

Insights into Attacks' Progression: Prediction of Spatio-Temporal Behavior of DDoS Attacks

Ahmed Abusnaina¹, Mohammed Abuhamad², DaeHun Nyang³,
Songqing Chen⁴, An Wang⁵, and David Mohaisen¹

¹ University of Central Florida, Orlando FL 32816, USA

² Loyola University Chicago, Chicago IL 60660, USA

³ Ewha Womans University, Incheon, South Korea

⁴ George Mason University, Fairfax VA 22030, USA

⁵ Case Western Reserve University, Cleveland OH 44106, USA

Abstract. DDoS attacks are an immense threat to online services, and numerous studies have been done to detect and defend against them. DDoS attacks, however, are becoming more sophisticated and launched with different purposes, making the detection and instant defense as important as analyzing the behavior of the attack during and after it takes place. Studying and modeling the Spatio-temporal evolvement of DDoS attacks is essential to predict, assess, and combat the problem, since recent studies have shown the emergence of wider and more powerful adversaries. This work aims to model seven Spatio-temporal behavioral characteristics of DDoS attacks, including the attack magnitude, the adversaries' botnet information, and the attack's source locality down to the organization. We leverage four state-of-the-art deep learning methods to construct an ensemble of models to capture and predict behavioral patterns of the attack. The proposed ensemble operates in two frequencies, hourly and daily, to actively model and predict the attack behavior and evolvement, and oversee the effect of implementing a defense mechanism.

Keywords: DDoS Attacks Prediction · Deep Learning.

1 Introduction

Distributed Denial-of-Service (DDoS) attacks are explicit malicious attempts to prevent legitimate users from accessing a service by sending an overwhelming amount of traffic to the service server. According to Netscout's annual worldwide infrastructure security report [15], the traffic generated for launching DDoS attacks exceeded 1 TBPS in size in 2019. On a more recent event, an attack of size 1.7 TBPS has been recorded. These attacks, if successful, result in a service shutdown that costs a provider an average of \$221,836 per attack [15].

The growing threat of DDoS attacks has inspired many recent research studies to contribute to the efforts toward the analysis and characterization of the attacks [16,17], including methods for the attacks detection and prediction [6,14]. These efforts have made the field of detecting DDoS attacks widely-explored

and resulting in highly-accurate detection systems [18,10,8]. However, there are limited studies that explore behavioral patterns and characteristics of the DDoS attacks during the progression of the attack and after the detection. Understanding the Spatio-temporal behavior and characteristics of the attack is crucial for defending against the attack, limiting its impact, and planing countermeasures to prevent it from occurring in the future. This study aims to contribute to this area by providing in-depth analyses and insights for modeling seven behavioral characteristics of DDoS attacks using deep learning-based methods. This analysis and modeling task takes place after the detection of the attack and continues as the attack progresses (in space and time). The Spatio-temporal analysis of DDoS behavior can be done by addressing various characteristics, such as the attack magnitude, botnet information, and attack source location.

This paper is dedicated to investigating several Spatio-temporal characteristics of the DDoS attacks, namely, attack magnitude, botnet family and ID, attack source locations including countries, ASNs, cities, and organizations. Due to the underlying nature of patterns to be extracted for separate characteristics, we leverage current state-of-the-art machine learning methods, including Deep Neural Networks (DNN), Long Short Term Memory (LSTM), Transformer, and Convolutional Neural Networks (CNN), to model separate characteristics and construct an ensemble of models to predict at different frequencies the behavioral patterns of DDoS attacks. The ensemble incorporates 14 different models, two for each characteristic, and operates in two frequencies, hourly-based, and daily-based frequencies, to actively monitor and account for the latest status of the attack while in progress. The ensemble is built and evaluated on a large-scale real-world dataset that includes 50,704 verified DDoS attacks launched by eleven botnet families and 674 botnet IDs on 9,026 targets from August 2012 to March 2013. This work sheds light on different aspects and patterns of DDoS attacks.

Contribution. This work presents an ensemble of models to predict the Spatio-temporal behavioral patterns of DDoS attacks. The contribution is as follows:

- **Modeling Spatio-temporal Characteristics:** Predicting seven different characteristics of the ongoing DDoS attacks using Spatio-temporal behavioral patterns of the attack, namely: *attack magnitude*, *botnet family*, *botnet ID*, *attack source country*, *ASN*, *city*, and *organization*, using large-scale real-world dataset of approximately nine million records of verified DDoS attacks.
- **Constructing Predictive Ensemble:** Implementing an ensemble of seven models based on four machine learning architectures, namely, DNN, LSTM, Transformer, and CNN, to actively predict the attack behavior on different operational frequencies (hourly and daily bases).
- **Addressing Unseen Attacks and Targets:** Evaluate the performance of the ensemble on a real-world large-scale dataset of known and unseen targets and DDoS attacks. The ensemble offers high accuracy over targets with no attacking history, and new represented DDoS attacks.
- **Addressing the Cold Start Problem:** We investigate the effect of cold start problem, i.e., modeling with insufficient information such as at the be-

gining of the attack. We show that the ensemble can achieve high accuracy even under the cold start situation.

2 Dataset Overview

2.1 Dataset Collection

The dataset is provided by the monitoring unit of a DDoS mitigation company [3]. Traces of malicious infected hosts were collected by collaborating with over 300 major Internet Service Providers (ISPs) globally monitoring attacks launched by specific malicious actors worldwide across America, Europe, Asia, Africa, and Australia. The activities of the participating hosts in the given botnet attacks, by either communicating with the infected infrastructure or launching the actual DDoS attack on the monitored targets, were monitored and analyzed over time. To this end, the traces of the traffic associated with various botnets were collected using different sensors on the Internet, in corporation with several ISPs, where the source of the collected traffic is an infected host participating in botnet attacks, and the destination is a verified targeted client. Afterward, malware botnets used in launching various attacks were reverse engineered and labeled to a known malware family using best practices (i.e., AMAL, a fully automated system for analysis, classification, and clustering of malware samples) [12,13]. The dataset consists of 50,704 verified DDoS attacks collected in the period of 08/29/2012 to 03/24/2013, a total of 207 days, targeting 9,026 clients, represented as hourly snapshots of each family activities over the monitored period, including the botnet information, targeted client IP, and the IPs of the hosts associated with the botnet attack.

2.2 Behavioral Characteristics of DDoS Attacks

We focus on three groups of characteristics: attack magnitude, botnet information, and attack source location. The following is a description of each group.

Attack Magnitude (AM). This attribute refers to the number of DDoS attacks launched by infected hosts on a specific target over a period of time, regardless of their malicious families and attack objectives. It is important to understand the magnitude of the attack to estimate and allocate a suitable amount of resources to counter the attack.

Botnet Information. The importance of knowing the attacking botnet families lies in implementing the correct defense against the attack since popular botnets have well-known attack patterns. Therefore, two characteristics have been extracted: *botnet family (BF)* and *ID*. The DDoS attacks reported in our dataset originated mainly from eleven popular botnet families: *dirtjumper*, *darkshell*, *ddoser*, *nitel*, *pandora*, *blackenergy*, *optima*, *aldibot*, *yzf*, *colddeath*, and *armageddon*. Botnet families may evolve over time. Therefore, new botnet generations are marked by their unique MD5 and SHA-1 hashes. We consider the botnet ID as a standalone characteristic, as the behavior of the botnet may change over

Table 1: Distribution of the botnet IDs over botnet families.

Family	# Botnet IDs	# Records
dirtjumper	251	6,902,882
darkshell	166	80,129
ddoser	102	37,172
nitol	43	20,411
pandora	41	1,397,027
blackenergy	28	95,330
optima	25	41,321
aldibot	9	269
yzf	6	113,923
colddeath	2	28,259
armageddon	1	906
Total	674	8,717,629

several generations. Table 1 shows the number of botnet IDs associated with DDoS attacks for each family. Note that the eleven botnet families have a total of 674 different botnet IDs, indicating the continuous evolvement of botnets over time. The number of records represents the instances of recorded DDoS attacks associated with infected hosts from a malicious botnet family.

Attack Source Location. It has been shown that botnets have strong geographical and organizational localities [2]. Therefore, such information can be used to predict future attack source locations and the shifting patterns of attackers across geographical locations to help in planning defenses and countermeasures. To this end, the hosts IP addresses were used to extract the attack source country (*CN*), city (*CT*), organization (*OG*) and (*ASN*), using the IP-to-region dataset and MaxMind online database [11]. In the monitored duration in which the dataset is collected, the attack source locations were distributed over 186 countries, 2,996 cities, 4,036 organizations, and 4,375 ASNs. The distribution of the infected hosts indicates the existence of worldwide botnet infections.

2.3 Dataset Splitting

The dataset is split into three parts as follows. ① *Training dataset*: The training dataset contains the traces and records of 80% (7220) of DDoS attacks' victims (i.e., targeted clients). For the purpose of predicting the behavioral patterns of the attacks during the attack progression, we considered the records that occurred at the first 80% of the attack duration for each victim (*target*) as the actual training dataset. ② *Known targets testing dataset*: This dataset contains the remaining records that occurred during the last 20% of the attack duration per target. This sub-dataset is used to evaluate the prediction models in modeling the behavioral pattern of DDoS attacks on targets with known history (by observing the earlier 80% of the attack duration). ③ *Unseen targets testing dataset*: This dataset consists of DDoS attack records of the remaining 20% (1806) of targeted clients that are not considered in the training dataset. The aim of this dataset is to evaluate the prediction models over targets with

Table 2: Overall characteristics of the dataset distribution.

Partition	# Targets	# Families	# IDs	# IPs	# Countries	# ASN	# Cities	# Org.
① Train Dataset	7,220	11	605	841,471	186	4,150	2,877	3,831
② Known Targets	7,220	11	606	158,230	179	3,275	2,275	3,024
③ Unseen Targets	1,806	10	248	234,113	151	2,571	1,800	2,382
Overall	9,026	11	674	880,451	186	4,375	2,996	4,036




Fig. 1: The general flow of the DDoS attacks prediction design. Here, T refers to the attacked target, whereas H represents one hour in the attack duration.

no attack history available to our model. Table 2 shows the distribution of the dataset characteristics over each partition of the dataset.

3 System Design

The proposed design aims to predict the seven different characteristics of the DDoS attack. The system design is shown in Figure 1.

3.1 Operational Frequency and Data Pre-processing

We adopted two operational frequencies to model and analyze behavioral patterns of DDoS attacks. The data pre-processing and handling follows the same manner in both approaches with slight modifications.

Operational Mode. For studying attack behavior manifested with the considered characteristics, data records were aggregated at different frequencies (i.e., *Agile* with hourly frequency and *Passive* with daily frequency). The agile mode requires six hours of data to be fully-functional at an hourly frequency, while the passive mode requires three days of information to be full-functional in modeling behavioral patterns at a daily frequency.

Data Processing and Sequence Generation. Addressing different characteristics of DDoS attacks captured by their records, the data is represented as

$\Phi_{\mathcal{X}} = \{\phi_1, \phi_2, \dots, \phi_t\} \in \mathbb{R}^{N \times T}$, where $\phi_{\alpha} \in \mathbb{R}^{1 \times T}$ is a vector of the attribute in hand (Φ) for the attack targeting $\mathcal{X}$ at a given time step α (e.g., ϕ_1 and ϕ_t represent the vectors of the first and last time step), T is the maximum length of the reported attacks, and N is the total number of targeted clients. For instance, addressing the *botnet ID* of an attack targeting $\mathcal{X}$, the data is represented as a matrix $ID_{\mathcal{X}} = \{id_1, id_2, \dots, id_t\} \in \mathbb{R}^{N \times T}$, where $id_{\alpha} \in \mathbb{R}^{1 \times T}$ is a vector of botnet IDs targeting $\mathcal{X}$ at a given time step α . We achieve such representation by the following steps. **Ⓐ Tokenization and Encoding:** We assign identifiers for unique elements (e.g., botnet IDs are assigned to unique identifiers when processing the *ID* attribute). Assuming an attack at target $\mathcal{X}$ in a time step α , the *ID* attribute is represented with a vector of all unique botnet IDs identifiers occurring in the attack record within α . For example, assuming the IDs appear in a certain attack record at the first time step are $\{id_{32}, id_{105}, id_{12}\}$, then, we present the vector as $ID_0 = \{id_{12}, id_{32}, id_{105}\}$. **Ⓑ Sequence Extraction:** The sequence of attribute behavior of DDoS attacks is extracted with different frequencies. Sequence extraction refers to the length of the previous time steps required to predict future steps. In the agile approach, we chose six-time steps (i.e., six hours) to be a sufficient time needed to predict future behaviors based on our experiment. For example, IDs sequences are generated as follows: $Seq_1 = \{ID_1, ID_2, \dots, ID_6\}$, $Seq_2 = \{ID_2, ID_3, \dots, ID_7\}$, and so on. Operating with the passive approach, we chose three time steps (three days) as sufficient information to predict daily future behavior. **Ⓒ Attribute Vector Padding:** The input data for each attribute are presented with different lengths based on the attribute magnitude at each time step. To allow efficient processing and tensor calculation, all vectors are padded to the maximum length enabling the packing of several attribute vectors in one sequence as well as packing several sequences in one batch. **Ⓓ Attribute Vector Embedding:** Attribute vectors are forwarded to an embedding layer in all deep learning-based models in our ensemble, to enable the compact representation of vectors. Vectors represented with attribute identifiers $\phi_{\alpha} \in \mathbb{R}^T$, where T is the maximum occurrence of unique identifiers in an attack, will be embedded to $\gamma_{\alpha} \in \mathbb{R}^{128}$, where 128 is the size of the vector embedding. We chose the size of the embedding based on several experiments that showed 128 is adequate to incorporate the information present in the attribute vector. Sequences are then viewed as matrices of $\Gamma_{\alpha} \in \mathbb{R}^{t_s \times 128}$, where t_s is the number of time steps.

Attack Magnitude. The approach to predict and study attack magnitude is different from the one adopted for other characteristics. The magnitude of the attack is calculated per targeted client at each time step and presented as one real value (instead of attribute vector). Thus, only step **Ⓑ** is required from the aforementioned approach, which aims to generate sequences of the calculated value of magnitude at each time step. To present the values of magnitude to the deep learning model, we normalize the values in the range of zero to one.

3.2 Prediction Models Architectures

Our approach adopts an ensemble of powerful classifiers to predict different behaviors of DDoS attacks including DNN, Transformer, LSTM, and CNN. We

chose different model architectures for modeling different tasks (i.e., characteristics behaviors) since certain architectures are proven to work better than the others in certain circumstances. In particular, the best performing deep learning architecture in predicting each DDoS attack characteristic is reported.

DNN for Attack Magnitude. The model architecture consists of four dense layers of size 1,000 units with ReLU activation function. Each dense layer is followed by a dropout operation with a rate of 30%. The last layer is connected to a sigmoid layer of size one signaling the normalized number of the attack magnitude (i.e., the scale of the magnitude from zero to one).

Transformer for Botnet Information. The model is adopted from the model proposed by Vaswani et al. [19]. It consists of stacked layers in both encoder and decoder components. Each layer has eight sub-layers comprising multi-head attention layers. The prediction is done by conducting a beam search with a length penalty ($\lambda = 0.6$). The Transformer is used to train two models performing two separate tasks, predicting botnet family and ID.

LSTM for Wide Attack Locality. The model consists of one LSTM layer with a size of 128 units. The LSTM layer is followed by a dense layer of size 128 and a dropout operation with a rate of 20%. Then, a dense layer with a sigmoid activation function is used to output the prediction of attack source locality. The LSTM is used to predict attack source country and ASN.

CNN for Specific Attack Locality. The model architecture consists of one convolutional layer with 64 kernels of size 1×3 convolving over the input vector, followed with a sigmoid output layer of size equals to the size of the addressed attribute vector (i.e., to predict the future status of the attack). The CNN architecture is used to predict the specific attack locality (i.e., city and organization).

4 Evaluation and Discussion

We report our results using two evaluation metrics, namely True Positive Rate (TPR) and True Negative Rate (TNR). TPR represents the number of correctly predicted elements over all the elements that occurred within the duration of the prediction. For instance, if the DDoS attack launched from four countries, of which, the prediction model predicts three correctly, the TPR is equal to 75% ($3/4$). TNR is referred to as $1 - (FP/N)$ where FP is the number of the incorrectly predicted elements and N represents all the elements that did not occur within the duration of the prediction. For instance, if the DDoS attack launched from four countries out of 186, and the prediction model incorrectly predicts two elements, the TNR is equal to 98.90% ($1 - (2/182)$). Note that TPR and TNR are preferred metrics in evaluating the systems as true indicators of performance in different scenarios. For example, achieving a TNR of 100% means zero false alarms. On the other hand, TPR indicates the precision of predicting attack behavior. Therefore, it's important for all models to maintain high TPR and TNR to ensure the usefulness of the classifier prediction.




Fig. 2: Evaluation of the prediction models over known and unseen targets.

4.1 Attack Characteristics Evaluation

Figure 2 summarizes the results of six attack characteristics using both *known* and *unseen* test targets when adopting the agile and passive approaches. The seventh attribute (i.e., the attack magnitude) is evaluated separately due to the data nature of the attribute. Two models were implemented for each attribute, one for each operational frequency (agile and passive operational modes), and evaluated on known and unseen targets.

Attack Magnitude. We evaluate the DNN model for predicting the attack magnitude using the *mean error* metric. Since the data observations were normalized, the output of the model indicates the magnitude as a fraction of the maximum recorded magnitude. Then, to calculate the magnitude, we multiply the model’s output by the maximum magnitude rounded to the decimal point. We report the results in Table 3. Here, the shift error is reported by the actual number of attacking hosts contributed to the attacks. For instance, the error rate of the agile model on *unseen* data is 0.0014% which is off by roughly 86 hosts from the actual number of the attacking hosts. Even though this number might seem large at first, it appears to be a good estimate knowing that the average attack magnitude on the agile data sampling rate (i.e., per the hour) equals 551 hosts (15.60% deviation). Similarly, the average shift rate for the passive approach is roughly 1,977 hosts for predicting the magnitude of *unseen* targets, which is also acceptable estimation knowing the average of magnitude is 15,394 hosts per day (12.97% deviation).

Botnet Family. Figure 2a shows the evaluation of the Transformer architecture trained to predict botnet family using different settings. The models achieve TPR of 95.97% and TNR of 99.65% for predicting botnet families of known targets on one-hour frequency, while maintaining TPR of 79.50% and TNR of 98.94% for unseen targets. The TPR score increases to 96.97% and 93.62% when using lower frequency (i.e., one-day) for known and unseen targets, respectively.

Table 3: Attack magnitude prediction evaluation.

Approach	Target	Mean Error Rate	Avg. Shift Error
Agile	Known	0.015%	∓ 88.64
	Unseen	0.014%	∓ 85.87
Passive	Known	0.012%	$\mp 1,734.40$
	Unseen	0.014%	$\mp 1,976.65$

Botnet ID. Figure 2b shows the evaluation of the performance of the Transformer-based prediction model on known and unseen targets for agile and passive operational frequencies. The model achieved a TPR of 90.16% and 76.42% for predicting known targets, and unseen targets with a TNR of 99.97% and 99.95%, respectively, using agile operational frequency. For passive operational frequency, the model achieved a TPR of 62.96% and 52.74% for predicting known targets, with a TNR of 99.95% and 99.93% for unseen targets, respectively.

Attack Source Country. Figure 2c shows the performance of the LSTM-based model on known and unseen targets for agile and passive operational frequencies. Using the agile approach, we achieved a TPR and TNR of 85.26% and 98.62% for known targets, and 83.83% and 99.95% for unseen targets, respectively. Similarly, the model achieved a TPR and TNR of 90.19% and 93.21% in predicting known targets attack source countries using passive frequency, and a TPR and TNR of 82.60% and 95.39% in predicting unseen target attack source countries.

Attack Source ASN. Figure 2d shows the evaluation of the LSTM-based models in predicting the attack source ASNs operating in two frequencies, agile and passive. The model achieved a TPR and TNR of 73.59% and 99.41% on known targets, and 65.68% and 99.96% on unseen targets, respectively, operating in agile frequency. Similarly, the model achieved a TPR and TNR of 27.06% and 97.53% on known targets, and 26.66% and 97.60% on unseen targets, respectively, on the passive approach. While the passive frequency-based LSTM model performance is low, it maintains a high TNR, reducing the false alarms.

Attack Source City. Using the agile frequency-based CNN model to predict the attack source city, we achieved a TPR and TNR of 72.23% and 99.72% for known targets, and 44.61% and 99.98% for unseen targets, respectively. For daily-based frequency, we achieved a TPR and TNR of 62.81% and 99.02% for known targets, and 17.39% and 99.34% for unseen targets, respectively. Figure 2e shows evaluation results of the CNN-based models. The high TNR (low false alarms) makes it possible to utilize the provided information by the model to implement a proper defense with high confidence.

Attack Source Organization. Figure 2f shows the evaluation of the performance of the prediction models on known and unseen targets for both operational frequencies. We achieved a TPR and TNR of 80.42% and 99.40% on known targets, and 84.48% and 99.72% on unseen targets, respectively, using agile frequency operational mode.

4.2 Discussion and Limitation

DDoS Attack Behavior Prediction. This work focuses on predicting the DDoS attack behavioral patterns after the detection of the attack. Therefore, the ensemble operates on top of the DDoS attack detection system. The purpose of the ensemble is to provide critical information and insights to help the targeted victims in designing and planning a proper defense mechanism.

- *Magnitude driven defenses:* DDoS attacks with a low magnitude will unlikely result in total denial of service, while high magnitude attacks can cause shutting down the service. Understanding the ongoing attack magnitude within a continuous time window allows a better decision making process and allocating resources to combat and mitigate the attack.
- *Botnet-based driven defenses:* Certain botnet families have repetitive attacking patterns. In addition, botnet families can collaborate to conduct a DDoS attack. Understanding the attack nature and behavior through its associate botnet families and IDs create a better awareness of how the attack will progress, and better defend against it.
- *Region-based driven defenses:* DDoS attacks have regional dependencies, as the infected hosts may be originated from the same region, or related regions. Understanding the regional distribution of the infected hosts, and the over-time shifting will provide better insights to implement region-based defenses.

First-hour Attack: The Cold Start. We implemented the ensemble to operate on the specified frequency using the available information aggregated using the sampling time while padding the unavailable sequence steps with zero-vectors. For example, assume an attack with only two-hours information is available, the agile approach will process the two-hours vectors and pad four-steps of zero-vectors to predict the third hour. This approach has shown to be effective in our experiments, especially for predicting botnet families and attack source countries. For instance, using six-hours information, the agile approach predicts the attack source cities with TPR of 72.23% and 44.61%, for known and unseen targets; while using only one-hour information results in TPR of 65.56% and 17.49% for the same settings, while maintaining a high TNR ($\approx 99\%$).

5 Related Work

DDoS Attacks Detection. DDoS attacks detection is well explored in different environments. Sekar et al. [18] proposed LADS, a triggered, multi-stage in-network DDoS detection system to overcome the scalability issues in detecting DDoS attacks over the large-scale monitored network. In addition, Chang et al. [1,2] performed an in-depth analysis of botnet behavior patterns. Their analysis showed that different botnets start to collaborate when launching DDoS attacks. Similarly, they conducted an in-depth analysis measurement study showing that bots recruitment has strong geographical and organizational locality.

Lu et al. [9] clustered botnet traffic into C&C channels using the K-mean clustering algorithm on large-scale network traffic payload signatures. In more recent work, Doshi et al. [4] distinguished normal traffic from DDoS traffic using limited packet-level features.

DDoS Attacks Behavior Prediction. Recent studies predicted different aspects of the attack behavior, such as Gupta et al. [7], where they estimated the number of bots involved in a flooding DDoS attack with high accuracy by calculating various statistical performance measures. In addition, Fachkha et al. [5] proposed a systematic approach for inferring DDoS activities, predicting DDoS attack characteristics, and clustering various targets to the same DDoS campaign. Furthermore, Wang et al. [20] designed three DDoS attacks models from temporal (attack magnitudes), spatial (attacker origin), and Spatio-temporal (attack inter-launching time) perspectives by analyzing 50,000 verified DDoS attacks. Even though recent studies investigated the attack detection and behaviors, only a few of them provided information that would assist the client in implementing a proper defense on the spot. Our design provides the victim with essential information that can be utilized to properly implement a magnitude-, region-, and malware-based DDoS attacks mitigation techniques and defenses.

6 Conclusion

This work proposes an ensemble approach for studying and predicting the behavioral characteristics of DDoS attacks. Toward this, we built an ensemble of deep learning models to predict seven behavioral characteristics of DDoS attacks, providing insights for handling such attacks. Evaluating our approach on a large-scale real-world dataset that contains records of more than fifty thousand verified attacks, the results of our approach show remarkable performance when operating on different sampling frequencies and under different settings. This success of efficient and accurate modeling of DDoS attack characteristics can help to implement proper defenses for mitigating the attack.

Acknowledgement. This work was supported by NRF grant 2016K1A1A2912757, NIST grant 70NANB18H272, and NSF grant CNS-1524462 (S. Chen), and by the Institute for Smart, Secure and Connected Systems at CWRU (A. Wang).

References

1. Chang, W., Mohaisen, A., Wang, A., Chen, S.: Measuring botnets in the wild: Some new trends. In: Proceedings of the 10th ACM Symposium on Information, Computer and Communications Security, ASIA CCS. pp. 645–650 (2015)
2. Chang, W., Mohaisen, A., Wang, A., Chen, S.: Understanding adversarial strategies from bot recruitment to scheduling. In: Proceedings of the Security and Privacy in Communication Networks - 13th International Conference, SecureComm. pp. 397–417 (2017)
3. Cymru, T.: Cymru. Available at [Online]: <https://www.team-cymru.com/> (2019)

4. Doshi, R., Apthorpe, N., Feamster, N.: Machine learning ddos detection for consumer internet of things devices. In: 2018 IEEE Security and Privacy Workshops (SPW). pp. 29–35. IEEE (2018)
5. Fachkha, C., Bou-Harb, E., Debbabi, M.: On the inference and prediction of ddos campaigns. *Wireless Communications and Mobile Computing* **15**(6), 1066–1078 (2015)
6. Gong, D., Tran, M., Shinde, S., Jin, H., Sekar, V., Saxena, P., Kang, M.S.: Practical verifiable in-network filtering for ddos defense. In: 39th IEEE International Conference on Distributed Computing Systems, ICDCS 2019, Dallas, TX, USA, July 7–10, 2019. pp. 1161–1174 (2019)
7. Gupta, B., Joshi, R., Misra, M.: Prediction of number of zombies in a ddos attack using polynomial regression model. *Journal of advances in information technology* **2**(1), 57–62 (2011)
8. Lee, K., Kim, J., Kwon, K.H., Han, Y., Kim, S.: Ddos attack detection method using cluster analysis. *Expert systems with applications* **34**(3), 1659–1665 (2008)
9. Lu, W., Rammidi, G., Ghorbani, A.A.: Clustering botnet communication traffic based on n-gram feature selection. *Computer Communications* **34**(3), 502–514 (2011)
10. Ma, X., Chen, Y.: Ddos detection method based on chaos analysis of network traffic entropy. *IEEE Communications Letters* **18**(1), 114–117 (2014)
11. MaxMind: Maxmind. Available at [Online]: <https://www.maxmind.com/> (2019)
12. Mohaisen, A., Alrawi, O.: Av-meter: An evaluation of antivirus scans and labels. In: Proceedings of the Detection of Intrusions and Malware, and Vulnerability Assessment - 11th International Conference, DIMVA. pp. 112–131 (2014)
13. Mohaisen, A., Alrawi, O., Mohaisen, M.: AMAL: high-fidelity, behavior-based automated malware analysis and classification. *Computers & Security*. **52**, 251–266 (2015)
14. Najafabadi, M.M., Khoshgoftaar, T.M., Calvert, C., Kemp, C.: A text mining approach for anomaly detection in application layer ddos attacks. In: Proceedings of the Thirtieth International Florida Artificial Intelligence Research Society Conference, FLAIRS. pp. 312–317 (2017)
15. Netscout: Netscout 14th annual worldwide infrastructure security report. Available at [Online]: <https://www.netscout.com/report/> (2019)
16. Rasti, R., Murthy, M., Weaver, N., Paxson, V.: Temporal lensing and its application in pulsing denial-of-service attacks. In: Proceedings of the IEEE Symposium on Security and Privacy, SP. pp. 187–198 (2015)
17. Rossow, C.: Amplification hell: Revisiting network protocols for ddos abuse. In: Proceedings of the 21st Annual Network and Distributed System Security Symposium, NDSS. (2014)
18. Sekar, V., Duffield, N.G., Spatscheck, O., van der Merwe, J.E., Zhang, H.: LADS: large-scale automated ddos detection system. In: Proceedings of the 2006 USENIX Annual Technical Conference, Boston, MA, USA, May 30 - June 3, 2006. pp. 171–184 (2006)
19. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems. pp. 5998–6008 (2017)
20. Wang, A., Mohaisen, A., Chen, S.: An adversary-centric behavior modeling of ddos attacks. In: Proceedings of the 37th IEEE International Conference on Distributed Computing Systems, ICDCS. pp. 1126–1136 (2017)