Mitigate DDoS Attacks in NDN by Interest Traceback

Huichen Dai, Yi Wang, Jindou Fan, Bin Liu
Tsinghua National Laboratory for Information Science and Technology
Dept. of Computer Science and Technology, Tsinghua University
{dhc10, yiwang09, fjd07}@mails.tsinghua.edu.cn, liub@tsinghua.edu.cn

***Abstract*—Current Internet is reaching the limits of its capabilities due to its function transition from host-to-host communication to content dissemination. Named Data Networking (NDN) – an instantiation of Content-Centric Networking approach, embraces this shift by stressing the content itself, rather than where it locates.**

NDN tries to provide better security and privacy than current Internet does, and resilience to Distributed Denial of Service (DDoS) is a significant issue. In this paper, we present a specific and concrete scenario of DDoS attack in NDN, where perpetrators make use of NDN's packet forwarding rules to send out Interest packets with spoofed names as attacking packets. Afterwards, we identify the victims of NDN DDoS attacks include both the hosts and routers. But the largest victim is not the hosts, but the routers, more specifically, the Pending Interest Table (PIT) within the router. PIT brings NDN many elegant features, but it suffers from vulnerability. We propose *Interest traceback* as a counter measure against the studied NDN DDoS attacks, which traces back to the originator of the attacking Interest packets. At last, we assess the harmful consequences brought by these NDN DDoS attacks and evaluate the Interest traceback counter measure. Evaluation results reveal that the Interest traceback method effectively mitigates the NDN DDoS attacks studied in this paper.

***Index Terms*—NDN, DDoS, PIT, Traceback**

I. Introduction

Internet was originally designed for host-to-host communication, but it is now overwhelmingly used for content distribution, which has gradually shown a poor match. To address this function transition, Named Data Networking [1] (NDN), an instantiation of the Content-Centric Networking [2] (CCN) approach, was recently proposed and widely regarded as a promising architecture for future networks. Quite different from the current IP-based network, this new paradigm is characterized by name-based routing and forwarding.

Security and privacy are among the fundamental requirements for NDN. In current Internet, Distributed Denial of Service (DDoS) attacks consume the resources of a remote host or network, thereby denying or degrading service to legitimate users. Such attacks belong to the most difficult security problems, since they are easy to launch, hard to prevent, and difficult to trace back. Therefore, NDN's resilience to DDoS attacks deserves our full attention. Though NDN's architecture provides better security and privacy support than current Internet does, DDoS attacks can also be launched in NDN.

This paper studies a specific and concrete scenario of NDN attacks, where perpetrators fill Interest packets with spoofed names to solicit *inexistent* content. Understanding these NDN DDoS attacks require some background knowledge of NDN. There are two kinds of packets in NDN – the *Interest* packet and the *Data* packet, and the Interest packet is the request and the Data packet is the response. NDN adopts a requester-driven communication model: a requester sends out an Interest packet to solicit the content it wants, which is specified by the name in the Interest packet, the content provider returns the content within a Data packet. While forwarding each Interest packet, a router will keep track of received but un-responded Interest packets, i.e., unsatisfied Interests, in the Pending Interest Table (PIT).

PIT brings NDN many significant features [3], however, it is also vulnerable and can be the target of NDN DDoS attacks. A DDoS attack in NDN can be launched in this way: an attacker may make use of the PIT to exhaust a router's memory and computing resources by sending out a huge number of Interest packets with spoofed names that solicit inexistent content. In this way, these Interest packets become attacking packets, they will never be satisfied and will stay in the PIT forever or until the PIT entries expire. If multiple distributed hosts (zombies or botnets) collaborate to attack the network by sending out numerous unsatisfiable Interest packets, the PIT size will grow rapidly and use up the memory and computing resource in routers. Furthermore, if distributed compromised systems send out Interest packets with spoofed names that share a common prefix, these Interest packets will be all forwarded to the same content provider corresponding to the name prefix, which is almost the identical replay of DDoS attacks in IP. The content provider is surely a victim of the attack, but it can deny the illegimate Interest packets by only using a filter, say Bloom filter, to distinguish Interest names that request inexistent content on the server. However, the PIT in the edge router directly connecting the content provider will be seriously suffering, because all the attacking Interest packets will traverse it to exhaust its resources. And because PIT is maintained by routers, a very obvious difference between DDoS in IP and in NDN is that, the victim of DDoS in IP is a specific host or server at the edge of network, while the largest victims of DDoS in NDN are routers inside network.

To mitigate the studied NDN DDoS attacks, this paper proposes a counter measures named *Interest traceback*, which resembles IP traceback very much. IP traceback [4]–[7] aims to find out the real originator of an IP packet (the attacking traffic may be filled with spoofed source addresses), it is meaningful in dealing with DDoS attacks. But IP traceback is non-trivial, it is difficult and expensive to implement. However, with the help of PIT, NDN supports Interest traceback at ease. We can readily trace back to the real originator of an Interest packet that requests the inexistent content by generating a spoofed Data packet to *satisfy* it, and this originator is a potential attacker. After the attacker is identified, we limit the rate of the interface connecting this attacker at its access router, in order to reduce the attacking Interest packets entering the network.

Especially, in this paper, we make the following contributions:

1) We study a specific and concrete scenario of DDoS attack in NDN, rather than stay at concept level. Detailed analysis

This work is supported by NSFC (61073171), Tsinghua University Initiative Scientific Research Program(20121080068), the Specialized Research Fund for the Doctoral Program of Higher Education of China(20100002110051).

on the NDN DDoS attack is provided as well, including identifying the largest victims, comparing with IP DDoS attacks, etc;
2) *Interest traceback* is proposed as a counter measure to deal with NDN DDoS attacks by tracing back to the originators of Interest packets;
3) We assess the harmful consequences of NDN DDoS attacks by experiments, and evaluation results show that our proposed counter measure effectively alleviate the attacks.

The remainder of our paper is organized as follows: Section II introduces the background of NDN. In Section III, we describe a scenario of NDN DDoS attacks in detail and propose a counter measure against it. Section IV presents the evaluation results, Section V surveys the related work and Section VI concludes our paper.

II. NDN Background

In order to better understand the Interest Flooding Attack, we need to first have a knowledge of the NDN background. However, NDN has a vast background that we cannot fully cover, here we will only introduce the aspects related to this paper.

NDN, a specific instantiation of the CCN paradigm, is a novel next-generation network architecture proposed by [1] recently. Different from current Internet practice, it makes content ("what") as its central role, rather than "where" the content is located. A critical distinction from IP is that, every piece of content in NDN network has an assigned name, and packets are routed and forwarded by names, rather than IP addresses.

A. NDN Names

NDN names are application-dependent and opaque to the network, but they all share the common characteristics – hierarchically structured and composed of explicitly delimited *components*. A typical example of an NDN name is the reversed domain names followed by a directory-style path, e.g., *org/ieee-infocom/2013/cfp.html*, where *org/ieee-infocom/* is the reversed domain name of *ieee-infocom.org*, and *2013/cfp.html/* is the content's directory path on the website server. '/' is the component boundary delimiter and not a part of the name; *org, ieee-infocom, 2013* and *cfp.html* are the 4 components of the name.

B. Requester-driven Communication Model

There are two kinds of packets in NDN – Interest packet and Data packet. In essence, Interest packet is the request and Data packet is the response. An Interest packet carries a name, or an identifier, that specifies the desired data; the name is also encapsulated in the Data packet, indicating the content of itself.

NDN adopts requester-driven communication model, in which a requester sends out an Interest packet for desired content, and a server returns the content within a Data packet.

C. Packet Forwarding

An NDN router has different forwarding processes for Interest packet and Data packet, which are illustrated by Fig. 1(a) and Fig. 1(b), respectively. From the figures we can see that an NDN router maintains three tables: Forwarding Information Base (FIB), Pending Interest Table (PIT) and Content Store (CS). FIB is the NDN routing table; PIT, as the name suggests, keeps track of un-satisfied Interest packets, as well as their incoming interfaces; and CS strategically caches Data packets to serve subsequent Interest packets requesting the same content.

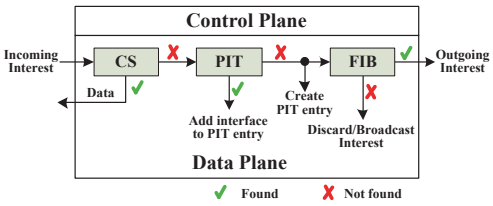

(a) Interest lookup and forwarding process.

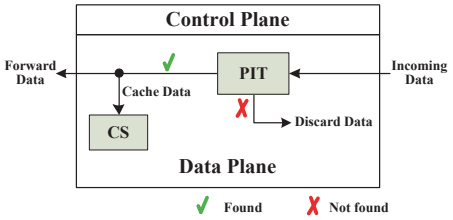

(b) Data lookup and forwarding process.

Fig. 1. Packet lookup and forwarding process.

Fig. 1(a) shows that once an Interest packet I arrives at interface i of an NDN router R, R:

1) consults CS if the desired content is present, if so, returns a copy in Data packet via i,
2) if not, looks up PIT to see if PIT has an entry for I. If so, adds i to that entry, and discards I,
3) otherwise, creates a PIT entry for I and add i to this entry, and
4) forwards I to the next-hop interface by looking up FIB.

When Data packet D returns, Fig. 1(b) shows that R:

1) forwards D over all the requesting interfaces in the corresponding PIT entry, and then deletes this entry,
2) caches D in the CS based on policies.

Based on the forwarding rules above, a Data packet travels back to the requester by taking the same path of the Interest packet, but in the reverse direction, which means symmetric routing. This property is very critical to our proposed counter measure against the NDN DDoS attacks.

III. Distributed Denial of Service in NDN

In this section, we introduce a specific and concrete scenario of NDN DDoS attacks in detail, and a counter measure called Interest traceback is subsequently proposed to fight against NDN DDoS attacks.

A. DDoS Attacks in NDN

NDN generalizes the Internet architecture by replacing the focus on where – addresses of hosts – with what – names of the content that users and applications care about. The packets in NDN are routed and forwarded by these names, i.e., name-based routing. Therefore, end hosts in a connection has no idea about the identity and location of each other, and the packets transmitted in the network only contain content names, rather than any addresses that identify the locations of end hosts. This anonymous communication property endows NDN inherent support of security and privacy. However, NDN does not eliminates DDoS attack, which resembles IP DDoS in the aspect that distributed attackers send out numerous packets to exhaust the victim's resources, thereby denying or degrading service to legitimate users.

Remember that in Section II, we introduced PIT, which keeps track of the unsatisfied Interest packets traversing a router, as well as their arrival interfaces. Unlike the stateless routing in IP,

PIT transfers the stateless routing to stateful routing in NDN, and this transferring brings NDN many significant features [3].

Nevertheless, an NDN router has to insert almost every incoming Interest packet into PIT, and remove each Data packet from PIT, resulting in a large-sized PIT with extremely high access frequency. Therefore, PIT consumes a large amount of memory and CPU time, which makes PIT very vulnerable since attackers can make use of this to exhaust memory and computing resources on a router. The attacking scenario is to send out numerous Interest packets with *spoofed* names by multiple compromised systems to solicit inexistent content – DDoS attacks. Since the names are spoofed, there is no content that satisfies these Interest packets will be returned. Thereby these Interest packets will stay in the PIT for as long time as possible, aiming to to exhaust the router's resources!

Generally, under this scenario, the NDN DDoS attack can be grouped into two categories, which closely couple with the forwarding rule of Interest packets. Both the two categories send out numerous Interest packets with *spoofed* content names, but they are differentiated by the ways how the spoofed names are composed. These Interest packets are forwarded within a domain.

1) First Category – Single-target DDoS Attacks: The first category of NDN DDoS resembles IP DDoS so much that it can be viewed as the equivalent of IP DDoS in NDN. It makes use of the Longest Prefix Match rule while looking up Interest names in the FIB. This category concatenates an existing prefix in the FIB and a forged suffix to compose a name, and fills the name in the Interest packet, as shown in Fig. 2. This Interest packet is then sent out, since its name has a matched prefix in FIB, it can be found in FIB. Therefore, this Interest packet will be at last forwarded to the destination content provider corresponding to the name prefix. However, because the name suffix is forged, the content provider does not have any content corresponding to this name. Consequently, the Interest packet will not be satisfied, and the Interest name will stay in the PITs along the path forever, or until the corresponding PIT entries expire.

The attacker repeats this process with a large amount of different spoofed suffixes. Moreover, distributed compromised system may collaborate to send out Interest packets with the same prefix and diverse suffixes, targeting at the same content provider, as illustrated in Fig. 3. As these Interest packets arrive at the same content provider, its load increases due to these spoofed Interest packets. And of course, this content provider is a victim of the attack. However, we believe that these attacks will not lead to harmful consequences on the content provider's system as severe as that of DDoS in IP. This is because content provider can distinguish these spoofed Interests from legitimate ones by only using a filter, say Bloom Filter, and then discard them. (The Bloom filter stores the names of existing contents on the content provider.) Moreover, these attacking requests cannot "lock" or occupy any memory resources on the content provider, which is the case for SYN flooding and LAND flooding in IP. The largest victims are, therefore, not the end hosts, but the routers, more specifically, the PIT.

These unsatisfied Interest packets in PIT will "lock" memory and computing resources on routers. The accumulated result is a PIT with extremely huge size and high CPU utilization. Surely we can set a timeout value for the PIT entries, but setting the timeout value is quite tricky. A too long timeout makes too many unsatisfied Interest packets stay in the PIT, and a too short timeout may delete a legitimate Interest from PIT before its responding packet returns. Moreover, if these Interest packets arrive at burst, timeout is still helpless to control the memory consumption of PIT.

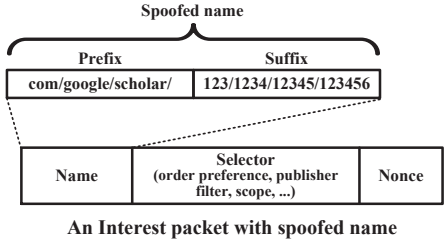

Fig. 2. An Interest packet with spoofed name: a legitimate prefix concatenated with a forged suffix. (We derive the format of Interest packet from [2].)

In a word, this category of NDN DDoS attack incurs extra load on both end hosts and routers, but the routers suffer much more! While this attack can indeed attack a specific content provider, the attack can be blocked using a Bloom Filter by the content provider, or can be alleviated by CDN techniques.

2) Second Category – Interest Flooding Attack: The second DDoS attack is launched by *flooding* Interest packets with *full* forged names by distributed compromised systems, which means these names cannot match any FIB entry in routers. We name it *Interest Flooding* attack, and this attack also takes advantage of the Interest forwarding rule: the Interest packets that are unmatched with FIB should be broadcast or discarded, as shown in Fig. 1(a).

In this paper, we assume that the un-matched packets will be broadcast, especially during the routing establishment stage, and this assumption is supported by the NDN/CCN proposal [1], [2]. Therefore, these Interest packets are duplicated and propagated throughout the entire network until they reach the hosts at the edge of the network. Fig. 4 illustrates that three compromised systems launches Interest flooding attacks, the attacking Interest packets are broadcast at each router, and the number of them is too large that they are not all shown in Fig. 4. Similarly, there is no corresponding content to satisfy these forged Interests. And as with the first category attack, content providers can block these forged Interests by Bloom filters, which are efficient and low-cost. However, Interest Flooding attack still incurs great burden on routers. On one hand, duplicating such a great amount of Interest packets cost a lot of computing resources; on the other hand, without responding Data packets, these Interest packets will stay in the PIT for as long time as possible, which will certainly exhaust memory and computing resources on routers, degrading routers' performance, or even making them crashed.

It seems that this kind of NDN DDoS attack – Interest Flooding attack – has no equivalent in IP, it is NDN-unique. We take the above two categories of NDN DDoS attacks as the basic ingredients, other NDN DDoS attacks can be viewed as a combination or mixture of them.

It's worth pointing out that NDN has no limit on the length of content names, the harmful consequences of both the first and second category of DDoS attacks can be worsened by very long spoofed names, because this will make PIT consume more memory and computing resources.

B. Counter Measures to NDN DDoS attack

Let's first examine the three major counter measures against IP DDoS in Section V. The first counter measure – resource management – is surely helpful for hosts in NDN, but not

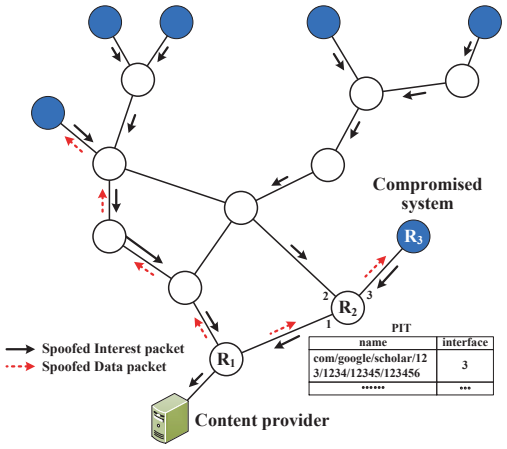

Fig. 3. Comprised systems send out attacking Interest packets targeting at the same content provider. Solid nodes represent compromised systems, and solid arrows stand for Interest packets with spoofed names, while dashed arrows means spoofed Data packets, indicating two of the Interest traceback paths.

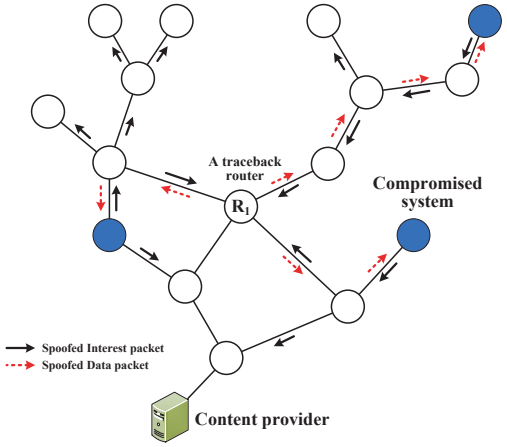

Fig. 4. Interest flooding attacks by compromised systems. Solid nodes represent compromised systems, and solid arrows stand for Interest packets with fully spoofed names (Interest packets are not all shown), while dashed arrows indicate three Interest traceback paths.

in urgent necessity because a simple filter can help to block the attacks. The second counter measure – IP filtering – is not applicable since Interest packets in NDN do not have any information about the source. The third counter measure – packet traceback, which is not directly supported in IP, but can be easily achieved in NDN with the help of PIT! PIT enables symmetric routing, which means Data packets take the same path of its requesting Interest packet, but in the reverse direction. Therefore, this property can be exploited to trace back an attacking Interest to its originator, dampening the NDN DDoS attacks from the source. We name this traceback technique in NDN as *Interest traceback*.

When the PIT size increases at an alarming rate or exceeds a threshold, Interest traceback process is triggered, and it works as follows: a router responds to the attack by tracing back to the Interest originators, i.e. the attackers, by generating spoofed Data packets to satisfy the long-unsatisfied Interest packets in the PIT. These spoofed Data packets are filled with the same forged names as in the Interest packets, as illustrated in Fig 5. The router sends out the spoofed Data packets, and they are forwarded back to the originator by looking up the PIT in intermediate routers. (As long as the Round Trip Time is less than the timeout value, the corresponding PIT entry will not be deleted from PIT in

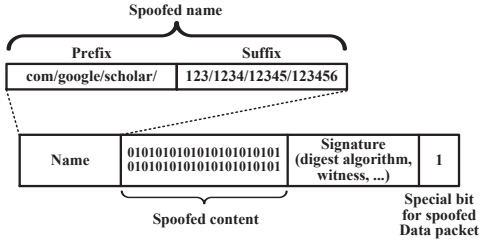

A Data packet with spoofed name and content

Fig. 5. A Data packet with spoofed name and content. The special bit is used to indicate whether this packet is spoofed for traceback or not. (We derive the format of Data packet from [2].)

intermediate routers.)

For example, edge router R_1 in Fig. 3 sends out a spoofed Data packet to satisfied the attacking Interest packet, and it is forwarded to router R_2. Router R_2 looks up the content name of the spoofed Data packet in its PIT, and forwards the spoofed Data packet to interface 3. At last, this spoofed Data packet is forwarded to the originator of the attacking Interest packet – R_3. Similarly, router R_1 in Fig. 4 sends out spoofed Data packets to trace back to the attacker as well.

When the spoofed Data packet arrives the interface via which the attacking Interest packet enters the network, the edge router is notified that the host directly connected with this interface is an attacker. Then, the edge router limits the incoming packet rate of this interface by dropping Interest packets, thereby mitigating the DDoS attack.

For brevity, Fig. 3 only shows two traceback paths (dotted arrows) to the attackers, the rest attackers can also be traced back in a similar way. Fig. 4 describes that a router traces back to the three attackers, the path are highlighted by dotted arrows.

IV. Evaluation

We divide the evaluation into two parts, the first part evaluates the harmful consequences of the DDoS attacks, the second part evaluates the counter measure against the DDoS attacks. Our experimental platform is: Xeon E5500 CPU, 2.27GHz, 15.9G RAM. The topology used is the Rocketfuel [8] topology for EBONE (AS1755), consisting of 172 routers and 763 edges.

A. Harmful Consequences of DDoS Attacks

The harmful consequences of DDoS attacks are mainly reflected by the number of PIT entries increased and CPU cycles consumed due to attack in routers.

We select a sub-topology from EBONE, in which we deploy 100 attackers that send out Interest packets at a rate of 1,000 per second. For the first category of NDN DDoS, we compose the Interest name by choosing a prefix from our Name Set (which has been repeatedly used our previous works [3], [9], [10]), and concatenating it with a forged suffix to generate names of around 1,000 bytes. Subsequently, we launch the attack and the attacking Interest packets are forwarded to the content provider corresponding to the chosen prefix. We will evaluate the number of PIT entries, PIT memory consumption and elapsed CPU cycles incurred by the DDoS attack on the *edge* router directly connected with the content provider, since it is under the most serious attack. The evaluation results are shown in Fig. 6, Fig. 7 and Fig. 8, respectively. Note that the PIT in our experiment is organized by a hash table. Fig. 6 and Fig. 7 respectively indicate the number of PIT entries and memory consumption of PIT incurred by the DDoS attack, not including the number of PIT entries and memory consumption of PIT due to normal

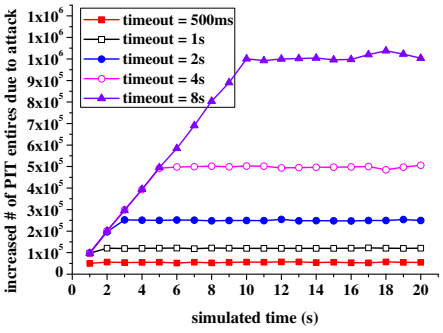

Fig. 6. Increased # of PIT entries due to DDoS attacks.

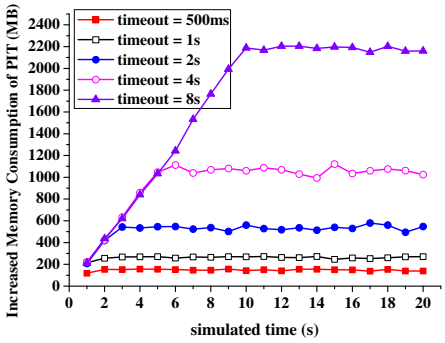

Fig. 7. Increased memory consumption of PIT due to DDoS attacks.

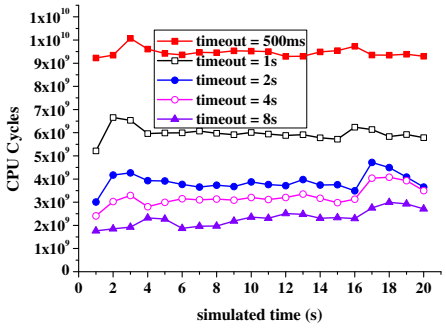

Fig. 8. CPU cycles consumed per second when the router is under DDoS attacks.

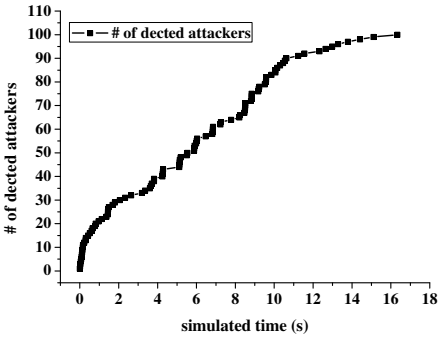

Fig. 9. # of identified attackers.

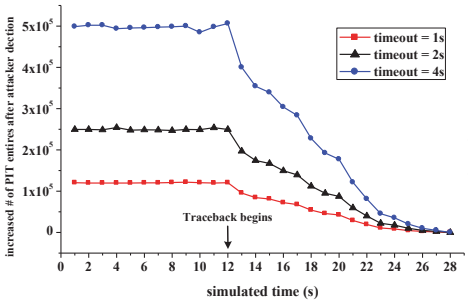

Fig. 10. Increased # of PIT entries due to DDoS attacks declines after Interest traceback begins.

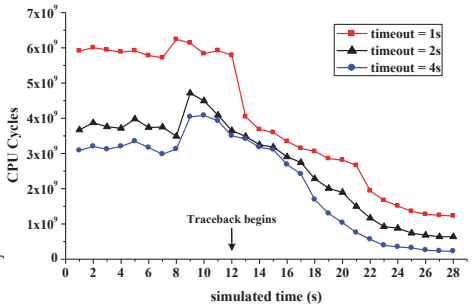

Fig. 11. CPU cycles consumed per second decline after Interest traceback begins.

traffic. Fig. 8 illustrates the CPU cycles consumed on the whole PIT (including PIT entries under normal traffic, around 1.5M entries and 412.78 MB in size [3]).

There are 5 curves in Fig. 6, 7 and 8, representing results of different timeout values for each PIT entry. The timeout values (T) include 500ms, 1s, 2s, 4s and 8s. Clearly, we can see that, the smaller the timeout value, the fewer the PIT entries and less memory consumption of PIT incurred by the DDoS attack. However, smaller timeout value means that the *timeout process* should be run at smaller granularity, which is very time-consuming since it needs to traverse the whole PIT and delete the timed-out entries. (There are smarter algorithms to accomplish this, but they all require auxiliary data structures, which consumes extra memory.) In our experiment, the *timeout process* executes per 100ms for $T = 500ms$, per 200ms for $T = 1s$ per 500ms for $T = 2s$, per 1s for $T = 4s$, and per 2s for $T = 8s$. Fig. 8 demonstrates that the smaller the timeout value, the much more the CPU cycles consumed. Moreover, the demanded CPU frequency is so high that commodity CPUs cannot even satisfy the requirement!

Next, we conduct experiment for the second category of DDoS attack. Interest names are all forged names of around 1,000 bytes in length, and they cannot match any prefix in the FIB. Therefore, these Interest packets will be flooded through the entire network. This time, each router is under serious attack, and results on each router are similar to the statistics in Fig. 6, 7 and 8.

B. Effect of Interest Traceback

When the PIT size increasing at an alarming rate or PIT size exceeds a threshold, it's obvious that this router is under a DDoS attack. We respond to the attack by tracing back to the Interest originators, i.e. the attackers, by generating spoofed Data packets to satisfy the long-unsatisfied Interest packets in the PIT. These spoofed Interest packets are filled with the same forged names as in the Interest packets. (As long as the Round Trip Time (RTT) is less than the timeout value, the corresponding PIT entry will not be deleted from PIT in the intermediate routers.) After identifying an attacker, the edge router limits its incoming packet rate by dropping Interest packets. In this experiment, we take an extreme action – limit the incoming rate to 0, i.e., no packet from the host is allowed into the network once it is identified as an attacker.

The results of the Interest traceback are presented in Fig. 9, Fig. 10 and Fig. 11. Fig. 9 shows the number of identified attackers over time, and we can see that, in our ideal simulation environment, all the 100 attackers are identified by Interest traceback at last. Fig. 10 and Fig. 11 describes that the number of PIT entries and consumed CPU cycles decrease sharply as more and more attackers are detected, which means that the Interest traceback method effectively mitigates the DDoS attack in NDN.

V. Related Work

This section is divided into two parts, the first part introduces the DDoS attack types and corresponding counter measures in IP, the second part describes previous researches on DDoS attacks in NDN.

A. IP DDoS attacks and counter measures

In IP, DDoS has been the most widely reported network attacks, and such attacks belong to the most difficult security problems, since they are easy to launch, hard to prevent, and difficult to trace back. DDoS attacks can be mounted in many types, typically including:

1) SYN flood: A SYN flood occurs when a host sends a flood of TCP/SYN packets, often with a forged sender address, causing the server to spawn a half-open connection.
2) LAND attack: LAND attack is similar to SYN flood, the difference is that the SYN packet's source address and destination address are both IP address of the targeted server, causing the server to lock up.

3) ICMP flood: It relies on mis-configured network devices that allow packets to be sent to all computer hosts on a particular network via the broadcast address of the network, rather than a specific machine.
4) Application level flood: Attacks targeting at the applications that cause server-running software to get confused and fill the disk space or consume all available memory or CPU time.

Faced with the DDoS attack, researchers strive to put forward their counter measures, which can be grouped into three categories: 1) resource management on the victim, 2) address filtering by an Intrusion Detection System (IDS), and 3) IP traceback.

Many previous works responding to DDoS attacks have concentrated on tolerating the attacks by alleviating their harmful effects on the victim by fine-grained resource management [11]–[14]. Although this approach is an effective stop-gap measure, it can not eliminate the problem nor can it dampen attackers. Therefore, it only cures the symptoms, not the disease, and cannot fundamentally solve the problem.

DDoS attackers often conceal the true origin of the attack by using spoofed source IP addresses, therefore, many routers employ a technique called ingress filtering [15] by an IDS to prevent DDoS attacking packets with spoofed IP addresses from propagating. While ingress traffic filtering reduces the success of source address spoofing, it can not prevent an attacker using a forged source address within the permitted prefix filter range.

Since ingress filtering only drops malicious packets, it is not likely to completely eliminate spoofed source IP addresses. As a result, techniques to trace back an attack to the true original source(s) are being developed – ideally stopping an attacker from the very source. Several traceback methods that are immune to spoofed source addresses have been proposed [4]–[7]. [4] and [5] describe a general purpose traceback mechanism based on probabilistic packet marking in the network, which allows a victim to identify the network path(s) traversed by attack traffic. [6] presents a Non-Intrusive IP traceback scheme which learns valid source addresses transiting network routers from sampled traffic under non-attack conditions for anomaly detection. [7] presents a hash-based technique for IP traceback that generates audit trails for traffic within the network, which effectively traces the origin of a single IP packet delivered by the network in the recent past.

B. Researches on DDoS attacks in NDN

There have been several published papers and technical reports regarding the NDN security problem, and many of them are aware of the DDoS attacks in NDN. Tobias Lauinger [16] presented many problems about security and scalability of Content-Centric Networking, including DoS by filling available memory of a router, which resembles our studied DDoS attack very much. Tobias Lauinger also mentioned a immature counter measure but did not verify its effectiveness.

Yoo Chung [17] also identified a kind of basic DDoS attack caused by broadcasting Interest packets, which can overfill the PIT in an NDN router and prevent legitimate Interest packets from being forwarded to where the content is. Matthias Wälisch [18] et al. contributed a collection of new attacks in NDN, including a subset of DDoS attacks, but neither [17] nor [18] provides any counter measure against them.

Paolo Gasti [19] et al. almost put forward an identical DDoS attack scenario as this paper does, but only tentative counter measures were described in that technical report, and no assessment or evaluation was presented, which makes the work very superficial.

All the above work just described different types or scenarios of DDoS attacks when they consider the security problem of NDN, but they did not really study and analyze them in depth, or propose any practical and effective counter measure against them. Furthermore, no quantitative analysis about the DDoS attacks themselves or the counter measures is provided to show the harm of the DDoS attacks or effectiveness of the counter measures. In contrast, this paper studies only a specific and concrete scenario of NDN DDoS attacks and estimates its damage to NDN routers. Moreover, an effective counter measure – Interest traceback – is proposed and evaluated.

VI. Conclusion

As an instantiation of the Content-Centric Networking approach, NDN tries to provide better security and privacy support than current Internet does. This paper presents a specific and concrete scenario of DDoS attacks in NDN, demonstrating the possibility of NDN DDoS attacks. Perpetrators make use of NDN's Interest packet forwarding rule to send out Interest packets with spoofed names as attacking packets. Afterwards, we identify the victims of NDN DDoS attacks include not only the hosts, but also the Pending Interest Tables (PIT) within routers. Moreover, the Pending Interest Table is the largest victim. We propose an effective counter measure called Interest traceback to fight against NDN DDoS attacks. At last, we assess the harmful consequences brought by NDN DDoS, and evaluation results demonstrate that the Interest traceback method effectively mitigates the NDN DDoS attacks.

References

[1] L. Zhang, D. Estrin, V. Jacobson, and B. Zhang, "Named Data Networking (NDN) Project," in *Technical Report, NDN-0001*, 2010.
[2] V. Jacobson, D. K. Smetters, J. D. Thornton, M. Plass, N. Briggs, and R. Braynard, "Networking named content," in *Proc. of CoNEXT*, 2009.
[3] H. Dai, B. Liu, Y. Chen, and Y. Wang, "On pending interest table in named data networking," in *Proceedings of ACM/IEEE ANCS*, Austin, Texas, USA, Oct 2012, pp. 211–222.
[4] S. Savage, D. Wetherall, A. Karlin, and T. Anderson, "Practical network support for ip traceback," in *Procedings of ACM SIGCOMM'00*, 2000, pp. 295–306.
[5] S. Savage and D. Wetherall, "Network support for ip traceback," *IEEE/ACM Transactions on Networking (TON)*, vol. 9, no. 3, pp. 226 – 237, 2001.
[6] V. L. L. Thing, M. Sloman, and N. Dulay, "Non-intrusive ip traceback for ddos attacks," in *Proceedings of ACM ASIACCS'07*, 2007.
[7] A. C. Snoeren, C. Partridge, L. A. Sanchez, C. E. Jones, F. Tchakountio, S. T. Kent, and W. T. Strayer, "Hash-based ip traceback," in *Procedings of ACM SIGCOMM'01*, 2001, pp. 3–14.
[8] Rocketfuel, "http://www.cs.washington.edu/research/networking/rocketfuel/."
[9] Y. Wang, Y. Zu, T. Zhang, K. Peng, Q. Dong, B. Liu, W. Meng, H. Dai, X. Tian, Z. Xu, H. Wu, and D. Yang, "Wire speed name lookup: A gpu-based approach," *To appear in Proceedings of NSDI'13*, 2013.
[10] W. Yi, H. Keqiang, D. Huichen, M. Wei, J. Junchen, L. Bin, and C. Yan, "Scalable name lookup in ndn using effective name component encoding," in *Proceedings of IEEE ICDCS'12*, 2012.
[11] O. Spatscheck and L. Peterson, "Defending against denial of service attacks in scout," in *Proceedings of the USENIX/ACM OSDI*, 1999, pp. 59–72.
[12] G. Banga, P. Druschel, and J. Mogul, "Resource containers: A new facility for resource management in server systems," in *Proceedings of the USENIX/ACM OSDI*, 1999, pp. 45–58.
[13] C. Meadows, "A formal framework and evaluation method for network denial of service," in *Proceedings of the IEEE Computer Security Foundations Workshop*, 1999.
[14] B. Schroeder and M. Harchol-Balter, "Web servers under overload: How scheduling can help," *ACM Transactions on Internet Technology (TOIT)*, vol. 6, no. 1, pp. 20 –52, 2006.
[15] P. Ferguson and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing," *RFC 2267, IETF*, 1988.
[16] T. Lauinger, "Security & scalability of content-centric networking," *Master's Thesis, Technischeät Universit Darmstadt*, 2010.
[17] Y. Chung, "Distributed denial of service is a scalability problem," *ACM SIGCOMM CCR*, vol. 42, no. 1, pp. 69 – 71, 2012.

[18] M. Wahlisch, T. C. Schmidt, and M. Vahlenkamp, “Backscatter from the data plane – threats to stability and security in information-centric networking,” 2012.
[19] P. Gasti, G. Tsudik, E. Uzun, and L. Zhang, “Dos & ddos in named-data networking,” *arXiv:1208.0952v2*, 2012.