Fundamental Scaling Laws of Covert DDoS Attacks

Amir Reza Ramtin[1], Philippe Nain[2], Daniel S. Menasche[3],
Don Towsley[1], Edmundo de Souza e Silva[3]
[1]UMass Amherst, USA, [2]INRIA, France, [3]UFRJ, Brazil

The Internet has become an indispensable commodity in the last several years. This achievement was parallel to the growth of sophistication that home networks have undergone, nowadays hosting a variety of devices such as PCs, tablets, mobile phones and specialized apparatus such as smart thermostats and other Internet of things (IoT) devices. While these devices offer users an array of services and conveniences, they come at the cost of increasing the attack surface of the home network.Because of the vulnerabilities of such devices, they have been increasingly used as the source of Distributed Denial-of-Service (DDoS) attacks.

DDoS attacks are difficult to prevent, because they are launched from a large number of infected devices connected to the Internet, collectively known as botnets. The attacker compromises devices by injecting malicious code (malware), which allows the attacker to perform actions at a later time using these devices as sources of harmful traffic without knowledge of the device's owner. The traffic generated by some botnets is typically composed of millions of small flows.

Despite all the continuing efforts to detect and mitigate these attacks, their number have not decreased. In fact, the number of DDoS attacks drastically increased in 2020 [2]. Roughly, DDoS attacks are produced by launching a burst of packets simultaneously from a very large number of devices towards a given target.

Needless to say, early identification of these attacks and their sources is of prime concern of companies. However, it is also imperative to discover if there are fundamental tradeoffs between the amount of damage an attacker can inflict to services and the attacker's ability to remain undetected. If these fundamental laws exist, they could shed some light concerning covertness versus damage and they could be used to help building effective DDoS countermeasures.

It should be evident that the objective of the attacker is to inflict as much damage as possible by generating enough traffic (for instance, generating a large amount of control packets) to wear out the victim's resources and, consequently, to disrupt user's services. The malicious traffic originates from home network devices with limited capacity. As such, the attack generated from a single home is far from sufficient to cause any damage. Then, necessarily, the attacker tries to use as many homes as possible, remotely activating a large number of controlled devices (the bots) that have been previously infected. Furthermore, it is advantageous for the attacker to remain *covert* (undiscovered) while attacking.

It is hard to differentiate attack traffic originating from a single home network from the regular home user traffic. This is probably why most network-based DDoS detection methods rely on detailed network traffic information (e.g., packet header data), which is in general computationally expensive and also raises concerns about user privacy. To avoid these drawbacks, methods based solely on metrics such as byte/packet counts should be preferred [6, 3]. A lightweight approach that employs network interface byte/packet counts also scales, and is oblivious to botnet-specific attack signatures and encryption.

Clearly, the larger the number of compromised devices in different home networks, the greater the amount of damage the attacker who controls these bots can potentially cause. The work of [6] proposes a method to detect an ongoing attack from a home-router without resorting to packet inspection, and also shows that the likelihood of detecting a DDoS attack can be improved with the number of participant bots in the attack. A fundamental question then arises: Can a DDoS attack be covert and if so, what is the damage it is expected to cause?

The covertness criterion considered in this paper was proposed in the context of low probability of detection (LPD) communications. Although there are a number of papers in this area [1, 8, 4, 5, 9], to the best of our knowledge the only other work that discussed covert DDoS attacks in home networks is [7], wherein we track a different setting where the admin can leverage several network traffic features to improve detection accuracy.

Goals. In this paper, we aim to avoid packet inspection as in [6]. Furthermore, it should be evident that the damage caused by common DDoS attacks (such as Mirai) is proportional to the number of infected devices (equivalently the number homes) participating in the attack. In addition, from the administrator's point of view, the number of false alarms should be kept to a minimum, since there is no point in detecting occasional attacks if the number of false alarms is unbearably high.

We then pose the following questions related to the attacker's ability to cause as much damage as possible and the likelihood to remain undetected: is there any fundamental limit on the damage an attacker can cause to the victim while remaining covert? if such limit exists, how is it related to the false alarm rate?

We stress that our goal is *not* to devise a deployable detection method but rather to discover fundamental laws that govern the covertness 'game' played between the attacker and admin and to understand the limits on the damage an

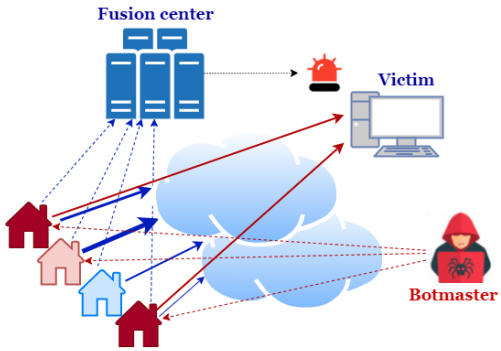

Figure 1: System outline

attacker can cause.

System description and model. To answer the above questions, we propose an analytical model to capture the essence of these attacks. The model comprises two components, characterizing regular traffic and traffic when an attack is underway. In particular, we focus on the simplest case wherein each component is associated with a single feature, such as packet counts observed per time slot.

The system we study is shown in Figure 1. In the figure, the blue dotted arrows represent measurement data collected at home routers and sent periodically to a data fusion center for analysis. The botmaster can issue commands to the infected homes (in red) but, as shown in the figure, the attacker can choose not to use all the homes he controls to initiate an attack, to cause significant damage and yet remain covert.

At a high level, we posit that an attacker is covert if admin running a *detector* (also known as a *classifier*) cannot determine if an attack is in progress by observing the traffic (byte or packet rate) from a set of homes. Formally, consider that admin runs an optimal statistical hypothesis test and uses it to compute the probability of false alarm (p_{FA}) and the probability of miss detection (p_{MD}). In this setting, the sum of errors $p_{FA} + p_{MD}$ lies in $[0, 1]$. Following the definition in [1], we then say that an attack is covert if the attacker has a strategy that makes the sum $p_{FA} + p_{MD}$ arbitrarily close to one.

Results. We establish that the amount of traffic that an attacker can issue while remaining covert grows as $O(\sqrt{n})$, where n is the number of compromised homes controlled by the attacker in the network (Figure 2). We also obtain conditions under which this bound is tight. We confirm these results using the real data collected at home-routers from a mid-sized ISP, with whom we partnered to gather statistics about baseline regular traffic. Our dataset includes packets and byte counts collected at thousands of home-routers over several months. Our analysis of the dataset shows that regular traffic can be modeled by a mixture of Gaussian distributions. We also use a dataset of attack traffic, generated by controlled experiments using real botnet code [6]. The traffic distribution of the attack traffic can also be approximately modeled by a mixture of Gaussian distributions.

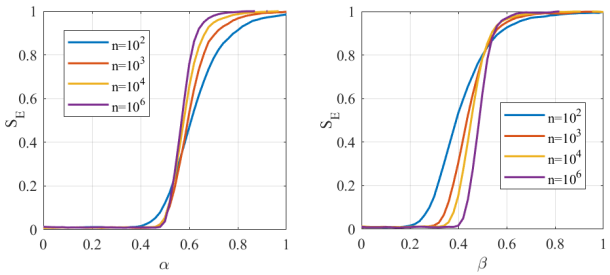

(a) all homes (b) fraction of homes

Figure 2: The error probability (S_E) as a function of the aggressiveness of the attacker. When the attacker issues attack from all homes (a), the average total traffic injected by the attacker is proportional to $n^{1-\alpha}$: as α grows the total attack traffic decreases and the probability of error transitions from 0 to 1. For values of n as small as 1,000, we already observe a sharp phase transition occurring around $\alpha = 0.5$, in agreement with the square root law. Then, we consider a variation where the attacker issues attack from a fraction of homes (b), given by $q(n) = n^{-\beta}$ ($\alpha = 0$). As β grows the total attack traffic decreases and S_E transitions from 0 to 1, with a phase transition at $\beta = 0.5$, agreeing with the square root law.

REFERENCES

[1] Boulat A. Bash, Dennis Goeckel, and Don Towsley. Square root law for communication with low probability of detection on AWGN channels. In *2012 IEEE International Symposium on Information Theory Proceedings*, pages 448–452. IEEE, 2012.

[2] Cloudflare. Network-layer DDoS attack trends for Q2 2020, 2020.

[3] Hamed Haddadi, Vassilis Christophides, Renata Teixeira, Kenjiro Cho, Shigeya Suzuki, and Adrian Perrig. Siotome: An edge-ISP collaborative architecture for IoT security. *Proc. IoTSec*, 2018.

[4] Ke-Wen Huang, Hui-Ming Wang, Don Towsley, and H. Vincent Poor. LPD communication: A sequential change-point detection perspective. *IEEE Transactions on Communications*, 2020.

[5] Bo Jiang, Philippe Nain, and Don Towsley. Covert cycle stealing in a single FIFO server. *ACM Transactions on Modeling and Performance Evaluation of Computing Systems (ToMPECS)*, 2021. To appear. arXiv preprint arXiv:2003.05135.

[6] G. Mendonça, G. H. A. Santos, E. de Souza e Silva, R. M. M. Leao, D. S. Menasche, and D. Towsley. An Extremely Lightweight Approach for DDoS Detection at Home Gateways. In *2019 IEEE International Conference on Big Data*, pages 5012–5021, 2019.

[7] Amir Reza Ramtin, Philippe Nain, Don Towsley, E. de Souza e Silva, and D. S. Menasche. Are covert ddos attacks facing multi-feature detectors feasible? *ACM SIGMETRICS Performance Evaluation Review*, 2021.

[8] Ramin Soltani, Dennis Goeckel, Don Towsley, and Amir Houmansadr. Fundamental limits of covert packet insertion. *IEEE Transactions on Communications*, 2020.

[9] Shihao Yan, Xiangyun Zhou, Jinsong Hu, and Stephen V. Hanly. Low probability of detection communication: Opportunities and challenges. *IEEE Wireless Communications*, 26(5):19–25, 2019.