



DDoS Mitigation Dilemma Exposed: A Two-Wave Attack with Collateral Damage of Millions

Lumin Shi¹, Jun Li¹(✉), Devkishen Sisodia¹, Mingwei Zhang¹,
Alberto Dainotti², and Peter Reiher³

¹ University of Oregon, Eugene, USA
lijun@cs.uoregon.edu

² Georgia Institute of Technology, Atlanta, USA
dainotti@gatech.edu

³ UCLA, Los Angeles, USA
reiher@cs.ucla.edu

Abstract. While mitigating link-flooding attacks on the Internet has become an essential task, little research has been done on how an attacker can further attack and abuse the mitigation solutions themselves. In this paper, we propose a two-wave attack with **collateral damage of millions** (or *Carom*), a new link-flooding attack that poses a mitigation dilemma for multiple simultaneously attacked networks, which must either endure the flooding attack or suffer unwanted side effects in mitigating the attack. Composed of practical components, the Carom attack aims to maximize the burden on attack mitigation systems and the collateral damage to defending networks, thereby wreaking havoc on large swaths of the Internet. After modeling real-world mitigation solutions, we evaluated the attack against the mitigation solutions with real-world datasets, showing the feasibility of the attack and quantifying the amount of damage it can inflict on today's Internet. We hope that this work can motivate the improvement of existing link-flooding mitigation solutions.

Keywords: Link-flooding attack · Distributed denial-of-Service (DDoS) · Collateral damage

1 Introduction

As the Internet continues to witness high-profile, large-scale distributed denial-of-service (DDoS) attacks, edge networks that receive DDoS traffic, defined as **first-wave victims** in this paper, only use DDoS mitigation solutions that are inadequate. The two typical DDoS mitigation protocols nowadays in use are remotely triggered black hole (RTBH) [21] filtering and BGP FlowSpec [29] which deploy traffic filters at routers in their upstream networks. The upstream networks will then distribute the filters at their traffic ingress

points to mitigate unwanted traffic before it propagates within the networks, hoping to unclog the links connected with the edge networks.

However, such mitigation protocols above have shortcomings. For example, RTBH removes *all* traffic, benign or malicious, towards a specified destination network prefix, which can generate massive amounts of collateral damage as it also filters traffic from legitimate sources, which we define as **second-wave victims**. While BGP FlowSpec enables fine-grained traffic filters with its extensive list of supported IP header fields, many networks implement them in their network routers or switches with little ternary content-addressable memory (TCAM) for storing filters [35]. Hence, networks often limit the number of filters their downstream customer networks can deploy, limiting the attack traffic they can cover.

On the other hand, fine-grained DDoS traffic filtering also face major challenges in order to effectively mitigate DDoS. First of all, as the victim is at the mercy of the attacker who decides how to construct each DDoS packet, the victim may not be able to identify attack flows accurately and thus derive filters that can effectively mitigate the attack. Further, as first-wave victims often use the source IP addresses of the attack traffic to match and filter attack traffic, the source-IP-based filtering method is only feasible when a victim has enough memory capacity on their routers *and* the attack does not employ IP spoofing (i.e., DDoS bots do not spoof their IP addresses). Sadly, to date, IP spoofing remains a significant problem on the Internet [25].

Given the above observations, we present a two-wave attack with collateral damage of millions, or *Carom*, a link-flooding attack that introduces a **mitigation dilemma** to many first-wave victims distributed among multiple autonomous systems (ASes) on the Internet. They must either endure the attack or disconnect themselves from a wide spread of second-wave victims during attack mitigation. The second-wave victims include disconnected networks that originate both attack and benign traffic and those networks that do not originate any attack traffic.

To the best of our knowledge, *Carom* is the first work that studies the procedure and consequences of launching a DDoS attack against multiple ASes simultaneously with practical techniques, and further exploiting the DDoS mitigation mechanisms in place to further trigger a second-wave attacks toward more victims. We also evaluate the attack against real-world mitigation solutions and demonstrate that the attack is practical and can inflict severe damage to today's Internet.

2 Related Work

Real-world DDoS attacks often employ simple, practical techniques. Lately, however, we begin to see real-world DDoS attacks that employ advanced attack techniques described in academic research [2]. Exemplified by the Coremelt attack [44] and the Crossfire attack [19], these advanced attacks do not need to directly flood targeted services but the network resource of the services. Or, exemplified by the pulsing attacks [20, 41, 45], they do not need to flood the targeted services constantly but deceive legitimate hosts to reduce their frequency

of issuing service requests. We describe and discuss these advanced DDoS attacks in this section. In particular, while the Carom attack, as we present in this paper, is feasible to launch on today’s Internet with severe consequences, we explain below that these advanced attacks are not so practical and face differing real-world challenges to launch successfully.

The Coremelt Attack [44]. This attack aims to overwhelm transit links on the Internet. It does not target any particular service but all services that rely on the overwhelmed transit links. In this attack, bots exchange ‘legitimate’ traffic among themselves, aiming to overwhelm the transit links that carry their traffic. The traffic is ‘legitimate’ since all traffic among the bots follow application protocols. Furthermore, the bots could mimic the traffic patterns of benign users to mask the attack. Hence, the attack is difficult to detect and mitigate.

However, the requirement to launch this attack is questionable on today’s Internet. The work uses inferred AS-level topologies [9] and a simplified link bandwidth model to evaluate the attack. It does not consider the rich, interconnected, router-level paths in transit networks nor the load-balancing schemes that utilize these rich connections. The work also assumes that there exist botnets that can generate sufficient ‘legitimate’ traffic to overrun core transit links. While not impossible, it is still extremely challenging to realize such an attack since most DDoS attacks today produce less than 100 Gbps attack bandwidth. Finally, the attack does not target any specific service, which reduces its practical value. Nevertheless, the work influenced another attack (i.e., the Crossfire attack [19]) on network resources that is more practical to launch.

The Crossfire Attack [19]. This attack sends traffic towards hosts in a selected set of networks to overwhelm their shared link. Unlike common DDoS attacks that send all attack traffic to the targeted victim services, bots in this attack distributes their traffic to multiple public-facing hosts who share the target upstream link, flooding and congesting the link and therefore all services downstream from the link. Here, the Crossfire attack has a more *focused target area* than the Coremelt attack [44]. Also note each host only receives partial attack traffic and each attack flow is of a low volume. The authors claim that real-world intrusion detection systems (IDS) will fail to sound alarms. Therefore, the DDoS mitigation will not be initiated.

While the Crossfire attack requires less attack resource than the Coremelt attack, the resource requirement is also not trivial as it needs to overrun an upstream link which is typically of a high bandwidth.

The Pulsing Attack [15, 20, 22, 26, 40]. A pulsing attack is to disconnect a network with periodic, short-lived traffic pulses. These pulses can cause congestion-aware flows to believe in traffic congestion, thereby reducing the sending rate of the congestion-aware flows. In addition, because the pulse duration is often short (e.g., 100s of milliseconds), DDoS detection systems may not detect such an attack since their traffic information feed (e.g., NetFlow/IPFIX) is too coarsely grained to spot the attack pulses. The main challenge of pulsing attacks is to synchronize the traffic pulses among DDoS bots. For example, a pulsing attack

typically requires a pulsing duration of 100s milliseconds, which is virtually impossible to accommodate across thousands, not to mention millions of bots. Indeed, Park et al. [37] show the synchronization difficulty in practice even with the latest pulsing attacks (e.g., CICADAS [20]) that claim higher feasibility. In the Carom attack, we apply pulsing durations that last for seconds rather than milliseconds; we trade the stealthiness of a pulsing attack for the attack feasibility in practice and the ability to observe the attack’s effectiveness.

3 Carom—A Two-Wave Attack with Collateral Damage of Millions

This section presents the procedure of the Carom attack. It includes the following phases:

- Reconnaissance. In this phase the attacker gathers information on bots, attack-source networks (i.e., networks that contain bots), and first-wave victims (i.e., networks that receive attack traffic).
- First wave. In this phase the attacker disconnects victims in cycles, each cycle attacking first-wave victims group by group.
- Second wave. Should a first-wave victim chooses to mitigate the attack, it is then forced to block its communication with what we call **second-wave victims**.

3.1 Reconnaissance

Before launching its attack, an adversary can gather all kinds of information useful to its attack, such as the available bandwidth between every bot and every victim, or list of the actively running services at victims. Below we focus on two specific tasks essential to Carom.

Verify IP Spoofing Capability. IP spoofing allows bots to send attack traffic with arbitrary source IP addresses. It is made possible when stub (edge) networks do not validate the source addresses of traffic leaving the networks. Carom leverages IP spoofing to increase its mitigation difficulty. Therefore, knowing the bots that have spoofing capability is crucial. We provide an elaborated discussion in Sect. 4. According to the Spoofer project [6], at the time of this writing, well over 22.5% of 8,067 autonomous systems (ASes) on the Internet allow IP spoofing consistently. The sampled $8k$ ASes account for less than 10% of all ASes on the Internet. In other words, the actual number of networks that allow IP spoofing can be drastically different. Despite years of research [24], the Internet continues to allow large-scale IP spoofing to happen [27].

Select Networks to Attack. Carom considers several factors when selecting its first-wave victims. First, the total attack bandwidth of the botnet should be larger than any selected first-wave victim’s link capacity. Albeit not always

accurate, with online databases such as PeeringDB [38], the adversary can infer a first-wave victim’s link capacity and build a list of first-wave victims that the botnet can disconnect.

The adversary may also configure Carom to attack networks that share the same network provider to increase the collateral damage during their attack mitigation. For example, the adversary can perform a traceroute scan to construct the router-level paths from each bot to each first-wave victim. The adversary can then group first-wave victims by their common network providers.

Carom also needs to rule out first-wave victims that major DDoS protection service (DPS) providers protect to avoid wasting attack traffic. Major DPS providers’ mitigation capacity ranges from several Tbps to 10s of Tbps [1, 12, 18, 32, 34, 39], which is sufficient to absorb recent large-scale DDoS attacks. An adversary can employ several approaches to learn DPS-protected networks. First, the adversary can leverage the traceroute results from the botnet to first-wave victims to search for routers’ hostnames or IP addresses associated with DPS providers. Second, the adversary may monitor for the round trip time changes of a first-wave victim to infer if anycast is invoked, as proposed by Sommese et al. [43]. Finally, the adversary can monitor the BGP announcements from DPS providers to know the network prefixes protected by DPS providers (i.e., such network prefixes do not belong to DPS providers).

3.2 First Wave via Moving Attacks

Carom can disconnect more first-wave victims than its botnet’s total attack bandwidth permits by moving from one group of victims to another. In essence, Carom can send attack traffic only towards a group of first-wave victims for a fixed amount of time, and then moves on to attack another group of first-wave victims. An attack cycle is completed whenever Carom begins to attack the first group again.

Carom relies on two conditions to change attack targets: *low attack effectiveness* and *pulse duration of an attack*. The first condition happens when a first-wave victim is under the protection from a DPS provider or, in rare cases, the victim’s upstream network. Carom utilizes bots from different geographical regions as vantage points to observe whether first-wave victims are overwhelmed during attack time. It can also use approaches in Sect. 3.1 to rule out victims that just subscribed to DPS providers upon the attack. The second condition, *pulse duration*, is a property that we extracted from pulsing attacks [15, 20, 22, 26, 40]. In general, a pulsing attack sends short-lived, bursty attack traffic (i.e., pulses) to a network at a frequency. We define the period with attack traffic as *pulse duration*, and the pulse volume as *pulse amplitude*. As the attack pulses can force congestion-aware flows to reduce their sending rate, the pulses can severely impact the user experience of real-time applications such as online gaming and conferencing calls. To address the feasibility issues of the pulsing attacks (discussed in Sect. 2), Carom launches attack pulses that last for seconds or a longer period of time.

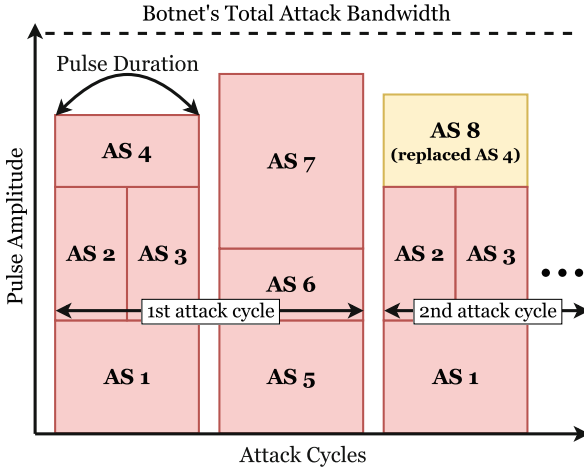


Fig. 1. A Carom attack example in action

We illustrate Carom in Fig. 1. In each attack cycle, a botnet attacks multiple attack groups. Each group contains a number of first-wave victims. For example, Carom completes an attack cycle as follows: The botnet first attacks the attack group that contains AS 1 to AS 4, it then attacks AS 5 to AS 7 in the second attack group. Note that the pulse duration does not have to be the same for each first-wave victim (e.g., AS 2 and 3 have a shorter pulse duration than other ASes). In other words, an adversary can increase the number of first-wave victims in each attack group by decreasing the pulse durations of some first-wave victims. Before the botnet attacks group 1 again, Carom replaces AS 4 with another first-wave victim (i.e., AS 8) since the attack effectiveness against AS 4 is low.

3.3 Second Wave by Exploiting DDoS Mitigation

The attack toward first-wave victims will trigger DDoS mitigation in place. Carom then tries to cause the DDoS mitigation to generate high collateral damage by dropping traffic from legitimate sources, i.e., second-wave victims. To do so, Carom generates DDoS traffic that is difficult for the DDoS mitigation to produce accurate filters. Below, we introduce three building blocks that Carom employs to generate DDoS traffic.

Indiscernible Packets. Attack packets are discernible when they contain shared characteristics that a mitigation system can leverage for accurate traffic filtering. Common link-flooding attacks consist of discernible attack packets. For example, in a DNS amplification attack, each amplified packet’s source port is 53 and has a DNS resolver’s IP address. In such an attack, a network can drop all traffic sourced from port 53 to mitigate the attack and only allow traffic from trusted DNS resolvers.

If a mitigation system cannot find shared characteristics of the attack packets, its mitigation accuracy suffers, such as when an attack generates packets with random packet headers (e.g., IP addresses, ports) and payloads. If a mitigation system deploys many filters that happen to cover benign traffic, it then may introduce a huge number of second-wave victims.

Traffic Dispersion. Carom disperses its attack packets among a set of destination IP addresses of a first-wave victim. Such a technique is also referred to as carpet bombing [11] by the network community. Traffic dispersion is the primary technique that enables the Crossfire attack [19], which aims to overwhelm a target network’s upstream link(s). However, the target network only receives part of the attack traffic. Hence, the authors of the Crossfire attack claim such an attack is difficult to detect. Unlike the Crossfire attack, we assume first-wave victims can always detect all attack traffic with perfect accuracy, and we focus on evaluating the mitigation dilemma a first-wave victim faces.

In case the network employs a stateful defense method, Carom disperses its attack traffic among the subnets that host services open to the Internet to increase the amount of attack traffic the network receives (see Sect. 4.3). Bots may also apply the traffic dispersion technique to its source IP addresses if they can spoof IP addresses.

In contrast, common DDoS attacks target one or a few IP addresses of the first-wave victim. Coarsely-grained filtering techniques, such as RTBH, are often sufficient to mitigate the attack with manageable collateral damage; RTBH removes all traffic (benign or malicious) towards the attacked IP addresses. In other words, such coarsely-grained mitigation is undesirable against traffic dispersion.

Stateful and Stateless Attack Traffic. Depending on the mitigation method of a first-wave victim, Carom generates attack traffic in either stateless or stateful mode. By default, Carom runs in the stateless mode where it sends TCP SYN or UDP packets with port numbers that the first-wave victim allows, and Carom discards the responses from the first-wave victim. However, in case the first-wave victim employs stateful mitigation solutions that filter traffic that does not comply with traffic protocols (e.g., TCP handshake), Carom will follow the protocol to bypass the mitigation. From there, Carom can generate traffic at a **congestion-unfriendly** rate to cause link congestion.

4 Mitigation Models Against Carom

We use three mitigation models (M1, M2, M3) to cover the mitigation methods in practice. For each mitigation model, we cover (1) the suitable Carom strategy against it, and (2) the conceivable collateral damage caused by the model with and without IP spoofing.

4.1 Stateless Coarse-Grained Mitigation (M1)

M1 employs RTBH and source-based RTBH (S/RTBH) to mitigate attack traffic by either destination or source IP addresses, respectively; they cannot take advantage of fine-grained DDoS detection results (e.g., attack flows presented in a 5-tuple format). Specifically, RTBH drops all traffic towards the IP addresses of a first-wave victim. If the RTBH is deployed by the *direct* upstream ASes of the first-wave victim, they effectively disconnect the first-wave victim from the Internet. On the other hand, if the first-wave victim can deploy RTBH filters at remote upstream ASes that are close to or at the attack-source networks, the collateral damage is reduced; it no longer blocks all networks from accessing the IP addresses. Unfortunately, Nawrocki et al. [31] show the latter scenario rarely happens in practice - remote ASes rarely accept inter-AS RTBH messages.

With S/RTBH, an AS can drop all traffic from the IP addresses that generate attack traffic. Because S/RTBH only drops traffic by source IP addresses, when deployed, the AS blocks all their customers networks (including the first-wave victim) from accessing the IP addresses that generate attack traffic. Therefore, S/RTBH is best deployed at locations close to the first-wave victim so the victim's neighboring customer networks remain unaffected. For example, if an upstream AS deploys S/RTBH filters at the egress port that directly connects to the first-wave victim, all other customer networks are not affected. However, in a link-flooding attack, it is too late to deploy a filter at the egress port, and the best practice is to filter attack traffic at its ingress points, which contradicts the example above and renders the method rarely practiced.

Carom Against M1

Since M1 cannot match and filter traffic based on layer-4 information, Carom does not need to craft traffic defined in Sect. 3.3. In other words, each bot can simply blast the same packet at a first-wave victim that employs M1. Instead, Carom employs the traffic dispersion technique (Sect. 3.3) to distribute the attack traffic among all IPs of the first-wave victim, which ensures that the victim has to include take all its IPs offline to mitigate the attack. In the unusual case where S/RTBH is enabled, Carom with IP spoofing can cause the network provider to disable all its clients from reaching the Internet. Carom detects the usage of M1 as follows:

- **RTBH:** Ask geographically distributed vantage points to connect with the IP addresses under attack. RTBH is employed when all such attempts fail to establish.
- **S/RTBH:** The adversary checks whether his/her bots can reach the first-wave victim's sibling networks (i.e., networks that share the same provider). S/RTBH is employed when none of the bots can communicate with the sibling networks.

4.2 Stateless Fine-Grained Mitigation (M2)

M2 includes but is not limited to BGP FlowSpec or BPF with eXpress Data Path (XDP) [17]. They offer fine-grained traffic mitigation, which can introduce fewer second-wave victims than M1. For example, a first-wave victim can ask its upstream network to forward only HTTP(s) traffic to one of its subnets. BGP FlowSpec is a filter dissemination protocol that allows its users to filter traffic by layer 3 and 4 packet header values. Its filters are often implemented in hardware routers that rely on low-latency memory (i.e., CAM/TCAM) for packet matching and filtering. BPF with XDP is an efficient software-based filtering solution. Specifically, XDP allows BPF to filter packets in kernel space before the kernel’s network stack constructs the packets.

While each filter in M2 generates virtually no second-wave victims, M2 is constrained by its filtering capacity, which forces it to use a limited number of filters to mitigate DDoS attacks. For example, a first-wave victim may only be able to deploy one thousand filters at its upstream ASes, which prevents the victim from mitigating an attack that involves thousands of bots. To mitigate the attack, the victim needs to use network prefixes to cover as many bot IPs as possible.

Carom Against M2

A first-wave victim may leverage M2 to employ strict network policies to prevent unsolicited traffic from congesting its network link. Carom generates attack packets selectively as described in Sect. 3.3 and disperses the attack packets as described in Sect. 3.3. The first technique allows Carom to bypass the network policies (if any) to ensure its traffic can reach the first-wave victim. The latter technique forces the first-wave victim to deploy more filters or use coarse-granular filters. The problem worsens when Carom employs IP spoofing: A first-wave victim can disconnect itself from the Internet with automated mitigation under the limited filter capacity.

Carom can detect M2 with the help of vantage points that are not part of the attack. Specifically, M2 is engaged if the vantage points can establish connections with the first-wave victim IPs during the attack.

4.3 Stateful and Stateless Mitigation (M3)

Stateful mitigation solutions are deployed in-line with the production traffic. They rule out traffic that does not follow transport or application protocols with little to no collateral damage. For example, `conntrack` drops TCP packets with incorrect states (e.g., sequence numbers). A web application firewall (WAF) may use a reCAPTCHA-like system to prevent bots from reaching first-wave victims directly. Note that bots may still bypass reCAPTCHA if they have ample computational resources. These solutions face performance issues at upstream networks to withstand large-scale attacks; they require immense computational and I/O resources. Indeed, during a large-scale attack, we often see under-provisioned

stateful solutions (43% reported DDoS attacks) to cause a denial of service on the protected networks [33].

A network can combine both stateless and stateful mitigation solutions to remove attack traffic. For example, the network may first use a stateless mitigation solution to remove obvious attack traffic and have the stateful mitigation solution to handle the remaining traffic. However, only DPS providers have the resources to deploy stateful solutions to absorb large-scale link-flooding attacks in practice.

Carom Against M3

The attack strategy against M3 is identical to the strategy against M2 (Sect. 4.2), and the explanation is as follows. An attack can involve 10s or even 100s of thousands of bots; these bots force M3 to track a devastating amount of flows. Mian et al. [30] and Corin et al. [13] show that the throughput of a software-switched network is severely reduced with only *thousands* of traffic filters. Depending on the packet volume, packet rate, and how packets are processed in software, the throughput reduction ranges from 30% to 80%.

Furthermore, with IP spoofing, Carom can render stateful mitigation solutions fruitless: IP spoofing not only allows bots to spoof addresses but to **bypass stateful firewalls for an extended period**. Specifically, a stateful firewall must allow the very first packet of a connection (i.e., SYN packet of a TCP connection) to reach its destination IP address. Only then can the firewall drop the same subsequent packets from the packet source. In other words, should each bot only sends *unique* SYN packets during the entire attack period, the firewall will forward all the attack packets.

5 Evaluation

5.1 Overview

The First-Wave Victims of Carom: Since the point of Carom is to put *multiple* first-wave victims in a mitigation dilemma, we study how many first-wave victims an adversary can disconnect simultaneously given different *magnitudes of attack power* and *attack scenarios*. (Note that the main first-wave victims are tier-3 ASes; they *source* information but rely on network links that ordinary DDoS attacks can overwhelm.) We approach the problem from three aspects as follows: **First**, we demonstrate how to infer the ASes that are protected by DPS providers; we developed a tool to track if an AS is protected by a DPS provider in near real time. Carom can rely on this tool to determine whether it should spend its attack power on the AS or not. **Second**, we use collected network information to infer the link capacity of ASes. A high inference accuracy allows Carom to better estimate the required attack bandwidth to overwhelm a group of first-wave victims. While an attacker can target a first-wave victim with a gradually increasing attack volume to estimate the required amount of

asn	providers	peers	siblings	country	org_name	info_type	info_traffic	info_ratio	info_scope
3043	0	4	0	US	Amphibian Me...	NSP	20-50Gbps	Balanced	Regional
3058	1	0	84	RU	Federal Stat...				
3061	2	0	6476	US	ProvDotNet L...	NSP		Not Disclosed	North America
3064	2	0	11149	US	Affinity Int...	Content	1-5Gbps	Heavy Outbound	Global

Fig. 2. A snapshot of the AS information database (with a reduced number of columns).

attack traffic of each first-wave victim more accurately, the attacker can leverage the link capacity inference to reduce reconnaissance time. **Lastly**, with the bandwidth inference result, we show the relationship between the magnitudes of attack power and the first-wave victims an adversary can disconnect. Specifically, we show the various types of possible victims with and without the moving target attack technique.

The Second-Wave Victims of Carom: Each first-wave victim is in a mitigation dilemma - it either: (1) does not mitigate the attack and suffer a poor network goodput or (2) mitigates the attack and (un)intentionally introduces second-wave victims (i.e., the IP addresses and ASes who do not carry any DDoS traffic). To make an informed decision, a first-wave victim needs to quantify the second-wave victims given its mitigation model.

For obvious ethical and legal reasons, we cannot launch real DDoS attacks against networks on the Internet. Therefore, we built a simulation to examine the scale of the second-wave victims under different magnitudes of attack power and mitigation resources. The simulation leverages real-world network measurement data and our DDoS mitigation survey results to synthesize the botnets and each first-wave victim’s defense model.

5.2 The First-Wave Victims of Carom

Experiment Setup. This section introduces our primary datasets and how we processed them to facilitate our understanding of the potential first-wave victims. First, we use the CAIDA AS relationship dataset [9] to find tier-3 ASes (i.e., ASes with no customer networks) on the Internet. For each tier-3 AS, we maintain counters to track its providers, peers, and sibling ASes. We then augment each AS number (ASN) with CAIDA prefix to ASN [7], CAIDA AS organizations [8], and PeeringDB datasets [38]. As a result, we built an AS information database, as demonstrated in Fig. 2. Each row is identified by an ASN and contains the information about the AS (e.g., neighbor AS counters, organization name, geographic location, traffic level, network type, and its network prefixes). Second, we collect and parse BGP updates from a total of 61 BGP route collectors from different geographic locations (using BGPKIT [4]) and monitor for the BGP announcements originated from major DPS providers [14]. We then manually select the corresponding ASNs of the major DPS providers using the AS information database above.




Fig. 3. The CDF of the number of DPS subscribers.

The Subscribers of DPS Providers. As part of the Carom design, it does not waste its attack power on ASes with over-provisioned resources, (i.e., the ASes who pay for anycast-based DDoS defense offered by major DPS providers), we developed a tool to conduct a week-long BGP-based measurement study to find such ASes. In other words, the study is to find the ASes that were under DDoS attacks and then protected by DPS providers. Note that anyone can use the tool and live BGP updates to track ASes under DPS protection in near real time.

Figure 3 applies the cumulative distribution function (CDF) to the number of DPS subscribers in a seven-day time window. It indicates how busy the DPS providers are given a time frame; we find the DPS providers initiated DDoS mitigation efforts to protect 46 ASes within a seven-day time window. Note, in the first day, we captured over 50% of all the protected ASes in the entire time window. Subsequently, we captured only 3 to 4 newly protected ASes each day. This implies that there is a limited set of ASes that are constantly protected by DPS providers.

The Accuracy of Link Capacity Inference. To know who are the potential first-wave victims in Carom, we infer the link bandwidth of the ASes on the Internet and find the ASes whose links can be fully disconnected by a botnet. At the time of this writing, 7,239 tier-3 networks disclosed their traffic level information on PeeringDB while there are more than 61,300 tier-3 networks on the Internet. In this work, we assume the traffic level of a network is 50% of the network’s link bandwidth. E.g., if a network’s traffic level is 5Gbps, then we assume its link capacity is 10Gbps. The assumption is a common convention accepted by the network community [16].

Because we do not have the resources to manually ask for each of the remaining networks about its network capacity, we infer each network’s traffic level instead. We used *scikit-learn* [5] and applied three classification methods: *decision tree*, *k-nearest neighbors (KNN)*, and *random forest*, on the features available




Fig. 4. The traffic level inference accuracy of three classification algorithms.

in our AS information database to infer the network capacity of each AS. Specifically, the selected AS features are *ASN*, *provider count*, *sibling count*, *peer count*, and *IP count*. We reduced 16 traffic levels available in PeeringDB to 4 traffic levels. To train and test each classifier, we used a 70%/30% data split for training and testing. We then evaluate each classifier 100 times with random states.

Figure 4 shows that the three classifiers can infer traffic levels with at least 77% median accuracy. Out of the three classification algorithms, KNN provides the best overall accuracy on our dataset. We use the trained KNN model to infer each AS’s link capacity for the remaining studies. Note that the inference accuracy can vary significantly as networks update their traffic levels on PeeringDB. For example, in 2018, Smith et al. [42] reported a 90% inference accuracy of link capacity of transit ASes using decision tree. The discrepancies in inference accuracies is due to the updated information in PeeringDB and the type of ASes that we are evaluating.

Attack Power vs. First-Wave Victims. With the link capacity inference results, we can now study the number of first-wave victims an adversary can introduce given different magnitudes of attack power with Carom. To maximize the number of first-wave victims, we first rank all tier-3 ASes by their link capacity from low to high, and then apply Carom against the ASes from the lowest end. For each first-wave victim, we use attack traffic that is worth 150% of the victim’s link capacity to disconnect it. Because we cannot launch actual DDoS attacks to measure the required attack bandwidth for each first-wave victim, we choose a constant factor (i.e., 150%) to compensate for the attack traffic transmission loss due to various network conditions such as early link congestion. The study contains four Carom profiles that ranges from one attack group to four attack groups. When Carom only has one attack group, Carom focuses all its attack power on the same set of first-wave victims. Meanwhile, when Carom has four attack groups, Carom rotates its attack power among 4 sets of first-wave victims, which introduces more first-wave victims.

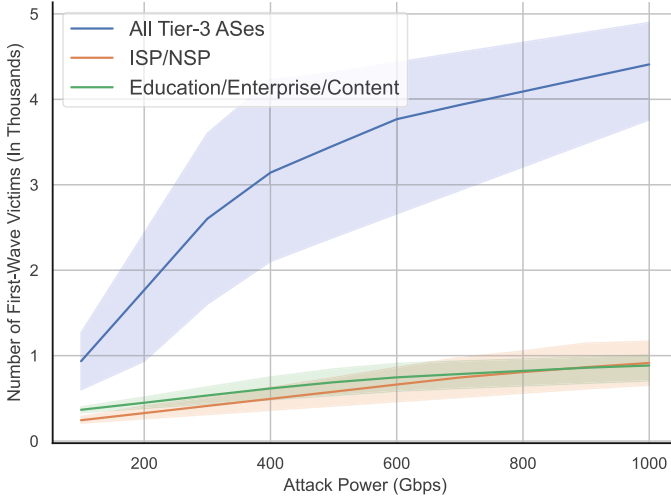


Fig. 5. Estimated first-wave victim ASes.

Figure 5 illustrates the numbers of potential first-wave victims by their network types. Specifically, the types are (1) tier-3 ASes, (2) Internet/network service providers (ISP/NSP), and (3) education/enterprise/content networks. Due to the limitation of PeeringDB dataset, only a subset of ASes disclosed their network types. Therefore, we expect the results of the latter two types to change as more ASes disclose their network types. With just 100 Gbps of attack throughput and Carom that attacks four groups of ASes, an adversary can attack up to 295 ISP/NSP ASes. Worse, the adversary can disconnect over 1,436 first-wave victims when he/she attacks any tier-3 ASes. When an adversary poses a botnet that can deliver 1Tbps attack throughput, it can periodically disconnect over 5,000 ($\approx 8.2\%$) tier-3 ASes with Carom that attacks four groups of tier-3 ASes. This study is an estimation of the potential number of first-wave victims. It does not consider the mitigation strategies each victim applies. We consider mitigation models in the evaluation of second-wave victims.

Figure 6 shows the number of first-wave victim IP addresses affected by the same attacks above. For example, with the 100 Gbps attack above, the 295 ISP/NSP first-wave victims own a total of 1.6M IP addresses. If we map each IP address to a household, the 100 Gbps DDoS attack can cause 1.6M homes to experience frequent Internet disconnection when the first-wave victims do not react to the attack. To put the number into perspective, such an attack could potentially disconnect more households than the Los Angeles city (1.38 million households) according to the U.S. Census Bureau [10]. With the same attack throughput, the attacker can disconnect up to 7,504,129 IP addresses when attacking any tier-3 ASes.




Fig. 6. Estimated first-wave victim IP addresses.

5.3 The Second-Wave Victims of Carom

Experiment Setup. The experiment is a simulation-based study that evaluates the second-wave victims caused by the first-wave victims who employ M2 under various attack scenarios. (Note that the attack strategy against M3 is the same as M2.) It is unnecessary to evaluate ASes that employ M1 (Sect. 4.1): Such an AS blocks all incoming traffic from all other ASes on the Internet; the second-wave victims are the entire Internet to the AS.

Because M2 employs fine-grained filters that matches source IPs of the bots in an attack, the source IP distribution affects what filters are generated hence critical to the experiment. Therefore, we built a set of synthesized Mirai botnets based on a real-world DDoS incident report [3]. Specifically, given a botnet size and the country distribution of the botnet, we map a set of bots to each country proportionally. Each bot is then assigned with an IP address of its belonging country and its uplink bandwidth using the Ookla’s speedtest dataset [36]. The sizes of the synthesized botnets range from 5K to 100K.

We assume each first-wave victim has a perfect DDoS detection accuracy which puts Carom in a disadvantageous situation. With the perfect detection result, each first-wave victim also employs a filter generation process to optimize for the attack traffic coverage within a fixed budget. In this study, the shortest filter prefix length the filter generation process can generate is $/8$, each first-wave victim in the simulation allows attack traffic occupying at most 30% of its link capacity, and for the sake of simplicity, we do not allow bots to spoof IP addresses to increase the mitigation difficulty.




Fig. 7. The second-wave victim IP addresses under varying botnet sizes and filter budgets.

Botnets vs. Second-Wave Victims. Figure 7 shows the numbers of second-wave victim IPs under different attacks and filter budgets (i.e., how many fine-grained filters a first-wave AS can deploy). In all attacks, we first see the number of second-wave victim IPs gradually increases as we increase the filter budgets initially. This is because the available filter budgets are too small for mitigating the attacks even we filter attack traffic using /8 prefix filters. Using the 5K-bot attack as an example, as we continue to increase the filter budget, the number of second-wave victim IPs reaches 1.6 billion, which corresponds to 43.2% of the public IP space or 34K to 42K of second-wave victim ASes, at its peak with 100 filters. The affected number of IPs then starts to decline as the filter budget continues to grow. With a 1K filter budget (20% of the botnet size), the number of second-wave victim IPs is reduced to $\approx 670\text{M}$, which corresponds to approximately 3K second-wave victim ASes.

Figure 8 demonstrates how the filter budgets affects the prefix lengths of generated filters under the same 5K-bot attack. We see that the filters become more specific as we increase the filter budget. With 2K filters, we no longer produce second-wave victims, this is because the rule generation process allows attack traffic to occupy at most 30% of an AS’s link capacity. In other words, all filters are generated at /32-level (individual IPv4 addresses).

5.4 Evaluation Summary

In this section, we demonstrate the feasibility in choosing first-wave victims from two aspects: First, to avoid attacking networks that are subscribed to DPS providers, we created and evaluated a tool that allows an attacker to track ASes under DPS protection in near real time. We found that over a seven-day period,

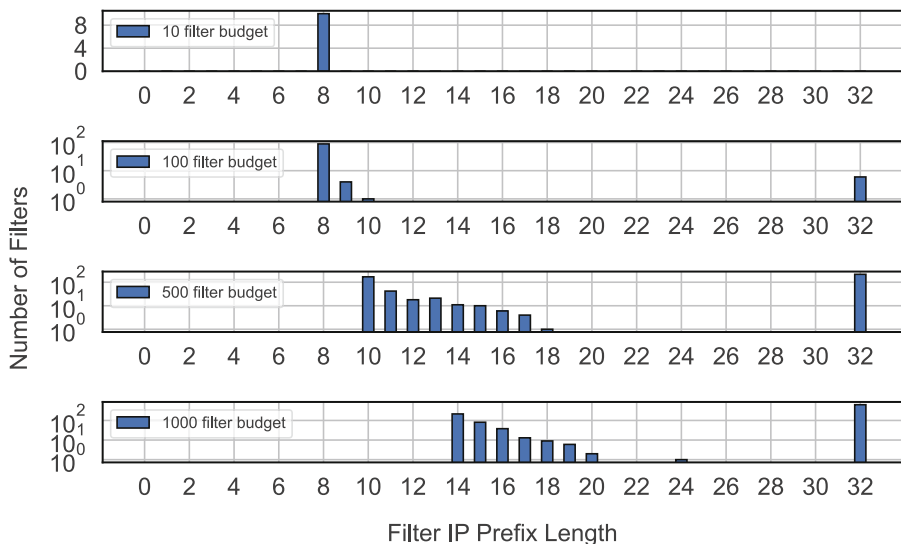


Fig. 8. The distribution of filter sizes with different filter budgets and 5K bots.

there is a small set of ASes under constant DPS protection. Second, to increase the probability that each set of first-wave victims can be fully disconnected, an attacker may attempt to infer link capacity of the potential first-wave victims. We proposed an inference method that applies classification techniques to publicly available data from PeeringDB. We show that an attacker can infer traffic levels with at least 77% median accuracy.

We then analyzed the number of first-wave victims an attacker may disconnect given different magnitudes of attack power, and found that with merely 100 Gbps of attack power attacking a total of four sets of first-wave victims, an attacker could disconnect up to 295 ISPs/NSPs (which total 1.6M IPs) or 1,436 tier-3 ASes (which total 7.5M IPs).

Lastly, we analyzed the impact that deploying stateless fine-grained filters (M2) has on second-wave victims. We show that the number of second-wave victims in fact (counterintuitively) increases as the filter budget increases to a certain point, and decreases after that point. For example, with a 5K-bot and budgets of 10 filters, 100 filters, and 1000 filters, the number of second-wave victims are 167M IPs, 1.5B IPs, and 70M IPs, respectively.

6 Conclusion

Since the first documented attack over 20 years ago, the Internet continues to face severe large-scale DDoS attacks. While some ASes can overprovision their computing and network resources to fight off large-scale attacks (e.g., by subscribing to a DPS provider), most ASes rely on insufficient solutions to mitigate

such attacks. The two-wave attack with **collateral damage of millions**, or *Carom*, exploits this mitigation insufficiency to wreak havoc on wide swaths of the Internet. Specifically, the attack leverages main ideas from prior attack research (i.e., pulsing and Crossfire) and the practical attack and defense constraints to impose a mitigation dilemma on its first-wave victims, which may then lead to a significant number of second-wave victims. Through experimentation on real-world data, we show that such an attack can simultaneously disconnect hundreds of ISP/NSP ASes or thousands of tier-3 ASes with a mere 100 Gbps attack. Furthermore, if these first-wave victims employ stateless fine-grained filtering to mitigate the attack, depending on the filter budget, they may end up disconnecting themselves from nearly 43.2% of the entire usable IPv4 space.

Given the potential devastation the Carom attack can unleash on today's Internet, we hope this work can raise the awareness of the network community and ultimately help spearhead the development and deployment of adequate DDoS mitigation solutions. For example, Carom can leverage IP spoofing to impose a higher mitigation capacity demand on the first-wave victims. Thus, a critical component to defend against it is to prevent spoofed packets from congesting a first-wave victim's network link. While initiatives such as CAIDA's Spoofer project [6] and MANRS [28] are calling network operators to implement ingress/egress filtering to prevent IP spoofing, a much higher AS participation rate is needed for an effective IP spoofing prevention. The filter budget of a DDoS mitigation directly affects the number of second-wave victims, too. The existing DDoS mitigation research has two complementary ideas: One is the machine-level filtering using programmable switches with efficient use of low-latency memories (e.g., TCAM/CAM) [23, 46], or using commodity servers and efficient packet processing pipelines such as XDP [17]. The other is the Internet-level filtering which is about building a filter distribution system so the participating ASes can each contribute some traffic filtering capacity.

References

1. Akamai: Prolexic Routed (2021). <https://www.akamai.com/us/en/products/security/prolexic-solutions.jsp>
2. Alcoz, A.G., Strohmeier, M., Lenders, V., Vanbever, L.: Aggregate-based congestion control for pulse-wave DDoS defense. In: Proceedings of the ACM SIGCOMM 2022 Conference (2022)
3. Antonakakis, M., April, T., Bailey, M., Bernhard, M., Bursztein, E., et al.: Understanding the Mirai botnet. In: 26th USENIX Security Symposium (USENIX Security 2017) (2017)
4. BGPKIT: BGP Data Analysis Tool Kit (2021). <https://github.com/bgpkit>
5. Buitinck, L., et al.: API design for machine learning software: experiences from the scikit-learn project. In: ECML PKDD Workshop: Languages for Data Mining and Machine Learning (2013)
6. CAIDA: State of IP Spoofing (2019). <https://spoofer.caida.org/summary.php>
7. CAIDA: Routeviews Prefix to AS mappings Dataset for IPv4 and IPv6 (2021). <https://www.caida.org/catalog/datasets/routeviews-prefix2as/>

8. CAIDA: The CAIDA AS Organizations Dataset (2021). <http://www.caida.org/data/as-organizations>
9. CAIDA: The CAIDA AS Relationships Dataset (2021). <http://www.caida.org/data/active/as-relationships>
10. Census Bureau: Census Bureau QuickFacts (2022). <https://web.archive.org/web/20220107213155/https://www.census.gov/quickfacts/fact/table/losangelescalitycalifornia,US/VET605219>
11. Cimpanu, C.: ‘Carpet-bombing’ DDoS attack takes down South African ISP for an entire day (2019). <https://www.zdnet.com/article/carpet-bombing-ddos-attack-takes-down-south-african-isp-for-an-entire-day/>
12. Cloudflare: Cloudflare Magic Transit (2021). <https://www.cloudflare.com/magic-transit>
13. Corin, R.D., Costanzo, A., Callegati, F., Siracusa, D.: Methods and techniques for dynamic deployability of software-defined security services. CoRR (2020)
14. Gartner: DDoS Mitigation Services Reviews and Ratings (2021). <https://www.gartner.com/reviews/market/ddos-mitigation-services>
15. Guirguis, M., Bestavros, A., Matta, I.: Exploiting the transients of adaptation for ROQ attacks on internet resources. In: Proceedings of the 12th IEEE International Conference on Network Protocols (ICNP) (2004)
16. Hassidim, A., Raz, D., Segalov, M., Shaqed, A.: Network utilization: the flow view. In: IEEE INFOCOM (2013)
17. Høiland-Jørgensen, T., et al.: The express data path: fast programmable packet processing in the operating system kernel. In: Proceedings of the 14th International Conference on Emerging Networking EXperiments and Technologies (2018)
18. Imperva: DDoS Protection for Networks (2021). <https://www.imperva.com/products/infrastructure-ddos-protection-services/>
19. Kang, M.S., Lee, S.B., Gligor, V.D.: The crossfire attack. In: 2013 IEEE Symposium on Security and Privacy (2013)
20. Ke, Y.M., Chen, C.W., Hsiao, H.C., Perrig, A., Sekar, V.: CICADAS: congesting the internet with coordinated and decentralized pulsating attacks. In: ASIA CCS 2016 (2016)
21. Kumari, W., McPherson, D.: Remote triggered black hole filtering with unicast reverse path forwarding (uRPF). Technical report (2009)
22. Kuzmanovic, A., Knightly, E.W.: Low-rate TCP-targeted denial of service attacks: the shrew vs. the mice and elephants. In: SIGCOMM (2003)
23. Liu, Z., et al.: Jaqen: A high-performance switch-native approach for detecting and mitigating volumetric DDoS attacks with programmable switches. In: USENIX Security Symposium (2021)
24. Lu, N., Zhang, J., Liu, X., Shi, W., Ma, J.: STOP: a service oriented internet purification against link flooding attacks. *IEEE Trans. Inf. Forensics Secur.* **17**, 938–953 (2022). <https://doi.org/10.1109/TIFS.2022.3152406>
25. Luckie, M., Beverly, R., Koga, R., Keys, K., Kroll, J.A., Claffy, K.: Network hygiene, incentives, and regulation: deployment of source address validation in the internet. In: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (CCS) (2019)
26. Luo, X., Chang, R.K.: On a new class of pulsing denial-of-service attacks and the defense. In: Proceedings of the Network and Distributed System Security Symposium (2005)
27. Majkowski, M.: The real cause of large DDoS - IP Spoofing (2018). <https://blog.cloudflare.com/the-root-cause-of-large-ddos-ip-spoofing/>

28. MANRS: MANRS for Network Operators (2022). <https://www.manrs.org/isps>
29. Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., McPherson, D.: Dissemination of flow specification rules. Technical report (2009)
30. Miano, S., Bertrone, M., Rizzo, F., Bernal, M.V., Lu, Y., Pi, J.: Securing Linux with a faster and scalable iptables. In: SIGCOMM CCR (3) (2019)
31. Nawrocki, M., Blendin, J., Dietzel, C., Schmidt, T.C., Wählisch, M.: Down the black hole: Dismantling operational practices of BGP blackholing at IXPs. In: Proceedings of the Internet Measurement Conference, IMC 2019 (2019)
32. NETSCOUT: Arbor Cloud DDoS Protection Services (2021). <https://www.netscout.com/product/arbor-cloud>
33. NETSCOUT: NETSCOUT's 14th Annual Worldwide Infrastructure Security Report (2020). <https://www.netscout.com/report>
34. Neustar: UltraDDoS Protect Mitigation Service (2021). <https://www.home.neustar/resources/product-literature/ddos-mitigation-service-product-literature>
35. Nvidia: Cumulus Linux 4.2 User Guide (2022). <https://web.archive.org/web/20220124023544/https://docs.nvidia.com/networking-ethernet-software/cumulus-linux-42/System-Configuration/Netfilter-ACLs/#hardware-limitations-on-number-of-rules>
36. Ookla: Speedtest by Ookla Global Fixed and Mobile Network Performance Map Tiles (2021). <https://github.com/teamookla/ookla-open-data>
37. Park, J., Nyang, D., Mohaisen, A.: Timing is almost everything: realistic evaluation of the very short intermittent DDoS attacks. In: Annual Conference on Privacy, Security and Trust (PST) (2018)
38. PeeringDB: PeeringDB. <https://www.peeringdb.com> (2021)
39. Radware: Cloud DDoS Protection Service (2021). <https://www.radware.com/products/cloud-ddos-services2>
40. Rasti, R., Murthy, M., Weaver, N., Paxson, V.: Temporal lensing and its application in pulsing denial-of-service attacks. In: 2015 IEEE Symposium on Security and Privacy (2015)
41. Shan, H., Wang, Q., Pu, C.: Tail attacks on web applications. In: ACM SIGSAC Conference on Computer and Communications Security (CCS) (2017)
42. Smith, J.M., Schuchard, M.: Routing around congestion: defeating ddos attacks and adverse network conditions via reactive BGP routing. In: IEEE Symposium on Security and Privacy (2018)
43. Sommese, R., et al.: MAnycast2: using anycast to measure anycast. In: Proceedings of the ACM Internet Measurement Conference. IMC 2020 (2020)
44. Studer, A., Perrig, A.: The Coremelt attack. In: Backes, M., Ning, P. (eds.) ESORICS 2009. LNCS, vol. 5789, pp. 37–52. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-04444-1_3
45. Sun, H., Lui, J.C.S., Yau, D.K.Y.: Defending against low-rate TCP attacks: dynamic detection and protection. In: Proceedings of the 12th IEEE International Conference on Network Protocols, ICNP 2004 (2004)
46. Zhang, M., et al.: Poseidon: mitigating volumetric DDoS attacks with programmable switches. In: Proceedings of NDSS (2020)