithms. Next, the three following joint probabilities are defined that can be expressed and computed in terms of the model parameters and the forward and backward variables defined above. These probabilities are used to readily derive the reestimation formulas to update the model parameters after collecting new observation sequences.

$$\zeta_t(m, n) \equiv \Pr[o_1^T, s_{t-1} = m, s_t = n | \lambda]$$

$$\eta_t(m, d) \equiv \Pr[o_1^T, s_{t-1} \neq m, s_t = m, \tau_t = d | \lambda]$$

$$\gamma_t(m) \equiv \Pr[o_1^T, s_t = m | \lambda]$$

Now, using the joint probabilities defined above, the model parameters can be reestimated by the following formulas:

$$\hat{\pi}_m = \gamma_1(m) \Big/ \sum_{m=1}^M \gamma_1(m)$$

$$\hat{a}_{mn} = \sum_{t=1}^T \zeta_t(m, n) \Big/ \sum_{t=1}^T \sum_{n=1}^M \zeta_t(m, n)$$

$$\hat{b}_m(k) = \sum_{t: o_t=k} \gamma_t(m) \Big/ \sum_k \sum_{t: o_t=k} \gamma_t(m)$$

$$\hat{p}_m(d) = \sum_{t=1}^T \eta_t(m, d) \Big/ \sum_{t=1}^T \sum_{d=1}^D \eta_t(m, d)$$

The model parameters are reestimated for each observation and after processing all observation sequences, the trained model can be used to compute

the likelihood of a new observation sequence by the following formula:

$$\begin{aligned} \Pr[o_1^T | \lambda] &= \sum_m \sum_d \Pr[o_1^T, (s_T, \tau_T) = (m, d) | \lambda] \\ &= \sum_m \sum_d \alpha_T(m, d) \end{aligned}$$

For our HsMM, the request values are the observable outputs and the hidden states represent different levels of resource consumption by users. In our implemented model, we use 5 hidden states where the model is always initialized in the first state. Also, in our model a transition can only happen from a lower state to a higher state.

Similar to the Markov chain model, we use requests made by legitimate users to estimate the parameters of the HsMM and then use the trained model to compute the likelihood of new request value sequences generated by users. The request value sequences generated by malicious users would be different from legitimate users and this will result in much lower likelihood values than those of legitimate users. As we will show in the next section, using the right threshold likelihood value, legitimate users and malicious users can be effectively distinguished.

5. Experimental Evaluation

We conduct experiments to evaluate the effectiveness of the proposed method for detecting malicious sources engaged in fraudulent use of cloud resources. This section provides a description of our experiments and presents the obtained results.

5.1. Dataset Description

Our experiments are based on request logs from a university department's public web server collected over 32 days from Nov 8, 2015 to Dec 9, 2015. We use the rules below to filter out requests that are irrelevant for our purpose:

- Requests that are not HTTP GET.
- HTTP GET requests with a response code other than 200 (OK).
- Requests with a user agent string indicating access from a non-user entity (e.g., Googlebot, wget, etc.).
- Request sources making requests using 10 or more different user agent strings. This is to remove aggregate request sources such as NAT boxes or web proxies making requests on behalf of their clients. About 90% of all request sources in the dataset only use a single user agent string.

- The access logs are split into two 16-day periods. We only include requests from users making at least 3 requests in one of the 16-day periods.

A request for an HTML document and the subsequent requests for fetching objects embedded in the same HTML document are combined and treated as a single request. Table 2 presents a summary of the normal dataset used for training and testing the proposed methodology.

Table 2. Summary of the normal experimental dataset.

Metric	Train dataset	Normal test dataset
Number of days	16	16
Total number of unique users	4,933	5,252
Total number of requests	36,466	36,474
Avg number of requests per user	7.39	6.94

To generate the normal training dataset, requests in the first half of the logs are grouped based on the request source, and then they are mapped to relative request values. The same process is applied to the requests in the second period of the logs to generate the test dataset representing users with normal resource usage behavior. Based on our observation of the user requests in the dataset, we choose to use the values from 1 to 5 to represent the relative values of user requests. Thus, the final dataset is a sequence of request values ranging from 1 to 5 for each user. Table 3 summarizes the mapping from the request size and the request intervals to relative values and the distribution of requests in terms of their values in our dataset. After processing, both the training data and the testing data could be represented by an array of requests as shown by the following examples:

```
11547,1,1,1,1,1,1,1,1,1,1,1,1
11548,1,1,1,2,1,1,1,1,2,2,1
36606,1,1,1,1
```

Each record starts with a unique id that differentiates request sources following a sequence of values that represents the resource categories. These records are used as inputs to obtain the estimated transition probabilities between different states of the Markov models.

Table 3. Mapping of request sizes to relative request values.

Request value	Request size	Percentage	Request interval	Percentage
1	< 500KB	86.7	< 20min	60.4
2	≥ 500KB and < 5MB	11.4	≥ 20min and < 1hr	10.3
3	≥ 5MB and < 50MB	1.9	≥ 1hr and < 4hr	9.4
4	≥ 50MB and < 500MB	0.1	≥ 4hr and < 12hr	1.1
5	≥ 500MB	5.4e-03	≥ 12hr	8.9

Although in our experiments we use an observation window of 16 days to profile the behavior of normal users and the same observation period is used for detection of malicious users, the proposed methodology

is only sensitive to the resource usage pattern of users and is not restricted to a specific observation period.

To conduct an EDoS attack, an attacker needs to specify the behavior of individual bots by defining the request rate and the requested resources in term of their resource consumption cost. By varying these two parameters, various attack strategies that an attacker is likely to adopt can be constructed and the effectiveness of the proposed detection methodology can be evaluated for those attack strategies.

5.2. Resource-intensive Attack Scenarios

We first consider attack strategies where the attacker focuses on making requests that result in high levels of resource consumption. For these attack strategies we assume that the attacker has a prior knowledge about the rate of requests made by legitimate users, and uses this knowledge to avoid suspicions by making requests with similar rates as legitimate users.

In sequel, we briefly describe a number of various attack scenarios that we use to generate synthetic malicious request sequences to evaluate the performance of the proposed attribution methodology. These attack scenarios are ordered in a decreasing order of attack effectiveness from the attacker's perspective. For the attack scenarios listed below the number of requests made by individual malicious sources is normally distributed with parameters $\mu=7$ and $\sigma=2$ which means the number of requests made by malicious users are not significantly different from those of legitimate users. The lengths of malicious request sequences drawn for this normal distribution is shared by all of the attack scenarios.

- **Scenario 1 (S1):** All malicious requests have a cost of 5.
- **Scenario 2 (S2):** All malicious requests have a cost of 4 or 5.
- **Scenario 3 (S3):** The request cost is 5 for 75% of malicious requests. The cost for the remaining 25% of requests is uniformly distributed between 1 and 4.
- **Scenario 4 (S4):** 75% of malicious requests have a cost of 4 or 5. The cost for the remaining 25% of requests is uniformly distributed between 1 and 3.
- **Scenario 5 (S5):** The request cost is 5 for 50% of malicious requests. The cost for the remaining 50% of requests is uniformly distributed between 1 and 4.
- **Scenario 6 (S6):** 50% of malicious requests have a cost of 4 or 5. The cost for the remaining 50% of

requests is uniformly distributed between 1 and 3.

5.3. Resource-intensive Results

For each attack scenario, we generate a dataset of malicious requests according to the description of that attack scenario. Each test dataset consists of generated malicious request sequences, combined with the request sequences from the normal test dataset. The normal test dataset is the same for all attack scenarios. Also, in our experiments, each test dataset contains the same number of normal and malicious sources (5,252). False Positive Rate (FPR), and False Negative Rate (FNR) are the metrics used for performance evaluation of the proposed attribution methodology under various attack scenarios. These metrics are briefly described in the following:

- **FPR:** The percentage of request sequences generated by legitimate users classified as malicious. Keeping the FPR under a low threshold is very important. Otherwise, legitimate users will be denied access to the protected service.
- **FNR:** The percentage of malicious request sequences generated by sources participating in an EDoS attack scenario not detected by the proposed method. Unlike an Intrusion Detection System (IDS) where it is very crucial not to miss any intrusions, because a single missed intrusion can result in system compromise, in our context, missed malicious sequences would only cause some billable fraudulent resource consumption.

For each attack scenario, the trained model is used for computing the log likelihood for all request sequences in the test dataset of that attack scenario. In general, the request sequences from legitimate users which are expected to receive higher log likelihood values. On the other hand, malicious request sequences representing a resource consumption behavior dissimilar to that of legitimate users are expected to be assigned lower log likelihood values. Once the log likelihoods are computed for all request sequences, the detection performance can be evaluated for different threshold values.

Table 4 shows experimental detection results for the attack scenarios based on the strategy of focusing on high cost requests using the simple Markov chain model. Table 5 shows the results for the same attack scenarios using the HsMM. The results are presented for several different threshold values to demonstrate the trade-off between higher false positive rates and lower false negative rates and vice versa. For each threshold value, the resulting FPR and the FNR for each of the six attack scenarios are presented. As shown, both models

can achieve low FPRs for reasonable FNRs. The HsMM however consistently outperforms the simple Markov chain model for the same FPR value.

Table 4. Experimental results for the Markov chain model for the attack strategy of focusing on high cost requests.

Threshold	FPR(%)	Attack Scenario FNR(%)					
		S1	S2	S3	S4	S5	S6
-20	2.32	0.00	2.55	0.97	5.79	5.14	13.58
-25	1.35	0.00	6.32	2.30	10.68	8.87	22.51
-30	0.91	0.00	7.81	3.77	14.58	13.40	31.51
-35	0.67	7.24	16.85	11.23	25.08	22.37	43.01
-40	0.46	7.24	21.63	13.08	31.70	28.64	52.36
-45	0.29	16.28	28.66	22.43	42.25	37.93	63.06

Table 5. Experimental results for the HsMM for the attack strategy of focusing on high cost requests.

Threshold	FPR(%)	Attack Scenario FNR(%)					
		S1	S2	S3	S4	S5	S6
-15	2.32	0.00	0.00	0.44	1.64	2.76	7.52
-20	1.26	0.00	2.57	2.02	5.62	6.84	15.69
-25	0.67	0.00	2.57	3.56	8.95	10.68	23.67
-30	0.51	7.24	9.04	10.95	16.70	20.13	33.80
-35	0.36	7.24	11.18	15.19	23.38	28.87	45.35
-40	0.21	16.28	18.55	23.59	33.45	38.42	57.60

For instance for a FPR of 0.67%, the FNR is 0.00% for the first attack scenario when the HsMM is used. This means that for a small FPR, all sources generating malicious requests according to the description of the first attack scenario (S1) are successfully detected. In comparison, the simple Markov chain Model produces a FNR of 7.24% for the attack scenario and the same FPR. In our experiments, the last attack scenario (S6) is the most challenging to detect as it is more similar to requests from legitimate users. But even for this attack scenario, for a FPR of 0.51%, still close to 70% of malicious request sequences are successfully detected using the HsMM.

It should be noted that the undetected malicious sequences are usually comprised of fewer requests compared to the successfully detected malicious sequences. For instance, when applying a threshold value of -25 for the attack scenario S6 using the HsMM, about 10% of malicious users are not identified as malicious. The average number of requests for these undetected malicious users is 3.98 versus 7.23 for the detected malicious users. This implies that the undetected malicious request sequences are less effective in terms of fraudulent resource consumption and the more effective malicious request sequences that are more aggressive in nature run higher risk of detection.

In our experiments, the FPRs are resulted by legitimate users in the test dataset generating significantly more requests than the other legitimate users. For instance, when applying a threshold value of -25 for the HsMM, 35 legitimate users out of 5252 users are incorrectly identified as malicious. On average, each

of these 35 users generates 72 requests. In comparison, the overall average number of requests for all legitimate users is only 7. Longer request sequences result in lower log likelihood values and therefore legitimate users generating a large number of requests contribute to some undesirable false positives. Sometimes a given website may have legitimate users that use the website in unusual ways. If unusual request patterns are expected from specific users, false positives can be avoided by ignoring requests from these known users. In our experiments, the legitimate users with an abnormally large number of requests are incorrectly identified as malicious.

Comparison with the naive detection approach: as mentioned previously, a naive detection approach based solely on the cumulative sum of request costs will suffer high positive rates. As a concrete example, to achieve a FNR of 0.00% for the S1 attack scenario, the naive approach will need to identify all users with a cumulative sum of request costs of 15 or more as malicious. This however will result in a very high FPR of 9.96% making this approach inapplicable in practice. In contrast, as shown in Tables 4 and 5, using the proposed attribution methodology, for the S1 attack scenario all malicious users can be successfully detected with FPRs less than 1%.

Table 6. Experimental results for the Markov chain model for the attack strategy of focusing on high number of requests.

Threshold	FPR(%)	Number of requests per source FNR(%)					
		S1	S2	S3	S4	S5	S6
-20	2.32	3.77	0.72	0.00	0.00	0.00	0.00
-25	1.35	17.02	3.41	0.32	0.00	0.00	0.00
-30	0.91	42.19	14.38	3.24	0.32	0.02	0.00
-35	0.67	67.82	34.22	12.20	2.65	0.30	0.02
-40	0.46	86.27	59.12	29.86	10.45	2.21	0.49
-45	0.29	95.00	79.38	51.71	25.46	8.85	2.70

Table 7. Experimental results for the HsMM for the attack strategy of focusing on high number of requests.

Threshold	FPR(%)	Number of requests per source FNR(%)					
		S1	S2	S3	S4	S5	S6
-15	2.32	6.51	2.11	0.76	0.15	0.02	0.00
-20	1.26	23.42	10.66	3.27	0.74	0.23	0.06
-25	0.67	49.58	20.91	9.69	4.23	1.52	0.74
-30	0.51	82.77	51.79	25.32	10.62	4.09	2.04
-35	0.36	94.94	79.59	54.72	24.68	12.32	6.78
-40	0.21	98.19	92.90	79.99	53.98	27.88	13.23

As evidenced by the obtained experimental results, attacks based on the strategy of focusing on requests with high resource consumption costs can not go undetected. The alternative for an attacker would be to attempt making requests with a similar distribution of request costs as legitimate users, but in larger quantities to increase the amount of fraudulently consumed resources.

Table 6 and 7 show experimental detection results for the attack strategy where the attacker focuses on

making larger numbers of requests that have the same distribution of request costs as legitimate users. FNRs are reported for various number of requests per source and the same threshold and FPRs as the previous experiments. As expected, malicious users are more likely to be detected when making higher number of requests. For instance, when malicious clients are making 70 requests, for a FPR of 0.67%, more than 90% of malicious sources are successfully detected by the HsMM. For the experiments involving long sequences of malicious requests, the performance of the two models are comparable and none of them consistently outperforms the other model for the same FPR value. However, when the number of requests per source is 80 or more, the simple Markov chain model seems to produce lower FNRs.

It should be noted that from the standpoint of an EDoS attacker, regular requests not causing high levels of resource consumption are not very helpful and this attack strategy only makes sense when malicious sources are able to make a significant number of regular requests and manage to remain undetected.

5.4. Time-periodic Attack Scenarios

Since EDoS attacks usually do not require high traffic rates, the bots that launch the attacks are programmed to be active from time to time by the bot masters to avoid detection mechanisms while incurring costs. Thus, malicious requests are more likely to appear in periodic patterns than the normal ones, especially over the long term. To present these attack strategies, we describe the following scenarios that we use to generate synthetic malicious request sequences. These scenarios are ordered in an increasing order of randomness in terms of the request intervals. For this group of attack scenarios, we follow the same normal distribution parameters to generate the number of requests for each individual attack sequence. Notice that we only keep the sequences whose lengths are larger than or equal to 10 to improve the efficiency of our model. Also, longer sequences favor the malicious users more from the profit perspective especially when there are only limited number of controlled bots.

For time intervals in range 1 to 5, we further categorize them into 3 different groups: (1) highly frequent interval, which contains 1; (2) medium frequent interval, which contains 2; and (3) intermission intervals, which include range 3 to 5. These terms will be utilized to describe the following scenarios.

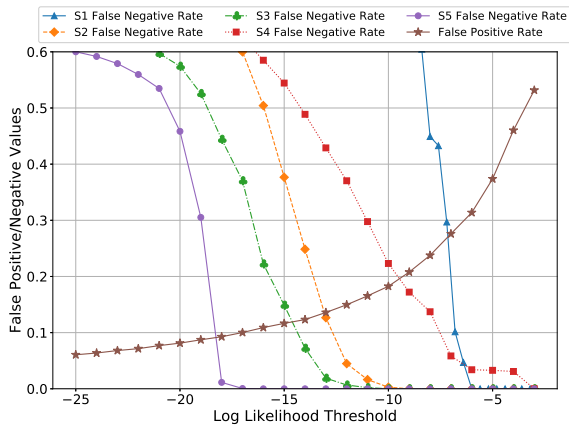
- **Scenario 1 (S1):** The malicious requests are composed of highly frequent (90%) and intermission intervals (10%) and intermission intervals are uniformly selected from range [3, 4, 5]

- **Scenario 2 (S2):** The malicious requests contain highly frequent intervals (75% * 90%), medium frequent intervals (25% * 90%) and intermission intervals (10%).
- **Scenario 3 (S3):** The malicious requests contain highly frequent intervals (50% * 90%), medium frequent intervals (50% * 90%) and intermission intervals (10%).
- **Scenario 4 (S4):** The intervals between malicious requests are generated following the *power lognormal distribution*, with parameters $p = 1.69$ and $\sigma = 3.99$, which is the same with the normal requests.
- **Scenario 5 (S5):** The interval values are randomly selected from range [1 - 5] following a uniform distribution.


5.5. Time-periodic Results

Then, we follow the same procedure to generate the malicious request datasets and use them to evaluate the accuracy of our proposed models with the metrics, FPR and FNR. We present the evaluation results in Fig 1. In both figures, we present the changes of FPR/FNR values (y-axis) with the increase of the log likelihood threshold values (x-axis). As aforementioned, the legitimate users are more likely to be assigned with higher log likelihood values and the malicious users are expected to get lower log likelihood values. Thus, the FNR values decrease and the FPR values increase as the increase of the threshold values. In addition, the intersected points of the FPR and FNR lines in the figure indicate the best performance of the models with the balanced FPR and FNR values.

We could observe that both models could achieve reasonably low FPR and FNR values. By comparing these two figures, we further find that the HsMM model outperforms the simple Markov chain model in all attacking scenarios. Another observation is that among all the scenarios, S1 appears to be the most difficult to be detected since it is closest to behaviors from legitimate users. This could be reasoned based on the fact that we group all time intervals larger than 12 hours in the same category so that there are no ways to distinguish the intermission interval lengths of the malicious requests and the legitimate ones. However, even in such circumstance, the HsMM model could still achieve about 85% accuracy in the best cases. In addition, it is interesting to find that, for scenario S4, the prior knowledge of the distribution of the request intervals of the normal users could not provide much leverage for the attacker. This points out the fact that the arrangement of time intervals between requests matters more than the value themselves.



(a) Experimental results for the Markov chain model for the attack strategy focusing on periodic request patterns.



(b) Experimental results for the HsMM for the attack strategy focusing on periodic request patterns.

Figure 1

One could imagine that the longer the request sequences are, the more likely it is for the periodic patterns to appear. Since the EDoS attacks usually gain more benefits from longer period of activities, it is reasonable to assume that the lengths of malicious request sequences are more likely to be larger. Thus, we further explore the efficiency of the HsMM model with the increasing length of the request sequences. The results are shown in Fig 2.

In this figure, we demonstrate the best FPR/FNR values (y-axis) changes with the increasing minimum length of the request sequences under all attacking scenarios. We could observe from this figure that even for the most challenging scenario S1, the FPR/FNR value would reach as low as less than 6% when the sequence length increases to 50. Another evident observation is that the performance of the model increases almost linearly with the length of the requests. This further demonstrates the power of our proposed

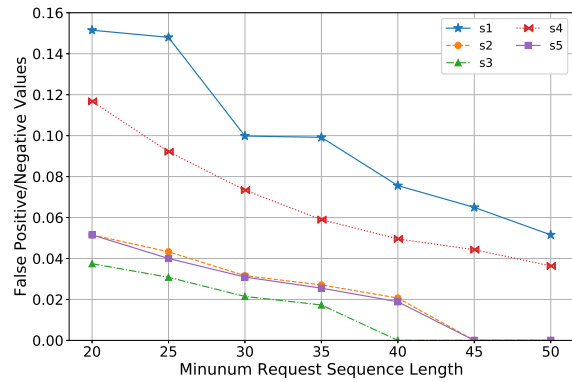


Figure 2. Experimental results for the HsMM for the attack strategy focusing on long-term requests

model in identifying malicious requests over the long term.

Comparison with the resource-intensive based detection approach: to make a fair comparison, we could compare the best performances of both models in Fig 1 with that in Table 4 and Table 5. We find that both models generate better detection results with the resource-intensive based approach when the malicious sequences are short. This indicates that the time-periodic patterns reveal as the time passes. From the attackers' perspectives, there are no economic incentives to generate only a few requests. Thus, the time-periodic based approaches could perform effectively in real-world attacking scenarios.

Another observation we have from comparing the results is that the threshold values for the time-periodic based approach are generally larger than those of the resource-intensive based approach. Also, the changes of the FPR/FNR values seem to be more drastic for the time-periodic based approach in Fig 1. These observations mean that it is practically more difficult to identify malicious behaviors with the time intervals since malicious activities appear to be close to the legitimate ones. However, the time-periodic based detection leads to better results, especially for the more challenging scenarios, as the malicious activities continue as presented by Fig 2.

In reality, the malicious users may apply combined strategies to deceive the detection systems. As a result, a promising direction of the effective detections is to combine these two approaches we proposed to build a comprehensive model to react to the ever-changing attacking schemes.

6. Discussion

The proposed EDoS attribution methodology directly considers the cost of user requests as the metric to

model the behavior of individual users. This makes it very challenging for malicious users involved in an EDoS attack to be effective in terms of fraudulent consumption of billable cloud resources and at the same time managing to remain undetected. As demonstrated experimentally, malicious users exhibiting anomalous resource consumption behavior can be quickly identified and prevented after making only a small number of suspicious requests.

As the attacking activity continues, the periodic patterns of the malicious users will appear to be more distinguishable for the models to achieve better identification accuracy. However, both approaches may fall short in handling the cold-start problems where little request logs are available. In general, it is a challenging problem since the attackers could emulate normal user behaviors to avoid being detected.

An attacker can attempt to optimize the requests made by individual participating bots by learning and mimicking the request pattern of top legitimate users in terms of higher usage footprint. However, it is unlikely for an attacker to be able to access historical data on requests of legitimate users or intercept communications to collect such data to optimize the behavior of participating bots. Even assuming the lack of such restrictions, applying the time-periodic based approach can still significantly limit the effectiveness of the participating bots. An attacker can try to compensate for the limited utility of individual bots by employing a much larger botnet. However, larger botnets could be very difficult to locate, rent and operate and may not be practical in practice.

The proposed EDoS attribution methodology relies on either resource usage or request interval footprint of users for detecting malicious sources. For future work, we plan to incorporate these two factors to improve the flexibility and detection accuracy of our model. Also, we could include other aspects of user behavior to further improve detection of malicious sources participating in EDoS attacks. For instance, the popularity of requested documents can be computed for each request cost bucket, and this can be considered when computing the likelihood of observed request sequences. In general, attackers are not expected to know the distribution of document popularity on victim websites. Focusing on requests involving documents with the highest resource consumption can result in deviation from the normal document popularity distribution and this additional metric can help to improve the detection performance.

From the defense's perspectives, we could either block the malicious actors by creating a blacklist or restrain the response throughput utilizing token-based mechanisms. The advantages of token-based mechanisms are threefold: (1) favoring toward requests on resource of smaller values, (2) minimizing possibilities of denying normal requests, and (3) customizing and

adjusting throughput limit for different users dynamically. With the advance of network softwarization, it is effortless to deploy such lightweight defense mechanisms in the network routers.

7. Conclusion

The consumers of public cloud services are charged for computing resources that they use. This pricing model exposes the cloud consumers to EDoS attacks where the adversary seeks to increase the financial burden of the victim service by making fraudulent requests that result in high consumption of billable resources.

We have presented a Markov-based anomaly detection scheme to profile the behavior of legitimate users in terms of their resource consumption. To detect users participating in an EDoS attack, the likelihood of request sequences generated by individual clients during a specific period of time is computed by the trained model. Users with likelihood values smaller than a threshold are identified as malicious. The effectiveness of the proposed attribution methodology for identifying malicious sources engaged in fraudulent use of cloud resources has been demonstrated using experimental evaluations for various attack scenarios. While most of previously proposed methods are only effective when malicious sources make significantly more requests than legitimate users, our proposed method is able to detect both malicious sources making an anomalous number of random requests, as well as more subtle malicious sources with a request rate similar to that of legitimate users but focusing on requests that are more costly for the victim.

Acknowledgments. We would like to thank anonymous reviewers for their comments. This work was supported in part by an ARO grant W911NF-15-1-0262 and a NSF grant CNS-1524462.

References

- [1] Alexa. <http://www.alexa.com/>.
- [2] Http archive. <http://httparchive.org/interesting.php?a=A11&l=Jan%2015%202016>.
- [3] Amazon ec2 pricing. <https://aws.amazon.com/ec2/pricing/>, 2016.
- [4] ALOMARI, E., MANICKAM, S., GUPTA, B., KARUPPAYAH, S., AND ALFARIS, R. Botnet-based distributed denial of service (ddos) attacks on web servers: classification and art. *arXiv preprint arXiv:1208.0403* (2012).
- [5] ARMBRUST, M., FOX, A., GRIFFITH, R., JOSEPH, A. D., KATZ, R., KONWINSKI, A., LEE, G., PATTERSON, D., RABKIN, A., STOICA, I., ET AL. A view of cloud computing. *Communications of the ACM* 53, 4 (2010), 50–58.
- [6] BEITOLLAHI, H., AND DECONINCK, G. Tackling application-layer ddos attacks. *Procedia Computer Science* 10 (2012), 432–441.

- [7] BURSZEIN, E., BETHARD, S., FABRY, C., MITCHELL, J. C., AND JURAFSKY, D. How good are humans at solving captchas? a large scale evaluation. In *Security and Privacy (SP), 2010 IEEE Symposium on* (2010), IEEE, pp. 399–413.
- [8] BURSZEIN, E., MARTIN, M., AND MITCHELL, J. Text-based captcha strengths and weaknesses. In *Proceedings of the 18th ACM conference on Computer and communications security* (2011), ACM, pp. 125–138.
- [9] ECKERSLEY, P. How unique is your web browser? In *Privacy Enhancing Technologies* (2010), Springer, pp. 1–18.
- [10] HUANG, K. T. D. Y., GRIERD, D. W. E. B. C., HOLT, T. J., KRUEGEL, C., MCCOY, D., SAVAGE, S., AND VIGNA, G. Framing dependencies introduced by underground commoditization.
- [11] IDZIOREK, J., AND TANNIAN, M. Exploiting cloud utility models for profit and ruin. In *Cloud Computing (CLOUD), 2011 IEEE International Conference on* (2011), IEEE, pp. 33–40.
- [12] IDZIOREK, J., TANNIAN, M., AND JACOBSON, D. Detecting fraudulent use of cloud resources. In *Proceedings of the 3rd ACM workshop on Cloud computing security workshop* (2011), ACM, pp. 61–72.
- [13] IDZIOREK, J., TANNIAN, M., AND JACOBSON, D. Attribution of fraudulent resource consumption in the cloud. In *Cloud Computing (CLOUD), 2012 IEEE 5th International Conference on* (2012), IEEE, pp. 99–106.
- [14] IDZIOREK, J., TANNIAN, M. F., AND JACOBSON, D. The insecurity of cloud utility models. *IT Professional*, 2 (2013), 22–27.
- [15] JUNG, J., KRISHNAMURTHY, B., AND RABINOVICH, M. Flash crowds and denial of service attacks: Characterization and implications for cdns and web sites. In *Proceedings of the 11th international conference on World Wide Web* (2002), ACM, pp. 293–304.
- [16] KARAMI, M., AND MCCOY, D. Understanding the emerging threat of ddos-as-a-service. In *Proceedings of the USENIX Workshop on Large-Scale Exploits and Emergent Threats* (2013).
- [17] KARAMI, M., PARK, Y., AND MCCOY, D. Stress testing the booters: Understanding and undermining the business of ddos services. In *Proceedings of the World Wide Web Conference (WWW)* (2016).
- [18] KHOR, S. H., AND NAKAO, A. spow: On-demand cloud-based eddos mitigation mechanism. In *HotDep (Fifth Workshop on Hot Topics in System Dependability)* (2009).
- [19] MITCHELL, T. M. Machine learning. 1997. *Burr Ridge, IL: McGraw Hill* 45 (1997), 995.
- [20] MOTOYAMA, M., LEVCHENKO, K., KANICH, C., MCCOY, D., VOELKER, G. M., AND SAVAGE, S. Re: Captchas—understanding captcha-solving services in an economic context. In *USENIX Security Symposium* (2010), vol. 10, p. 3.
- [21] NARESH KUMAR, M., SUJATHA, P., KALVA, V., NAGORI, R., KATUKOJWALA, A. K., AND KUMAR, M. Mitigating economic denial of sustainability (edos) in cloud computing using in-cloud scrubber service. In *Computational Intelligence and Communication Networks (CICN), 2012 Fourth International Conference on* (2012), IEEE, pp. 535–539.
- [22] OIKONOMOU, G., AND MIRKOVIC, J. Modeling human behavior for defense against flash-crowd attacks. In *Communications, 2009. ICC'09. IEEE International Conference on* (2009), IEEE, pp. 1–6.
- [23] RYAN, M. D. Cloud computing security: The scientific challenge, and a survey of solutions. *Journal of Systems and Software* 86, 9 (2013), 2263–2268.
- [24] SOOD, A. K., AND ENBODY, R. J. Crimeware-as-a-service—A survey of commoditized crimeware in the underground market. *International Journal of Critical Infrastructure Protection* 6, 1 (2013), 28–38.
- [25] SQALLI, M. H., AL-HAIDARI, F., AND SALAH, K. Edos-shield—a two-steps mitigation technique against edos attacks in cloud computing. In *Utility and Cloud Computing (UCC), 2011 Fourth IEEE International Conference on* (2011), IEEE, pp. 49–56.
- [26] SUBASHINI, S., AND KAVITHA, V. A survey on security issues in service delivery models of cloud computing. *Journal of network and computer applications* 34, 1 (2011), 1–11.
- [27] TAKABI, H., JOSHI, J. B., AND AHN, G.-J. Security and privacy challenges in cloud computing environments. *IEEE Security & Privacy*, 6 (2010), 24–31.
- [28] THING, V. L., SLOMAN, M., AND DULAY, N. A survey of bots used for distributed denial of service attacks. In *New Approaches for Security, Privacy and Trust in Complex Environments*. Springer, 2007, pp. 229–240.
- [29] WANG, H., XI, Z., LI, F., AND CHEN, S. Abusing Public Third-Party Services for EDos Attacks. In *10th USENIX Workshop on Offensive Technologies (WOOT 16)*. 2016.
- [30] WEN, S., JIA, W., ZHOU, W., ZHOU, W., AND XU, C. Cald: Surviving various application-layer ddos attacks that mimic flash crowd. In *Network and System Security (NSS), 2010 4th International Conference on* (2010), IEEE, pp. 247–254.
- [31] YU, S.-Z. Hidden semi-markov models. *Artificial Intelligence* 174, 2 (2010), 215–243.
- [32] YU, S.-Z., AND KOBAYASHI, H. An efficient forward-backward algorithm for an explicit-duration hidden markov model. *Signal Processing Letters, IEEE* 10, 1 (2003), 11–14.