Shield: DoS Filtering Using Traffic Deflecting

Erik Kline Alexander Afanasyev Peter Reiher
Laboratory for Advanced Systems Research
UCLA Computer Science Department
{icebeast, afanasev, reiher}@cs.ucla.edu

***Abstract*—Denial-of-service (DoS) attacks continue to be a major problem on the Internet. While many defense mechanisms have been created, they all have significant deployment issues. This paper introduces a novel method that overcomes these issues, allowing a small number of deployed DoS defenses to act as secure on-demand shields for any node on the Internet. The proposed method is based on rerouting any packet addressed to a protected autonomous system (AS) through an intermediate filtering node—*a shield*. In this way, all potentially harmful traffic could be discarded before reaching the destination. The mechanisms for packet rerouting use existing routing techniques and do not require any kind of modification to the deployed protocols or routers. To make the proposed system feasible, from both deployment and usage points of view, traffic rerouting and outsourced filtering could be provided as an insurance-style on-demand service.**

***Index Terms*—DDoS, Filtering, IP Anycast, BGP, Traffic deflection**

I. Introduction

As the Internet grows, malicious users continue to find intelligent and insidious ways to attack it. Many types of attacks happen every day, but one particular kind—denial-of-service (DoS) attacks—remain the most common, accounting for more than a third of all malicious behavior on the Internet in 2011 [1]. The main goal of these attacks is literally to deny some or all legitimate users access to a particular Internet service, harming the service as a whole. In the extreme case, when the attack is aimed at the core Internet infrastructure (e.g., attacks on the root DNS servers [2]), the whole Internet could be jeopardized. There is a clear need for comprehensive, cheap, and easily deployable DoS protection mechanisms.

Attackers may have different motivations (extortion, vengeance, or simple malice) and the goal of a DoS attack could be achieved in many ways. Thus, there is a wide variety of attack methods available [3] and a growing number of proposed defense mechanisms to stop or mitigate them. Many of the proposed DoS defenses are both clever and potentially effective [4]. However, the most common question with DoS defenses is how to deploy them.

Some defenses require deployment in core routers [5], but the tier 1 ASes that own these routers have little incentive to do so. The economic model of all transit providers, including tier 1 providers, consists of charging for the amounts of forwarded traffic. Thus, such providers are extremely cautious with any kind of filtering, as they risk the loss of money or even customers. In addition, unless fully deployed by every major ISP, core defenses generally provide very limited protection.

Other defenses require deployment on the edges of autonomous systems (ASes), hoping to catch malicious traffic before it goes to the outside world. Unfortunately, it is difficult to detect a DoS attack, and especially a distributed DoS attack, at the source. These attacks can have a multitude of carriers (e.g., infected nodes) that generate small amounts of virtually undetectable malicious traffic, making edge defenses useless. Moreover, there are basically no incentives for ASes to deploy such filtering mechanisms. Edge filters prevent attacks from leaving an AS, but provide little protection for the filtering AS. Until everybody implements a similar defense, the ASes that deploy filters will not gain anything; their users will still be vulnerable. This results in the current status quo, where an AS usually allows almost any traffic with minimal or no filtering at all. Ingress filtering [6] is the most common type of available and widely deployed anti-attack measure [7]. Unfortunately, this filtering defends only against IP spoofing and provides little help for deflecting other types of attacks.

Finally, defenses can be deployed near every victim in the form of traffic analysis tools, firewalls, and anti-virus software. In general, to be able to absorb an attack, these defenses require the victim to be highly over-provisioned (i.e., the network should have a large enough bandwidth, or the content should be sufficiently replicated). Clearly, not every possible victim has the resources to over-provision, which nullifies the effect of such defenses.

In this paper we propose a promising security model for the Internet that can leverage virtually any kind of previously proposed mechanism, without facing the deployment problem. Our solution relies on the existing

978-1-4577-1394-1/11/$26.00 ©2011 IEEE

routing techniques—BGP routing [8], IP Anycast [9], [10], IP tunneling [11], and others—to divert traffic during a DoS attack from a direct route to a route that contains special DoS filters (Figure 1). When an attack occurs, all packets destined for the target AS are forced by means of the routing system to pass through specially deployed filtering nodes, which we call *shields*. During an attack, shields (1) pretend to be a valid origin for the attacked prefix (in the same way IP Anycast or a multiple origin AS works), (2) perform a dedicated traffic analysis and filtration of malicious traffic, and (3) deliver all legitimate traffic to a real destination.

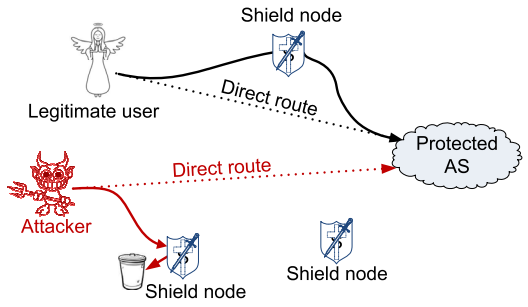

Fig. 1. Diverting traffic flow from a direct route to pass through filtering nodes (shields)

One of the key elements in the proposed solution is the on-demand nature of the filtering mechanisms. The shields will divert traffic and perform traffic analysis and filtering for a protected IP prefix or AS *only* during an active attack. This helps resolve the contradiction between the requirement to protect and that of not disturbing normal network functionality by either increasing delays or reducing available bandwidth.

An on-demand DoS attack defense solution is beneficial because the occurrence of a future DoS attack is difficult or even impossible to predict. Not every Internet server or AS is exposed to a DoS attack at the same time; any attack, whatever its duration, is temporary and will eventually cease. Thus, always-on solutions will waste resources by analyzing harmless traffic most of the time. Our shields solve this problem by allowing an effective insurance-style sharing of defense resources among a large number of Internet users.

The proposed system is a method to deploy DoS defense filters; it is not a design of a DoS defense filter. Many existing high-quality filters could be efficiently implemented on the shield nodes, providing first-class protection for participating parties.

II. Shield Design

Protection against DoS attacks is a very complex and contradictory problem. One of the most important aspects of DoS defense is where such a defense (e.g., traffic analyzer and filter) should be deployed. This decision largely defines deployment and exploitation cost, as well as the level of protection achieved. Defenses could be implemented locally at the server or AS level. Unfortunately, there are serious drawbacks to the local deployment model. Local defenses may only give a limited protection to the outside world from the attackers inside an AS (a node in a botnet may not generate a detectable volume of malicious traffic) an a limited ability to mitigate certain types of low-rate application-targeted attacks. As a result, until all ASes deploy the defenses, any outside attacker could successfully perform a flood-based DoS attack by overwhelming capacities of local links.

An alternative deployment strategy is implementing DoS defenses inside the part of the Internet core that sees and forwards virtually all traffic—inside the 50 largest ASes that forward more than 95% of all Internet traffic [12]. Although core defenses mitigate to some degree the problem of partial deployment and may provide a comprehensive protection against any volume of DoS attacks, they are unlikely to be implemented and deployed in the foreseeable future. Core ASes move traffic as fast as possible (a couple of nanoseconds spent on each forwarded packet) and will not deploy any service that may harm throughput.

Our system takes a different approach. Instead of deploying defenses along the direct route (e.g., at the core routers or AS edges), our system outsources defenses to dedicated shields that can be shared among a large number of Internet users. This outsourcing could be accomplished in a number of ways, but there are two basic components: traffic deflection towards the shield nodes and final delivery of legitimate traffic from shields to the true destinations. It could be implemented within the existing global routing system, using traffic trapping and black-holing for route deflection towards shields, stretched-path forwarding, source routing, and different types of tunneling for final delivery of legitimate traffic to the real destination. In the following sections we discuss these and other implementation strategies, their feasibility, and the level of DoS defenses they could provide.

A. Traffic redirection

If there are so many problems with deployment of defense mechanisms along the direct route, why not

alter the routing to make traffic go through dedicated filters located elsewhere? Instead of moving the filters to the traffic, move the traffic to the filters. The question here is how to divert traffic and make it flow in the direction of the filter, not in the direction of an AS under attack. Our answer to this question contains two somewhat related solutions, applicable in the existing global routing infrastructure: *traffic trapping* and *traffic black-holing*.

The first technique includes deployment (physical or tunnel-based) of shields at important topological points in the Internet, such as Internet Exchange Points (IXPs). During an attack, appropriately deployed shields can start announcing shorter paths for the protected prefix (e.g., by forging the AS-PATH attribute of the BGP protocol [8]), effectively trapping the traffic for this prefix (Figure 2).

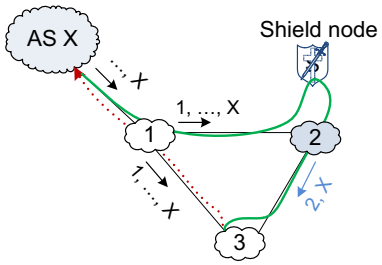

Fig. 2. Traffic trapping by shortening the AS-PATH attribute: node 2 announces to node 3 a trap route {2, X}

Unfortunately, there are several shortcomings that seriously limit applicability of traffic trapping. First, AS-PATH altering may be harmful and introduce routing loops, though additional restrictions and cautious per-neighbor advertisement decisions can significantly reduce, if not eliminate, this potential damage. Second and probably most critical, is the requirement for substantial deployment. To be effective, shields should be deployed on every major IXP—otherwise, the efficiency and level of DoS protection will be significantly jeopardized. Finally, traffic trapping does not really provide a fully comprehensive filtering solution. If an attacker is lucky enough to be co-located or be very close, it can still send malicious traffic directly to the victim, bypassing all traffic traps.

Luckily, traffic trapping is not the only method that can be employed for traffic rerouting. There is a much less restrictive alternative that allows any level of deployment (i.e., from one to millions of shields) at any place on the Internet. Instead of shortening AS-PATHs, during an active attack, shields can become the only origins (from a BGP's point of view) for the attacked prefix. The exact procedure could be as follows. When an AS detects an attack for a prefix, it informs shields that they need to enable filtering for the prefix. Receiving this solicitation, shields, on behalf of the AS, begin announcing the attacked prefix as if it is their own, and the AS itself withdraws all original routes for the prefix in question (Figure 3). This way, the AS has effectively black-holed traffic for the attacked prefix, which is an extreme case of traffic trapping.

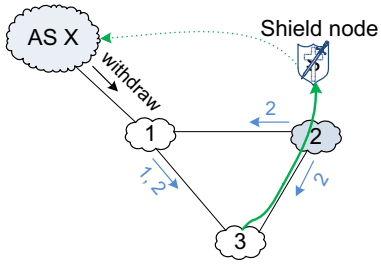

Fig. 3. Traffic black-holing by anycasting protected prefix: origin withdraws the real route and shields act as route origins

In both cases, once shields determine that the attack has ceased, they release filtering resources and revert back to the original routing. In the traffic black-holing case, they also notify the protected AS that filtering will be terminated soon. This allows AS to return its original routing announcements to the global routing system and proceed with receiving data over the standard shortest paths. In special circumstances, operators are able to manually terminate protection service, effectively returning routing to the pre-protection state.

B. Legitimate traffic delivery

An astute reader may have realized a problem with this scheme. While the traffic has been diverted from the original route and effectively filtered, how does the destination actually get the filtered (i.e., legitimate) traffic? In the case of traffic trapping the answer to this question is quite easy. Shields can internally maintain the real shortest-path next hop and forward all legitimate traffic to it. Traffic could get trapped many times along the route, but it will always be getting closer and closer, and eventually will get to the final destination.

The matter is not that simple for the black-holing case. Since the original prefix was withdrawn, shields cannot simply place the packets on the pipe and expect them to get to the host. In fact, since shields are advertising the prefix as their own, the packets will get routed back to the same or another nearby shield. Therefore, we need an additional mechanism to get the packets from the shields to the true destination.

There is no ideal solution that satisfies all possible exploitation and security requirements, but several alternative schemes offer different implementation complexity, deployment feasibility, and level of DoS protection.

The simplest scheme is to use a secret routable IP prefix for IP-IP tunneling purposes. The ultimate goal is to use such a prefix to tunnel legitimate traffic towards the real destination, without exposing the original prefix to direct attacks. However, this is obviously a single point of failure. If the attacker learns the secret prefix, the destination will be again exposed to direct attacks.

This simple scheme could be improved by coupling with dynamic IP prefixes to make it harder or impossible to discover the destination's Achilles' heel. However, if we require an AS to announce a new secret prefix before traffic black-holing or to change the announced prefix during an active defense, the secret could be easily inferred from BGP updates. To overcome this shortcoming, it is possible to maintain a sufficient set of obscurely announced secret prefixes (i.e., with falsified identities). Each one of these secret prefixes could be used to tunnel traffic during an attack, providing the ability to withdraw any exposed prefix without compromising the defense.

A different approach is to use the source routing option of IP [13]. Before an AS withdraws a route, shields can discover a valid hop-by-hop path and source-route all legitimate packets along this path during the defense period. Unfortunately, this solution also has several limitations. First, the attackers can source route malicious traffic as well, circumventing the defense. Second, the method relies on source routing, which is neither widely deployed [14] nor recommended as permanently enabled for security reasons [15].

A future solution is to use recently proposed advanced routing techniques, such as avoidance routing [16]. On a very high level, avoidance routing allows everybody to specify how requested traffic transits the Internet, such that specific routers and areas can be avoided. With this system in place, the attacked destination will withdraw the prefix globally and then announce it only for routers on direct paths to shields (i.e., avoiding all routers not on a direct path). As a result, all routers on the Internet would route to shields, except the routers on direct paths between shields and the attacked AS. Thus, all traffic of interest would be forwarded to shields, and shields would be able to naturally forward it to the real destination, without any need of supplementary mechanisms. A limitation is that the method does not protect against attacks originating from users or routers that happen to be on a direct path. However, this would be a small fraction of the Internet, greatly reducing the scale of a possible attack. While these routing techniques are still in their infancy, they will be helpful for this problem once deployed.

III. Deployment Incentives

There are many incentives to deploy and use Shield. The first major incentive is that Shield can be deployed literally anywhere on the Internet. This offers a great business opportunity, even for a small startup. The second major incentive is that the system is absolutely agnostic to the filtering mechanism used. Any filtering method could be employed, providing any kind of balance between cost and DoS defense efficiency.

Further, by deploying Shield in multiple strategic topological places on the Internet, the load during attacks will be effectively balanced. That is, each Shield node diverts and filters only a fraction of traffic, and only from closely located sources, as in IP Anycast. Besides load balancing, it affords opportunities to provide multi-grade DoS defenses—small ISPs may be satisfied with a single shield node, while large content producers and ISPs may need all available shields.

The final and probably the most important incentive is that the system is on-demand. On-demand service is a natural way to share resources among a large number of users. Not all clients would get attacked simultaneously, and thus the defense would still be usable without needing to linearly scale resources with the number of users. This leads naturally to an insurance-style business model.

This final point is not just an incentive for defense providers, but also for users of the service. Deploying and maintaining effective DoS defenses (especially somewhere in the Internet core) could be tremendously expensive and not many can afford it, even if it is doable. With our model, businesses could still benefit from first-class DoS defenses, paying a nominal fee for the DoS protection insurance service. Theoretically, this will not only make efficient defenses affordable, but will also significantly reduce the threat level on the Internet. If attacks cannot succeed, there is little interest in initiating them.

IV. Open Questions

The proposed model has several open research questions and potential shortcomings. Here we discuss these questions and some possible solutions, although many issues are still open for research, analysis, and debate.

The first important exercise one must conduct whenever they propose a defense scheme is to think about how the defense itself could be used by an attacker. There are several ways an attacker might try to use shields

maliciously. The first possibility is to launch a very small attack on a target, causing it to request filtering. Since shields require tearing down and establishing new routes, there could be a short period of dysfunction. If the attacker performs such attacks on a periodic basis without carrying out a large-scale attack, shields would be constantly entering and exiting the filtering state. Due to BGP damping [17], routing tables may require minutes to stabilize, which may effectively deny service during these stabilization periods.

One possible fix is to define a dynamic DoS-defense triggering threshold. Such a threshold could not only take into account current traffic patterns, but also a recent history of past defense actions. During periodic low-level attacks the system could persist longer in either the defense or non-defense states, requiring the attackers to devote more resources to the attack. The victim could also wait to withdraw its prefix until the shields have announced it, lowering the time that the victim is in a black hole.

Another possible avenue for attack is to target all the customers of the shielding company simultaneously. The crux of any insurance model is the assumption that all policy holders cannot make claims simultaneously. However, this system differs from car insurance, for example, in one major way. With car insurance, all policy holders would have to simultaneously have car damage to make claims, which is extremely improbable. In our system, if an attacker has sufficient knowledge and resources, it can cause some damage to all policy holders, resulting in massive claims for a DoS defense service.

Mitigating this problem is generally an issue of deployment and service provisioning. The company that runs the service must determine how many customers can be supported simultaneously, client priorities, and potential risks. They must then make a standard cost/benefit analysis on how much more insurance they can sell beyond the available resources.

Another interesting question is the effect of shields on Internet routing. Constantly withdrawing and announcing routes would lead to route flapping, which is highly undesirable. The simplest solution is to limit a maximum rate of withdraws and announcements of all participating nodes. Further, because periodic withdrawing and announcing may result in temporary DoS as described above, we would want to limit the rate at which we use it. Thus, once filtering has been entered, to prevent damage to Internet routing, filtering should not be terminated for some period of time, even if the attack has ceased.

V. Related work

One of the most successful technologies on the Internet is Content Delivery Networks (CDNs), such as Akamai [18]. These technologies allow for scalable content delivery by pushing content out to multiple caches distributed throughout the Internet. When someone requests the content, the request, instead of going to the source, likely goes to a nearby cache. This allows many more customers to be serviced simultaneously than the content provider could natively handle. These networks can also be used for DoS defense in the same manner. When someone launches an attack, the CDN distributes this attack amongst a multitude of CDN nodes, requiring an attacker to take down the entire CDN in order to DoS the target.

Clearly, CDNs and Shield work in a similar manner. In Shield, we also distribute the attack to multiple shield nodes across the Internet. Thus, an attacker needs to take down Shield to take down the victim. A fundamental difference between CDNs and Shield is how traffic is actually handled. In a CDN, the traffic is generally served by a local cache, which works well only for cacheable content. On the contrary, in Shield, traffic is served by the content provider, so it can serve any type of cacheable and non-cacheable traffic, including live video streams and online video games. The primary disadvantage of Shield, compared to a CDN, is that the rerouting phase will need to occur before the defense is fully active. At best, this will provide a minimal latency increase; at worst, the service could be unavailable for a short period. In contrast, a CDN uses its caching technology to serve traffic as fast as possible.

There are a large number of existing DoS technologies, many of which have significant deployment issues. Some focus on IP-spoofing defenses in the core, such as route-based filtering [19] and IDPF [20]. These technologies require routers to keep track of packet characteristics (e.g., incoming interface) and filter mismatches. The solutions have varying levels of success, but they all require deployment on routers, which is a major hurdle. Other technologies, such as TVA [21] and StackPi [22], require routers to modify packets with marks or capabilities to ensure packets come from legitimate hosts.

Some technologies do not validate packets, but simply attempt to reduce the incoming flow. StopIt [23] deploys a new infrastructure service at each AS. When a node is under attack, it sends out a StopIt request. This request is disseminated to routers along the forwarding path and eventually to the source AS. If the source AS is well-behaved, it stops the traffic itself. Otherwise, filtering

nodes along the path will filter the traffic out. Once again, this requires major changes to the core.

There are many technologies that live on the edge, and attempt to filter out attack traffic. D-Ward [24] attempts to detect traffic at the source, and filter out the outgoing attack. As we discussed earlier, detecting at the source can be extremely difficult, and source-based defenses require a large-scale deployment to be effective.

DefCOM [25] shows that a heterogeneous network of filtering nodes, each node looking for different characteristics, is a valuable tool to combat DoS. However, getting the traffic into a DefCOM network is a major challenge—one that could be solved using Shield.

A system with some similarities to ours in SOS [26]. SOS works by creating a perimeter around a protected destination, and only allowing traffic to reach the destination through that perimeter. SOS then uses tunnels, internal routing, and filtering to get legitimate traffic to the destination. It is an always-on system that requires significant infrastructure investments to be effective, so there must be enough incentives and benefits to deploy such an always-on service. In general, always-on DoS defense is applicable only for critical and emergency services. In contrast, Shield is an on-demand system that makes more sense for smaller-scale services that can tolerate minor downtime, without significant investments in a DoS-defense infrastructure.

VI. Conclusion

In this paper we have proposed Shield, a scheme for deploying DoS defenses. Shield works by rerouting traffic to filtering nodes on demand, using either traffic trapping or black-holing techniques, which do not require any modifications to existing routing protocols. These filtering nodes—shields—filter out illegitimate traffic, then forward the legitimate traffic to the destination. We have discussed different mechanisms for forwarding legitimate traffic, as well as addressed possible questions and limitations of our system. Finally, we proposed an insurance-based DoS protection model to encourage deployment and wide use of Shield.

References

[1] Trustwave SpiderLabs, "The Web hacking incident database. Semiannual report. July to December 2010," 2011.

[2] R. Naraine, "Massive DDoS attack hit DNS root servers," *InternetNews.com*, October 2002, http://www.esecurityplanet.com/trends/article.php/1486981/Massive-DDoS-Attack-Hit-DNS-Root-Servers.htm.

[3] J. Mirkovic and P. Reiher, "A taxonomy of DDoS attack and DDoS defense mechanisms," *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 2, pp. 39–53, 2004.

[4] T. Peng, C. Leckie, and K. Ramamohanarao, "Survey of network-based defense mechanisms countering the DoS and DDoS problems," *ACM Computing Surveys (CSUR)*, vol. 39, no. 1, 2007.

[5] E. Kline, M. Beaumont-Gay, J. Mirkovic, and P. Reiher, "RAD: Reflector attack defense using message authentication codes," in *Proceedings of Annual Computer Security Applications Conference (ASAC)*, 2009, pp. 269–278.

[6] P. Ferguson and D. Senie, "Network ingress filtering: Defeating denial of service attacks which employ IP source address spoofing," RFC 2827, May 2000.

[7] R. Beverly and S. Bauer, "The spoofer project: Inferring the extent of source address filtering on the internet," in *Proceedings of USENIX SRUTI*, 2005, pp. 53–59.

[8] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," RFC 4271, January 2006.

[9] C. Partridge, T. Mendez, and W. Milliken, "Host Anycasting Service," RFC 1546, November 1993.

[10] H. Ballani and P. Francis, "Towards a global IP anycast service," in *Proceedings of SIGCOMM*, vol. 35, no. 4, August 2005, pp. 301–312.

[11] D. Farinacci, T. Li, S. Hanks, D. Meyer, and P. Traina, "Generic Routing Encapsulation (GRE)," RFC 2784, March 2000.

[12] J. Mirkovic and E. Kissel, "Comparative evaluation of spoofing defenses," *IEEE Transactions on Dependable and Secure Computing*, pp. 218–232, 2009.

[13] J. Postel, "Internet Protocol," RFC 791, September 1981.

[14] R. Govindan and H. Tangmunarunkit, "Heuristics for Internet map discovery," in *Proceedings of INFOCOM*, vol. 3, 2000, pp. 1371–1380.

[15] F. Gont, "Security assessment of the Internet protocol version 4 (draft-ietf-opsec-ip-security-07.txt)," Internet Draft, April 2011, http://tools.ietf.org/id/draft-ietf-opsec-ip-security-07.txt.

[16] E. Kline and P. Reiher, "Securing data through avoidance routing," in *Proceedings of NSPW*, 2009, pp. 115–124.

[17] C. Villamizar, R. Chandra, and R. Govindan, "BGP Route Flap Damping," RFC 2439, November 1998.

[18] "Akamai," http://www.akamai.com.

[19] K. Park and H. Lee, "On the effectiveness of route-based filtering for distributed DoS attack prevention in power-law Internets," in *Proceedings of SIGCOMM*, vol. 31, no. 4, 2001, pp. 15–26.

[20] Z. Duan, X. Yuan, and J. Chandrashekar, "Constructing inter-domain packet filters to control IP spoofing based on BGP updates," in *Proceedings of INFOCOM*, 2006.

[21] X. Yang, D. Wetherall, and T. Anderson, "A DoS-limiting network architecture," in *Proceedings of SIGCOMM*, 2005, pp. 241–252.

[22] A. Perrig, D. Song, and A. Yaar, "StackPi: New packet marking and filtering mechanisms for DDoS and IP spoofing defense," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 10, pp. 1853–1863, October 2006.

[23] X. Liu, X. Yang, and Y. Lu, "To filter or to authorize: Network-layer DoS defense against multimillion-node botnets," in *Proceedings of SIGCOMM*, 2008, pp. 195–206.

[24] J. Mirkovic, G. Prier, and P. Reiher, "Attacking DDoS at the source," in *Proceedings of ICNP*, 2002, pp. 312–321.

[25] J. Mirkovic, M. Robinson, and P. Reiher, "Alliance formation for DDoS defense," in *Proceedings of NSPW*, 2003, pp. 11–18.

[26] A. Keromytis, V. Misra, and D. Rubenstein, "SOS: Secure overlay services," in *Proceedings of SIGCOMM*, 2002, pp. 61–72.