IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received September 16, 2018, accepted October 14, 2018, date of publication November 9, 2018,
date of current version November 30, 2018.

Digital Object Identifier 10.1109/ACCESS.2018.2877710

Understanding Internet DDoS Mitigation from Academic and Industrial Perspectives

YUAN CAO[1], (Member, IEEE), YUAN GAO[1], (Member, IEEE), RONGJUN TAN[1], QINGBANG HAN[1], AND ZHUOTAO LIU[2]

[1]College of Internet of Things Engineering, Hohai University, Changzhou 213022, China
[2]Electrical and Computer Engineering Department, University of Illinois at Urbana–Champaign, Urbana, IL 61801, USA

Corresponding author: Qingbang Han (20111841@hhu.edu.cn)

This work was supported in part by the National Natural Science Foundation of China under Grant 61601168, Grant 61401146, Grant 11574072, and Grant 11274091, in part by the Key Research Project of Jiangsu under Grant BE2016056, and in part by the Fundamental Research Foundation of Shenzhen under Grant JCYJ20170302151209762.

ABSTRACT Defending against distributed denial of service (DDoS) attacks in the Internet is a fundamental problem. One practical approach to addressing DDoS attacks is to redirect all destination (e.g., via DNS or BGP) to a third-party, DDoS protection-as-a-service provider (e.g., Cloudflare and Akamai), which is well provisioned and equipped with proprietary filtering mechanisms to remove attack traffic before passing the remaining traffic to the destination. Although such an approach is appealing, as it requires no modification to the existing Internet infrastructure and can scale to handle very large attacks, recent industrial interviews with more than 100 interviewees from over 10 industry segments reveal that this approach alone is not sufficient, especially for large organizations (e.g., Web hosting companies and government) that cannot afford to allow third-parity security-service providers to terminate their network connections. Instead, these organizations have to rely on their ISPs to filter attack traffic. In this paper, we discuss the challenges faced by the ISPs in order to disrupt the Internet security-service market and sketch our solutions, powered by smart contracts.

INDEX TERMS DDoS attacks, defense, deployability.

I. INTRODUCTION

The Internet provides an open environment in which any host can communicate with any other host. As a result, security services have traditionally been deployed at each host, rather than inside the network, allowing each host to specify its own security policies, in accordance with the end-to-end principle. Unfortunately, traffic control cannot be accomplished solely by the end host, because modern traffic control algorithms expect the sender and receiver to cooperate to stay within the capacity of each link on the path from the sender to the receiver.

Because the Internet does not enforce any flow control requirements apart from the end hosts, a number of attacks have been developed to overwhelm Internet end systems. The most significant of these attacks is the volumetric Distributed Denial-of-Service (DDoS) attack, representing over 65% of all DDoS attacks. In a volumetric DDoS, many attackers coordinate and send high-rate traffic to a victim, in an attempt to overwhelm the bottleneck links close to the victim. Typical Internet links use RED and drop-tail FIFO queuing disciplines, which provide nearly-equal loss rates to all traffic. Consequently, saturated links impose equal loss rates on attacking and legitimate traffic alike. While legitimate traffic tends to back off to avoid further congestion, attack traffic need not back off, so links saturated by a DDoS attack are effectively closed to legitimate traffic. Recent DDoS attacks include a 620 Gbps attack against Krebs' security blog [1], a 1 Tbps attack against OVH [2], a French ISP, and various attacks powered by the Mirai botnet [3], [4].

Over the past few decades, both industry and academia make a considerable effort to address this problem. Academia have proposed various approaches, ranging from filtering-based approaches [5]–[10], capability-based approaches [11]–[14], overlay-based systems [15]–[17], systems based on future Internet architectures [18]–[20] and other variance [21]–[23]. Meanwhile, many large DDoS-protection-as-a-service providers (e.g., Akamai, Cloudflare), some of which are unicorns, have dominated the market. These providers massively over-provision data centers for peak attack traffic loads and then share this capacity across

2169-3536 © 2018 IEEE. Translations and content mining are permitted for academic research only.
Personal use is also permitted, but republication/redistribution requires IEEE permission.
See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

many customers as needed. When under attack, victims use DNS or BGP to redirect traffic to the provider rather than their own networks. The DDoS protection-as-a-service provider applies a variety of techniques to scrub this traffic, separating malicious from benign, and then re-injects only the benign traffic back into the network to be carried to the victim.

Despite such effort, recent industrial interviews with 100 people from over 10 industry segments that are potential DDoS target indicate that DDoS attacks have not been fully addressed [24]. At the very high level, the reasons are twofold: (i) academic solutions that provably have desirable security properties (e.g., per-sender fairness [13], per-AS fairness [25]) incur significant deployment overhead so that few of them have ever been deployed in the Internet; (ii) current DDoS protection-as-a-service providers relying on empirical filtering rules that are insufficient against sophisticated attacks. Inspired by such industrial feedback, in this paper, we provide systematical analysis about these academic DDoS prevention proposals, focusing on understanding the deployment requirement and challenges of each solution. We also analyze the common industrial DDoS prevention solutions, shedding light on why they work now and why they may fail work in the future. We hope that our analysis can offer useful insights for future innovation in DDoS prevention area.

II. FILTERING-BASED SYSTEM

Filtering-based systems are among the earliest work to address DDoS attacks. In this section, we discuss four representing work in this area, including IP Traceback [5], [6], AITF [7], Pushback [8], [9] and StopIt [10]. At the very high level, these systems aim to defend against DDoS attacks by filtering attack traffic, ideally, at the Autonomous Systems (ASes) close to senders. To this end, these solutions relies on a mechanism to differentiate attack traffic and legitimate traffic.For instance, IP Traceback systems adopt packet marking algorithm to construct the path that carries attack traffic so as to block attack flows. One simple marking algorithm is that each router marks packets with some probability. After receiving enough number of attack packets, with high probability, all routers on the path have marked packets so that the victim can see all possible paths. As a result, the victim can block packets traversing certain paths if it receives too much traffic on these paths (possible attacks). The shortcomings of IP Traceback systems are that on one hand the marking algorithm may converge slowly and on the other hand it faces scaling issues to reconstruct all routes when defending against largescale geographically distributed attacks (which are common in nowadays attacks [3]).

However, the real issue, besides the convergence and scalability concern, prevents such systems from being widely adopted in the Internet is their significant deployment hurdles. As the Internet has involved into a giant distributed system operated by relatively independent ASes, requiring all routers owned by different ASes in the Internet to mark packets on behalf of a DDoS target is impractical because

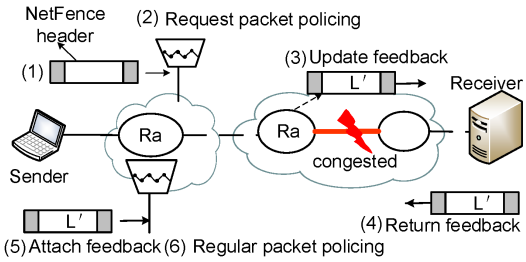

FIGURE 1. The NetFence architecture. NetFence relies on the congestion feedback mechanism to police the traffic for the bottleneck link.

remote ASes that have no business relationship with the DDoS target have little incentive to upgrade their routers.

Different from the IP Traceback system, the AITF system [7] only constructs AS level path for packets and aggregates all traffic traversing the same series of ASes as one flow. Then the victim informs the remote AS close to the attack source to stop certain flow when it suspects attacks. Although arguably AITF has less deployment overhead than the IP Traceback system since only the ingress and egress routers of each AS need to mark packets. However, the DDoS victim still has limited ability to enforce deployment at remote ASes. Further, AITF has high false positive rates to block legitimate traffic since legitimate traffic may be aggregated into the same flow as attack traffic.

In the Pushback system [8], [9], the way to construct the attack path is that congested routers inform their upstream routers to limit the rate of certain type of traffic sending to the victim. Recursively, routers sitting close to the attack source will block the attack traffic. However, the blocked traffic type may include legitimate traffic as well. Further, when upstream routers are located in remote ASes, the Pushback system has the same deployment problem.

StopIt [10] assumes that the victim can detect attack flows. To stop attack, the victim would need to install filters at other remove ASes that are close to the attack source to drop attack traffic. Apparently, StopIt has the same deployment issue as the above filtering-based systems. However, the assumption that the victim can detect attack traffic is actually intriguing since it has the potential of using endhost-based traffic analysis to build up the filters (although StopIt [10] did not explore such a direction).

To sum up, filtering based system often assume a correct way to identify attack traffic and then require remote ASes to install filters to stop attack traffic. However, given the complexity of current Internet ecosystem, requiring deployment from ASes that are unrelated to the victim is unpractical.

III. CAPABILITY-BASED SYSTEMS

Different from filtering-based systems, capability based systems aim to ensure that a sender needs to receive explicit permission before being allowed to contact the destination. When one packet is traversing the network, all routers on its path add its own signatures at the packet header. The series of

signatures are served as a capability which will be returned to the sender by the receiver. The senders have to include capabilities in their future packets so that the routers can verify these capabilities and forward packets accordingly (packets without capabilities are de-prioritized). Capability based systems divide the packets into three categories: privileged packets, request packets and best effort packets. The privileged packets are sent with valid capabilities and the request packets are sent to establish the capabilities. Best effort packets or legacy packets are sent from the legacy hosts who are not capability-aware. Typically, the best-effort traffic cannot be protected.

A. DEPLOYMENT-INTENSIVE APPROACHES

To bootstrap, the capability based systems, such as SIFF [11] and TVA [12], require each source to acquire capabilities over the request channel before further sending packets to the destination. However, attackers could flood the request channel to prevent the capability setup packets from reaching the destination. Such a bootstrap problem is named the Denial of Capability (DoC) attacks [26]. The Portcullis [25] protocol mitigates the DoC attacks by splitting the bandwidth according to the per-computation fairness principle. Specifically, sources have to solve a computation puzzle to send request packets. Therefore, attackers will bound computation resources cannot always overflow the request channel and legitimate users will finally get their bandwidth share at the capability request channel (with high probability).

The second problem is that the colluding senders and receivers located on two sides of a link can grant each other capabilities so that they can still flood the link with privileged packets. NetFence [13] is proposed to achieve per-sender fairness under the colluding scenario (Fig.1). NetFence designs a congestion policing feedbacks mechanism that allows bottleneck router to stamp congestion feedback in packet headers. The access routers who receive these feedbacks will police the traffic for the congested link. Specifically, the access router keeps a rate limiter for each sender that goes through the congested link and updates these rate limiters based on an AIMD (Additive Increase Multiplicative Decrease) algorithm [27]. As proved in [13], each senders rate limit at the bottleneck link will finally converge to fairness.

Although the per-sender fairness result is strong, NetFence relies on unpractical assumption and incur significant upgrade in the Internet core. To begin with, NetFence requires that the source spoofing is impossible. Towards this end, it needs the universal deployment of the Passport [28] protocol. However, considering that another source spoofing elimination protocol Ingress Filter [29] has not been universally deployed after almost three decades [30], it is highly doubtful that Passport which imposes even much larger deployment overhead than Ingress Filter will ever be deployed. Thus, assuming that the Internet-wide source spoofing has been eliminated is unrealistic. Second, each access router needs to share a key with all congested routers, which requires a very complex key management system. Furthermore, all routers need to be able to perform cryptographic process for traversing packets, which requires router upgrade and introduces additional packet processing overhead. Finally, similar to filtering-based systems, the victim has limited ability to enforce deployment when routers are located at remote ASes.

B. DEPLOYMENT-FRIENDLY APPROACHES

These deployment issues have drawn academic attentions. For instance, CRAFT [31] and Mirage [22] are proposed towards easy deployment. In its design, a CRAFT router emulates TCP states for all traversing flows to ensure that no one can get more share than what TCP allows. To maintain the TCP state machine, the CRAFT router only relies on self-created capabilities and does not require any further cooperation from remote routers or ASes. Thus, it requires a much smaller deployment sphere compared with other solutions. However, since TCP have many standards and some traffic (e.g., video flows [32]) may even not use standard TCP protocols, CRAFT is not compatible with real Internet environment and may limit future transport protocol innovation (for instance, Google have deployed a new transport protocol named TCP BBR [33]).

Mirage [22] is a puzzled based defense system designed for securing Web applications. It borrows the idea of frequency hopping in wireless communication to allow the victim server to change their IP addresses. Each time attackers want to send traffic to the victim, they have to solve a computational puzzle to obtain the new address of the victim. For deployment, Mirage relies on the large IPv6 address space to perform effectively. Further, on the client end, it only requires the client to install a Javascript plugin on browsers, which is much more light-weight compared with those solutions that requiring client network stack modification.

Although Mirage makes a great effort towards deployable DDoS solutions, its underlining per-compute fairness is not effective against large scale DDoS attacks. There is a dilemma of setting the difficulty level of puzzles: given simple puzzles, attackers with huge amount of resource can easily flood the victim whereas difficult puzzles may force legitimate clients to use all their CPU cycles to solve puzzles (an extreme example could be Bitcoin mining). To the best of knowledge, no puzzle based systems have ever been deployed in practice for DDoS defense.

C. READILY-DEPLOYABLE SOLUTION

MiddlePolice [14] is the first readily deployable DDoS prevention mechanism. As plotted in Fig.2, MiddlePolice proposes to deploy traffic policing units only on cloud providers that have commercial relationships with the victim. It removes deployment requirement from all other unrelated entities, including remote ASes and clients. Essentially, MiddlePolices deployment model is the same as existing DDoS prevention-as-a-service vendors, which has been proved to be successful.

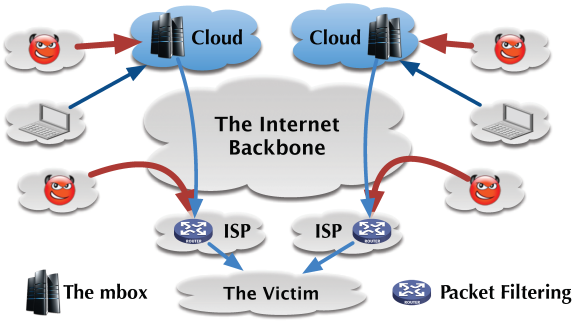

FIGURE 2. MiddlePolice requires deployment only from the cloud to be readily deployable in the Internet.

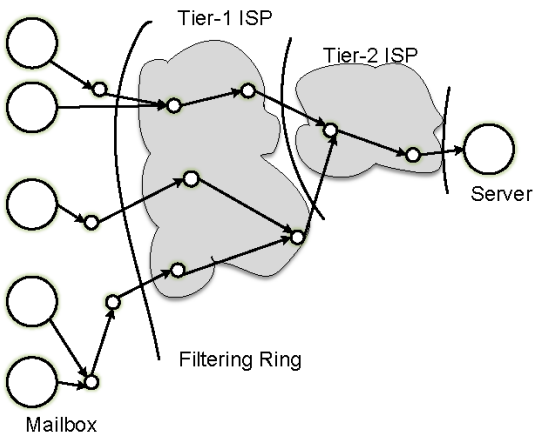

FIGURE 3. The Phalanx architecture. Phalanx allows sender to construct random paths via mailboxes to reach the victim server. Attackers that do not know the path can only disrupt a small fraction of legitimate traffic.

Besides the deployment, MiddlePolice is also the first work advocating enforcing victim-defined traffic control during DDoS mitigation. In prior solutions, it the network (e.g., access routers in NetFence [13]) that determines the traffic policing algorithm. As a result, if a single in-network prevention mechanism is ever deployed, a single fairness regime will be enforced, forcing the victims to accept the choice made within the network. However, there is no single fairness metric that is effective in all situations. For instance, per-computation fairness is unfair to mobile users and per-sender fairness is subject to source spoofing. Further, the victim may want more advanced bandwidth sharing policies besides fair share. For instance, the victim may want to always allow premium users to get services even in the case of DDoS attacks. Ideally, DDoS prevention approaches should allow the victim server to enforce self-preferred traffic control polices that are consistent with their business and missions. In fact, this aligns with the Internet end-to-end design principle that pushes functionality to the network edge. As such, MiddlePolice proposes to allow the victim to define self-preferred traffic control algorithms, which not only ensures that the victim can stop attack traffic, but also enabling the victim to wisely allocate the remaining network bandwidth among legitimate clients.

MiddlePolice is not without challenge. One concern is that due to the different technical and sophistication level, some victims may have not clear picture about their network traffic so that they cannot propose meaningful traffic control policies. Another concern is the privacy: it may be privacy-invasive for some large victims (e.g., government) to redirect all its traffic to the cloud if they cannot fully trust the cloud. Umbrella is readily deployable and privacy-preserving, yet designed primarily for protecting the immediate links between a victim and its ISPs [47].

IV. OVERLAY-BASED SYSTEMS

Overlay-based systems propose to address DDoS attacks by building overlay networks atop the Internet. For instance, Phalanx [15] leverages the power of swarm to bat the bots (Fig.3). Specifically, it builds an overlay network consisting of nodes called mailboxes. Each sender randomly chooses a sequence of mailboxes to construct its path to the victim server. Since attackers do not know the exact path taken by each sender, they can disrupt at most only a fraction of the legitimate traffic. Different from Phalanx that aims to increase the path diversity using the overlay network, SOS [16] proposes to use the overlay to allow the victim server to differ the authorized and unauthorized packets. In particular, the SOS architecture enables a victim server to secretly pick several nodes in the overlay network as forwarding proxies. These proxies are typically only disclosed to legitimate users. As a result, the victim can deploy filters to drop any traffic that does not traverse these proxies.

The strongest metric of overlay-based systems is that they can, theoretically, scale to defend against DDoS attacks with any size by replicating more nodes in the overlay network. However, these solutions fail to propose valid incentives for running as an overlay. To give some sense about how hard it is to build a large scale overlay network, the Tor network, the most popular anonymity network, has recruited only about 7000 overlays even though Tor offers legitimate incentives for people to run overlays (i.e., protecting online privacy). Additionally, these proposals fail to give convincible evidence that redirecting traffic through the overlay would not impact networking performance.

V. FURTHER INTERNET ARCHITECTURE

Since the design of original Internet architecture does not place enough emphasis on security, the research community has proposed some clean-slate new Internet architecture that can solve various security problems, including DDoS attacks. One of the pioneering proposal is the Accountable Internet Protocol (AIP) [19]. AIP tries to associate each entity a selfcertifying ID so that any action in the Internet can be traced back to its executors. Arguably, this makes it easier for the victim to enforcement punishment in the case of misbehavior. In fact, AIP shares the similar goal as these source spoofing elimination protocols such as the aforementioned Ingress Filter [29] and Passport [28]. However, AIP

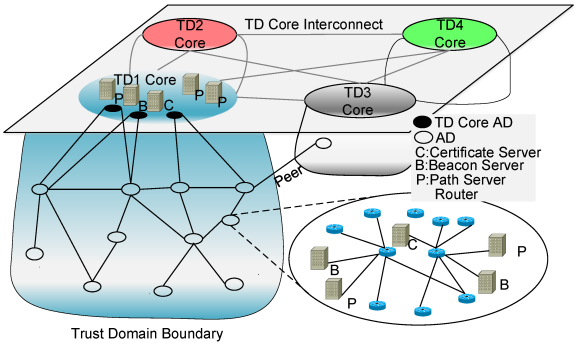

FIGURE 4. The SCION architecture. SCION divides the Internet into Trust Domain to achieve reliable and controllerable end-to-end path construction.

requires fundamental changes of the Internet protocol and architecture.

SCION [18] (and many other works built upon SCION like SIBRA [34]) is one of most well known future Internet architecture proposals. SCION, illustrated in Fig.4, separates Autonomous Domains (ADs) into groups of independent routing sub-planes, called trust domains (TDs), which then interconnect for global connectivity. One may consider one TD as a top tier ISP. Trust domains provide isolation for routing failures and misconfiguration happened in different ADs. To obtain network routing information, TD cores periodically propagate a Path Construction Beacon (PCB). Upon receiving PCB, an AD will add its own public key into PCB. Therefore, PCB records all possible paths that an AD can use to reach the TD cores. Each AD needs to pick k paths and registers them with the TD cores. To construct end-to-end paths, the source AD composes its path to the TD Core (up-path) with the path from the TD Core to the destination AD (down-path). Therefore, SCION allows explicit path selection by source, which gives end-hosts strong control for both inbound and outbound traffic.

As a clean-slate architecture, SCION has many desirable security properties that the current Internet architecture does not have. One great metric of SCION is that each domain can evolve independently, and we have heard news that some large ISPs in Switzerland have tried SCION in their controlled network. Unfortunately, we have not experienced much globalscale deployment yet, and we do not anticipate that a complete transition is near the future.

Software Defined Networks (SDN) and Network Functions Virtualization (NFV) are new technologies for higher performance networks with increased band-width and scalability. SDN is a new network architecture that can improve cloud manageability, scalability, controllability and dynamism. In SDN, the network control plane is decoupled from the data plane. The network routing rules and the resources are centrally programmed and controlled [35]. In parallel, NFV was proposed to decouple the network software functions and routing rules from the underlying infrastructures and supports low cost, high efficient and elastic network function setup and capability increment [36]. In NFV environmental, SDN is capable of implementing the virtualized network functions through a cost-effective and high-speed way that the computing resources are elastically allocated to accomplish the undergoing network functions. The Combination of these two promising paradigms (SDN and NFV) helps flexibly manage and deploy of the virtual resource allocations to respond to dynamic conditions and needs.

However, since all the services and policies in the SDN are distributed and managed from a single centralized controller, the SDN controller can become the vulnerable spot for the adversaries. If the controller is attacked, the control of the whole network is lost. The huge data plane could be taken offline if the control plane is overloaded by the junk even and huge session density. Recently, a few works have been proposed to exploit the SDN and NFV advantages such as dynamic packet forwarding rule configuration, global traffic analysis traffic monitoring and so on to detect and prevent the DDoS attacks [37], [38]. In [39], an elastic DDoS defense system Bohatei has been designed and implemented. It leverages the NFV capabilities to handle 5000 Gbps attack and the flexibility of the SDN to be resilient to the dynamic adversaries with minimum latency. Reference [40] also proposed and developed a solution that integrates NFV and SDN seamlessly to mitigate the DDoS attacks in a decentralized and coordinated way.

Information Centric Networking (ICN) is also future Internet technology that focuses on contents rather than infrastructures. In contrast to the traditional network, it has to guarantee the security of the contents instead of the communication channels. The ICN architecture is vulnerable to various kinds of routing and caching attacks. As any user in ICN can publish contents, the adversaries can publish available or unavailable contents online and request the malicious routes, which are forwarded to their neighbours so that the ICN architecture can be overwhelmed. Besides, any user can cache contents and send them to subscribers from the nearest copy. The adversaries can deliver malicious requests to force ICN caches to store unpopular contents and evict the popular ones and hence an attacker pollutes ICN caching system. DDoS attacks may occur if the above attacks are performed in a large scale distributed manner [41].

In order to close up the security hole of ICN, a few works have been proposed. For routing attacks, the satisfaction-based pushback strategy is proposed to resist the routing DDoS attack in [42]. Compagno *et al.* [43] defined a threshold to limit request satisfaction ratio (RSR) and Pending Interest Table (PIT) space in Named Data Networking (NDN) architecture. For caching attacks, a low-cost solution is proposed to detect the cache pollution and malicious content in [44]. The detection threshold is caculated by the number of references to each content and its changes. The Cacheshield was presented in [45], which employs a shield function to discard the malicious contest in the routers. A raking algorithm for the content publishers and subscribers was also

TABLE 1. Property summary of some representative solutions from each category. "*O*(*N*) states" means that the number of states maintained by a router increases with the number of attackers. "Cryptography" means that a router needs to support cryptography operation, e.g., MAC computation. "Puzzle" means that the mechanism requires computational puzzle distribution.

	Pushback[8]	SIFF[11], TVA[12]	Netfence[13]	Phalanx[15]	Mirage[22]	SIBRA[34]	MiddlePolice[14]
Source upgrades	**No**	Yes	Yes	Yes	Yes	Yes	**No**
Dest. upgrades	**No**	Yes	Yes	Yes	Yes	Yes	Yes
AS Deployment	Remote and unrelated	Remote and unrelated	Remote and unrelated	Remote and unrelated	**Commercially related**	Remote and unrelated	**Commercially related**
Router support	$O(N)$ states	Cryptography; $O(N)$ states for [12]	$O(N)$ states; Cryptography	$O(N)$ states	Larger memory	**None**	**None**
Traffic control policies	None	None	Per-sender fairness	None	Per-compute fairness	Per-AS fairness	**Victim-selectable policies**
Other requirements	**None**	New header	New header; Passport[28]	New header	Puzzle; IPv6 upgrade	Redesign the Internet	Traffic understanding; Cloud Trust

proposed in [46] to forbid the publication spam. ICN will be incrementally deployed as a promising future network architecture. However, the security should be seriously taken care of in ICN. The solutions are expected to achieve lightweight, privacy, low latency and plausible deniability for users.

In Table 1, we summarize property of some representative solutions from each category.

VI. COMMON INDUSTRIAL PRACTICE

Since few academic proposals have ever been deployed, industrial approaches to DDoS mitigation become significantly important to keep business online. These DDoS prevention service providers ask their customers to redirect all customer traffic to their well-provisioned data centers, use proprietary algorithms to scrub likely-malicious traffic, and then re-inject the scrubbed traffic back towards the original destination. Since large service providers typically have geographically distributed data centers, they are particularly good at filtering the large-but-obvious-to-catch DDoS attacks; that is, although the volume of attack traffic is very large, it is fairly easy to identify the attack traffic. Fortunately, most of current volumetric DDoS attacks (such as SYN flooding and NTP reflection) belong to such category.

However, as increasing number of Internet-of-Things (IoT) devices are connected to the Internet, we would expect to see larger scale and more sophisticated DDoS attacks in the near future. Although service providers could potentially scale their cloud to handle even larger DDoS attacks, they will face at least three open challenges when defending against more sophisticated DDoS attacks. We start with a simple strategic attack can completely bypass the filtering of DDoS prevention providers deployed on the cloud. Current vendors have the victim use DNS to redirect custom traffic to service providers cloud. However, more advanced attackers can first try to figure out the real IP address of the victim, and then directly send attack traffic to the victim without relying on DNS. As a result, attack traffic will not go through service providers' cloud, completely bypassing their defense.

Although it is possible for current DDoS prevention providers to fix the first problem, it will become increasingly difficult for them to filter sophisticated DDoS attacks simply based on their proprietary filtering algorithms. Large DDoS prevention vendors typically have years of operation experience and therefore develop a set of empirical filtering rules to scrub obvious DDoS attack traffic. Thus, they often advocate that their algorithms can catch over 99% DDoS attack traffic. This number may be true since these large-but-obvious-to-catch DDoS attacks contribute huge amount of traffic volume. However, since vendors only have a limited set of empirical rules and each rule is stateless (otherwise the vendors themselves are vulnerable to state exhaustion attacks), any attack vector that has not been experienced and learned by the vendors will pass such filtering. Thus, when DDoS attacks enter the new era of being cleverer and more sophisticated, relying on static filtering rules will be insufficient.

The third challenge is the direct cause of the second problem: since existing DDoS prevention vendors only focus on and are particularly good at scrubbing large-but-obvious-to-catch attacks, in case of sophisticated DDoS attacks, how customers should handle the attack traffic that is not scrubbed by these vendors. We believe the third challenge is the ultimate question that the next-generation DDoS prevention technique should address.

VII. CONCLUSION

In this paper, we systematically analyze representative academic proposals and common industrial practice for DDoS prevention. We focus on understanding the deployment requirement and challenges of each academic solution. In particular, we categorize these academic proposals into filtering based approaches, capability-based approaches, overlay-based systems and systems based on future Internet architectures. In general, although these proposals provide provably secure properties, such as fairness, they often require extensive and unenforceable deployment in the Internet, which is the major reason why we witness little progress on real-world deployment of these systems.

Therefore, industrial DDoS prevention solutions (often referred to as Cloud or CDNs) are critical for keeping business online. However, these solutions typically suffer from another challenge: although they are good at defending against extremely large volumetric DDoS attacks, they are not effective against sophisticated DDoS attacks that are backed up by numerous unsecured IoT devices. We hope that this concise survey paper provides a clear picture of the landscape of DDoS prevention, inspiring the community to bridge the academic research with industrial practice.

REFERENCES

[1] B. Krebs. *KrebsOnSecurity Hit With Record DDoS*. Accessed: 2017. [Online]. Available: https://krebsonsecurity.com/2016/09/krebsonsecurity-hit-with-record-ddos/

[2] O. Klaba. *OVH Hit With 1 Tbps DDoS*. Accessed: 2017. [Online]. Available: https://twitter.com/olesovhcom/status/778830571677978624

[3] M. Antonakakis *et al.*, "Understanding the Mirai botnet," in *Proc. USENIX Secur. Symp.*, 2017, pp. 1092–1110.

[4] M. Antonakakis, "Understanding the Mirai botnet," in *Proc. 26th USENIX Conf. Secur. Symp.*, Berkeley, CA, USA: USENIX Association, 2017, pp. 1093–1110. [Online]. Available: http://dl.acm.org/citation.cfm?id=3241189.3241275

[5] S. Savage, D. Wetherall, A. Karlin, and T. Anderson, "Practical network support for IP traceback," in *Proc. ACM SIGCOMM*, 2000, pp. 295–306.

[6] D. X. Song and A. Perrig, "Advanced and authenticated marking schemes for IP traceback," in *Proc. IEEE INFOCOM*, vol. 2, Apr. 2001, pp. 878–886.

[7] K. J. Argyraki and D. R. Cheriton, "Active Internet traffic filtering: Real-time response to denial-of-service attacks," in *Proc. USENIX ATC*, 2005, pp. 135–148.

[8] R. Mahajan, S. M. Bellovin, S. Floyd, J. Ioannidis, V. Paxson, and S. Shenker, "Controlling high bandwidth aggregates in the network," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 32 no. 3, pp. 62–73, 2002.

[9] J. Ioannidis and S. M. Bellovin, "Implementing pushback: Router-based defense against DDoS attacks," in *Proc. USENIX NSDI*, 2002, pp. 1–8.

[10] X. Liu, X. Yang, and Y. Lu, "To filter or to authorize: network-layer DoS defense against multimillion-node botnets," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 4, pp. 195–206, 2008.

[11] A. Yaar, A. Perrig, and D. Song, "SIFF: A stateless Internet flow filter to mitigate DDoS flooding attacks," in *Proc. IEEE Symp. Secur. Privacy*, May 2004, pp. 130–143.

[12] X. Yang, D. Wetherall, and T. Anderson, "TVA: A DOS-limiting network architecture," *IEEE/ACM Trans. Netw.*, vol. 16, no. 6, pp. 1267–1280, Dec. 2008.

[13] X. Liu, X. Yang, and Y. Xia, "NetFence: Preventing Internet denial of service from inside out," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 4, pp. 255–266, 2011.

[14] Z. Liu, J. Hao, Y.-C. Hu, and M. Bailey, "MiddlePolice: Toward enforcing destination-defined policies in the middle of the Internet," in *Proc. ACM-CCS*, 2016, pp. 1268–1279.

[15] C. Dixon, T. E. Anderson, and A. Krishnamurthy, "Phalanx: withstanding multimillion-node botnets," in *Proc. USENIX NSDI*, 2008, pp. 45–58.

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

[17] D. G. Andersen, "Mayday: distributed filtering for Internet services," in *Proc. USNIX USITS*, 2003, pp. 31–42.

[18] X. Zhang, H.-C. Hsiao, G. Hasker, H. Chan, A. Perrig, and D. G. Andersen, "SCION: Scalability, control, and isolation on next-generation networks," in *Proc. IEEE Symp. Secur. Privacy*, May 2011, pp. 212–227.

[19] D. G. Andersen, H. Balakrishnan, N. Feamster, T. Koponen, D. Moon, and S. Shenker, "Accountable Internet protocol (AIP)," in *Proc. ACM SIGCOMM*, 2008, pp. 339–350.

[20] D. Naylor *et al.*, "XIA: Architecting a more trustworthy and evolvable Internet," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 3, pp. 50–57, 2014.

[21] M. Walfish, M. Vutukuru, H. Balakrishnan, D. Karger, and S. Shenker, "DDoS defense by offense," in *Proc. ACM SIGCOMM*, 2006, pp. 1–12.

[22] P. Mittal, D. Kim, Y.-C. Hu, and M. Caesar. (2011). "Mirage: towards deployable DDoS defense for Web applications." [Online]. Available: https://arxiv.org/abs/1110.1060

[23] Y. Gilad, A. Herzberg, M. Sudkovitch, and M. Goberman, "CDN-on-demand: An affordable DDoS defense via untrusted clouds," in *Proc. NDSS*, 2016, pp. 1–15.

[24] Z. Liu, H. Jin, Y. Hu, and M. Bailey, "Practical proactive DDoS-attack mitigation via endpoint-driven in-network traffic control," *IEEE/ACM Trans. Netw.*, vol. 26, no. 4, pp. 1948–1961, Aug. 2018, doi: 10.1109/TNET.2018.2854795.

[25] B. Parno, D. Wendlandt, E. Shi, A. Perrig, B. Maggs, and Y.-C. Hu, "Portcullis: Protecting connection setup from denial-of-capability attacks," in *Proc. ACM SIGCOMM*, 2007, pp. 289–300.

[26] K. Argyraki and D. Cheriton, "Network capabilities: The good, the bad and the ugly," in *Proc. ACM HotNets-IV*, 2005, p. 140.

[27] D.-M. Chiu and R. Jain, "Analysis of the increase and decrease algorithms for congestion avoidance in computer networks," *Comput. Netw. ISDN Syst.*, vol. 17, no. 1, pp. 1–14, Jun. 1989.

[28] X. Liu, A. Li, X. Yang, and D. Wetherall, "Passport: Secure and adoptable source authentication," in *Proc. USENIX NSDI*, 2008, pp. 365–378.

[29] D. Senie and P. Ferguson. (2000). *Network Ingress Filtering: Defeating Denial of Service Attacks Which Employ IP Source Address Spoofing*. [Online]. Available: https://tools.ietf.org/html/bcp38

[30] (2011). *Everyone Should be Deploying BCP 38! Wait, They Are*. [Online]. Available: http://www.senki.org/everyone-should-be-deploying-bcp-38-wait-they-are

[31] D. Kim, J. T. Chiang, Y. C. Hu, A. Perrig, and P. Kumar, "CRAFT: A new secure congestion control architecture," in *Proc. ACM CCS*, 2010, pp. 705–707.

[32] P. Ameigeiras, J. Ramos-Munoz, J. Navarro-Ortiz, and J. M. Lopez-Soler, "Analysis and modelling of YouTube traffic," *Trans. Emerg. Telecommun. Technol.*, vol. 23, no. 4, pp. 360–377, 2012.

[33] N. Cardwell, Y. Cheng, C. S. Gunn, S. H. Yeganeh, and V. Jacobson, "BBR: Congestion-based congestion control," *ACM Mag. Queue*, vol. 14, no. 5, p. 50, 2016.

[34] C. Basescu *et al.*, "SIBRA: Scalable Internet bandwidth reservation architecture," in *Proc. NDSS*, 2016, pp. 1–16.

[35] Y. Xu and Y. Liu, "Ddos attack detection under SDN context," in *Proc. IEEE INFOCOM*, Apr. 2016, pp. 1–9.

[36] *Network Functions Virtualisation Introductory White Paper*. Accessed: 2012. [Online]. Available: http://portal.etsi.org/NFV/NFV_White Paper.pdf

[37] Q. Yan, F. R. Yu, Q. Gong, and J. Li, "Software-defined networking (SDN) and distributed denial of service (DDoS) attacks in cloud computing environments: A survey, some research issues, and challenges," *IEEE Commun. Surveys Tuts.*, vol. 18, no. 1, pp. 602–622, 1st Quart., 2016.

[38] K. Futamura *et al.*, "vDNS closed-loop control: A framework for an elastic control plane service," in *Proc. IEEE IEEE Conf. Netw. Function Virtualization Softw. Defined Netw. NFV-SDN*, Nov. 2015, pp. 170–176.

[39] S. K. Fayaz, Y. Tobioka, V. Sekar, and M. Bailey, "Bohatei: Flexible and elastic DDoS defense," in *Porc. 24th USENIX Secur. Symp.*, Washington, DC, USA: USENIX Association, 2015, pp. 817–832.

[40] L. Zhou and H. Guo, "Applying NFV/SDN in mitigating DDoS attacks," in *Proc. IEEE Region 10 Conf. TENCON*, Nov. 2017, pp. 2061–2066.

[41] M. Wählisch, T. C. Schmidt, and M. Vahlenkamp, "Backscatter from the data plane—Threats to stability and security in information-centric network infrastructure," *Comput. Netw.*, vol. 57, no. 16, pp. 3192–3206, 2013.

[42] A. Afanasyev, P. Mahadevan, I. Moiseenko, E. Uzun, and L. Zhang, "Interest flooding attack and countermeasures in named data networking," in *Proc. IFIP Netw. Conf.*, May 2013, pp. 1–9.

[43] A. Compagno, M. Conti, P. Gasti, and G. Tsudik, "Poseidon: Mitigating interest flooding DDoS attacks in named data networking," in *Proc. IEEE 38th Conf. Local Comput. Netw. (LCN)*, Oct. 2013, pp. 630–638.

[44] M. Conti, P. Gasti, and M. Teoli, "A lightweight mechanism for detection of cache pollution attacks in named data networking," *Comput. Netw.*, vol. 57, no. 16, pp. 3178–3191, 2013.

[45] M. Xie, I. Widjaja, and H. Wang, "Enhancing cache robustness for content-centric networking," in *Proc. IEEE INFOCOM*, Mar. 2012, pp. 2426–2434.

[46] N. Fotiou, G. F. Marias, and G. C. Polyzos, "Fighting spam in publish/subscribe networks using information ranking," in *Proc. 6th EURO-NGI Conf. Next Gener. Internet*, Jun. 2010, pp. 1–6.
[47] Z. Liu, Y. Cao, M. Zhu, and W. Ge, "Umbrella: Enabling ISPs to offer readily deployable and privacy-preserving DDoS prevention services," *IEEE Trans. Inf. Forensics Security*, to be published.

YUAN CAO (S'09–M'14) received the B.S. degree from Nanjing University in 2008, the M.E. degree from The Hong Kong University of Science and Technology in 2010, and the Ph.D. degree from Nanyang Technological University in 2015. He is currently an Assistant Professor with the College of Internet of Things Engineering, Hohai University. His research interests include hardware security, ASIC physical unclonable function, and analog/mixed-signal VLSI circuits and systems.

YUAN GAO received the B.E. degree in communication engineering and the M.E. and Ph.D. degrees in communication and information system from Hohai University, China, in 1998, 2005, and 2016, respectively. She is currently an Associate Professor with the College of Internet of Things Engineering, Hohai University. Her research is in the areas of wireless communication and networks with an emphasis on cooperative communication, signal processing, and security in cognitive radio networks.

RONGJUN TAN received the B.E. degree in communication engineering form Hohai University, China, in 2018, where she is currently pursuing the master's degree with the College of Internet of Things Engineering. Her research is in the areas of wireless communication and networks with an emphasis on physical layer security.

QINGBANG HAN received the Ph.D. degree from Tongji University. He is currently a Full Professor with the College of Internet of Things Engineering, Hohai University, and the Director of the Underwater Information Perception Center. His research interests include information acquisition and information processing. He is also a regular reviewer of many international journals. He is a member of the China Institute of Acoustics Society of Detection, the Acoustics Association of Jiangsu Province, and the China Geophysical Society Engineering Survey Professional Committee.

ZHUOTAO LIU received the B.S. degree from Shanghai Jiaotong University in 2012 and the Ph.D. degree from the University of Illinois at Urbana–Champaign in 2017. He is currently with the NetInfra Team, Google Research. His primary research interests include system security and privacy, data center networking, and smart-contract-powered network security services.