International Journal of Information Security
https://doi.org/10.1007/s10207-020-00500-z

REGULAR CONTRIBUTION

ChoKIFA+: an early detection and mitigation approach against interest flooding attacks in NDN

Abdelmadjid Benarfa[1] · Muhammad Hassan[2] · Eleonora Losiouk[2] · Alberto Compagno[3] · Mohamed Bachir Yagoubi[1] · Mauro Conti[2]

Abstract

Several ongoing research efforts aim to design potential Future Internet Architectures, among which Named-Data Networking (NDN) introduces a shift from the existing host-centric Internet Protocol-based Internet infrastructure towards a content-oriented one. However, researchers have identified some design limitations in NDN, among which some enable to build up a new type of Distributed Denial of Service attack, better known as Interest Flooding Attack (IFA). In IFA, an adversary issues not satisfiable requests in the network to saturate the Pending Interest Table (PIT) of NDN routers and prevent them from properly handling the legitimate traffic. Researchers have been trying to mitigate this problem by proposing several detection and reaction mechanisms, but all the mechanisms proposed so far are not highly effective and, on the contrary, heavily damage the legitimate traffic. In this paper, we propose a novel mechanism for IFA detection and mitigation, aimed at decreasing the memory consumption of the PIT by effectively reducing the malicious traffic that passes through each NDN router. In particular, our protocol exploits an effective management strategy on the PIT, through which the Malicious Interest (MIs) already stored in the PIT are removed and the new incoming MIs are dropped. In addition, the proposed countermeasure provides an additional security wall on the edges of the network to detect and mitigate the attack as early as possible and improve the network health, i.e., routers PIT occupancy during IFA. To evaluate the effectiveness of our work, we implemented the proposed countermeasure on the open-source ndnSIM simulator and compared its effectiveness with the state of the art. The results show that our proposed countermeasure effectively reduces the IFA damages both in terms of preserved legitimate traffic and availability of routers PIT. Considering the legitimate traffic, the amount of Benign Interests preserved by our approach increases from 5% to 40% with respect to the preservation guaranteed by the state-of-the-art solutions. Concerning the routers PIT availability, our approach guarantees that the 97% of the PIT size is left free for handling the legitimate traffic.

Keywords NDN · DDoS attack · IFA · Congestion · PIT management

1 Introduction

The existing Internet architecture was deployed in the early 70s to address the elementary communication requirements which were obligatory at that period. In particular, it was designed to guarantee a reliable host-to-host interconnectivity. However, over the past few years the Internet has increasingly shown a poor match with its initial design and it started moving towards a content distribution and retrieval paradigm. Today, the Internet produces immense and ever-growing traffic volumes, mainly due to the multiple

✉ Abdelmadjid Benarfa
a.benarfa@lagh-univ.dz

Muhammad Hassan
hassan@math.unipd.it

Eleonora Losiouk
elosiouk@math.unipd.it

Alberto Compagno
acompagn@cisco.com

Mohamed Bachir Yagoubi
m.yagoubi@lagh-univ.dz

Mauro Conti
conti@math.unipd.it

[1] Laboratory of Computer Science and Mathematics, University of Laghouat, Laghouat, Algeria

[2] Department of Mathematics, University of Padova, Padua, Italy

[3] Cisco Systems, Issy-les-Moulineaux, France

Published online: 08 May 2020

distribution of the same popular contents, which inevitably introduce high costs for network operators such as additional Content Delivery Networks (CDN) deployment. Numerous solutions have been proposed in recent years to narrow the gap between the Internet design and its current usage [32,41]. Among the different solutions proposed under the Future Internet Architecture (FIA) program sponsored by the National Science Foundation (NSF), Name Data Networking (NDN) [46] is the most promising architecture. NDN follows the Content-Centric Networking (CCN) approach, it uses caches in the network, multiparty communication through replication, and models of interaction that decouple senders and receivers [4,22]. By explicitly addressing the data (content) instead of the physical location of the hosts in the network, NDN converts data into the first-class entity. Thus, instead of a direct connection between two hosts, an NDN consumer directly requests the name of the content by issuing an interest. The network then handles the request by efficiently finding and retrieving back the closest copy of the relevant content. This decoupling of time and space, between request resolution and content transfer, enables NDN to guarantee security by design [48], resilience to disruptions, content distribution, mobility [11], and storage [43], as native features belonging to the network architecture.

The purpose of this paper is to propose a detection and mitigation solution against one of the most significant NDN attacks: the *Interest Flooding Attack* (IFA). IFA is an NDN-customized DDoS attack, in which the adversaries send a high number of spoofed interest packets to consume the memory resources of on-path routers. As a result, routers drop all subsequent incoming legitimate requests, thus making the overall network unresponsive.

1.1 Motivation and contribution

IFA has been extensively addressed by previous works, either in terms of detection or in terms of mitigation [3,12,13,38,39]. However, all the existing solutions are affected by one or more among the following limitations:

- A variable effectiveness of the detection approach according to the location: the attack detection is not robust and accurate, if applied close to content providers and victims, and if the volume of adversarial traffic is large [35];
- The amount of damage applied on the legitimate traffic: the legitimate traffic is likely to be damaged, since most of the proposed countermeasures [3,12,38] limit the rate of incoming traffic and is not able to differentiate between Benign Interests (BIs) and Malicious Interests (MIs), thus resulting in unfair punishments;
- The overhead introduced on the routers that both detect and mitigate the attack: during the attack detection, most of the approaches used by routers are likely to encounter harmful consequences;
- The network overhead caused by collaborative approaches: the proposed collaborative mechanisms [3,12,13,35] introduce unnecessary overhead due to the extra messages exchanged among routers.

To improve the detection and mitigation against IFA, by overcoming the above-mentioned limitations of the state-of-art solutions, we propose an efficient countermeasure, named as Choose To Kill IFA (ChoKIFA), which mitigates the damages caused by IFA by differentiating the malicious traffic from the legitimate one, and by reducing the former. To distinguish between BIs and MIs, the proposed protocol applies on each router an active queue management scheme, i.e., CHOose and Keep for responsive flows, CHOose and Kill for unresponsive flows (CHOKe) [33]. For each incoming interest, ChoKIFA evaluates several conditions before saving it in the router PIT by differentiating and penalizing only the malicious ones, thus preventing the propagation of the attack in the network. The preliminary results concerning the effectiveness of ChoKIFA are reported in [7]. In this paper, we propose an enhanced version of ChoKIFA, named as *ChoKIFA+: An Early Detection and Mitigation Approach against Interest Flooding Attacks in NDN*. In ChoKIFA+, the edge routers, which consumers are directly connected to, provide an additional security wall to detect the attack as early as possible and to limit the damage of the overall network. With respect to ChoKIFA, ChoKIFA+ further improves the network health status by reducing the routers PIT occupancy. The major contributions of this work are as follows:

- Design and implementation of a detection and mitigation mechanism against IFA, named as ChoKIFA [7], that relies on the CHOKe approach. ChoKIFA penalizes the malicious traffic by both dropping the new incoming MIs and removing the ones already stored in the PIT, without delay and without requiring information about the global network state;
- Design of an additional security mechanism applied on edge routers, named as ChoKIFA+. This approach aims to quickly identify and block IFA at the edge routers, thus improving the network health during the attack;
- Evaluation of the efficiency and effectiveness of ChoKIFA and ChoKIFA+, both implemented on the ndnSIM[1] simulator [2], by comparing them with the state-of-the-art mitigation approaches [3]. The results show that ChoKIFA+ effectively mitigates the IFA effects by guaranteeing an interest satisfaction rate (ISR) up to 98% and by reducing up to 40% the number of false positives. Moreover, comparing to ChoKIFA, ChoKIFA+ further

[1] ndnSIM implements the NDN protocol stack on NS-3 simulator.

improves the network health during IFA by reducing the PIT size up to 99 %.

1.2 Organization

The remainder of this paper is organized as follows. We present the NDN architecture in Sect. 2. Section 3 is reserved to present IFA. The related work on IFA and DDoS countermeasures exploiting AQM techniques is presented in Sect. 4. Section 5 briefly describes the proposed protocol including system, adversary model and working methodology of ChoKIFA. In Sect. 6, we present the implementation, evaluation and comparison of proposed countermeasure against IFA and state-of-the-art mitigation mechanisms. Finally, Sect. 7 concludes the paper.

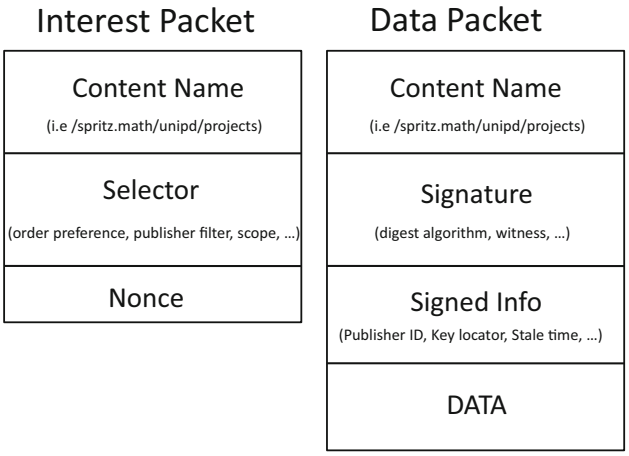

Fig. 1 NDN packet types

2 Background

Today, the Internet has shifted from connections between hosts to a global distribution and retrieval of contents in huge amount. This clear difference between the present Internet architecture and its current usage uncovers all its limits. To this end, numerous research efforts aim to move the current Internet towards a new architecture, called Information Centric Networking (ICN) [41]. In particular, ICN presents the data as a "first class" entity, by focusing on *what* is the content rather than *where* is the content. Among diverse ICN styles, NDN [44,45] and CCN [22] projects have earned significant recognition in both Academia and Industry.

2.1 NDN architecture components

NDN is based on a *requester-driven* communication model between two types of hosts: clients, called consumers, and servers, identified as producers. In contrast to IP, where a content is explicitly exchanged between nodes, NDN consumers request content pieces from the network, without being aware of producer location. In particular, when an NDN client requests a specific content, it sends out a unique interest associated with that content, where the interest is identified by a Uniform Resource Identifier (URI) with a routable name scheme, e.g., `/example.com/video4u/examples.mp3`.

Figure 1 shows the format of the two NDN packet types, i.e., *Interest* and *Data* packets. The most imperative information in both packets is the content name, which is the identifier of the requested content and of the associated data. The other two fields in the interest packet are selector and nonce. The selector allows to: (i) specify whether the name is the full name including the digest, or the full name excluding the digest, or the content name is known to be in a range of legitimate components (i.e., `MinSuffixComponents/MaxSuffixComponents`); (ii) provide the name of the key used to sign the corresponding data packet (i.e., `PublisherPublicKeyLocator`); (iii) choose whether to exclude the list and/or ranges of name components from the corresponding content packet (i.e., `ExcludeFilter`).

Besides the content name, the data packet includes the Signature, the Signed Info and the Data itself. The Signature and the Signed Info are used to verify the integrity of the received content, which is signed by the content provider [11].

2.2 NDN forwarding process

NDN introduces additional responsibilities on the routers by enabling router-side content caching and interest aggregation [22]. Each NDN router implements three foremost data structures, named as:(i) Forwarding Information Base (FIB); (ii) Pending Interest Table (PIT); (iii) Content Store (CS).

Figure 2 illustrates the functionality and key components of an NDN router. In particular, the FIB is used to perform a lookup operation to determine the interfaces which the incoming interests should be forwarded to (e.g., it includes the interest prefix and the interface number). The second lookup table is the PIT, which comprises of all the entries (e.g., interest prefix, arrival interface) of outstanding forwarded interests. Once an interest for a specific content is received, NDN router first checks whether the demanded content already exists in the CS. If the content is not available in the CS, the router looks in the PIT for a pending interest issued for the same content. If there is no entry, the router forwards the interest towards its destination and adds a new entry in the PIT, with the associated arrival interface. On the contrary, if there is already an entry in the PIT, additional interests requesting the same content will not be forwarded, but only added to the existing entry. Later, when the requested content arrives, all the pending interests for it are satisfied

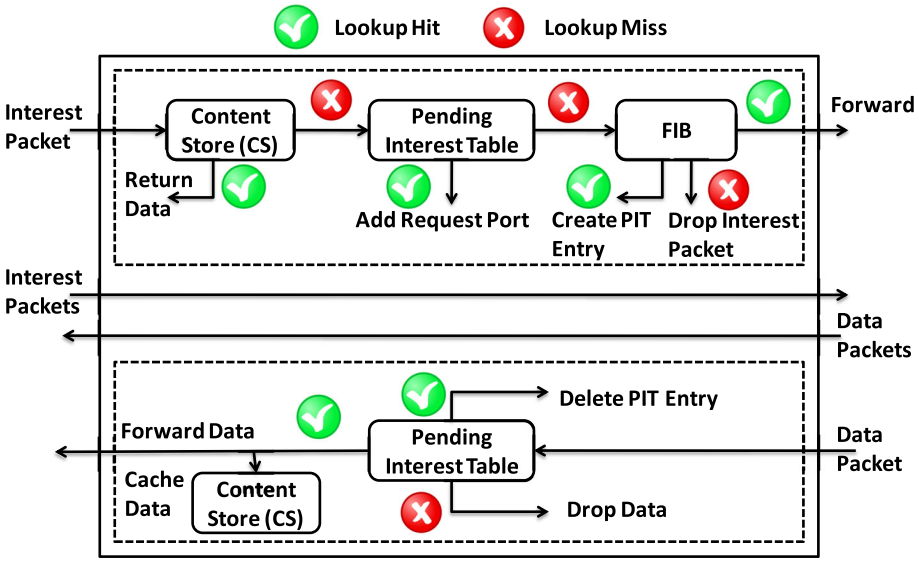

Fig. 2 NDN node architecture

just by sending a copy of the content back to all the consumers who issued an interest before. Finally, the requested data packet is forwarded to the consumer just by traversing the reverse trail of the associated interest [46].

In NDN, if no data are returned for a given interest (e.g., a router not forwarding an interest or a content producer not having the associated content), no error packets are generated. The router PIT entries for such unsatisfied interests are removed after the expiration time has passed. As a result, a consumer can determine whether to reissue the same interest after a certain timeout.

3 Interest flooding attacks

In *Interest Flooding Attack* (IFA), the adversary aims at exhausting routers memory and resources by flooding them with a huge amount of requests for contents that are not available in the network and preventing them from handling the legitimate requests from benign consumers. To achieve his aim, the adversary exploits two essential NDN features [35]: (i) the interest forwarding strategy based on the longest prefix match and (ii) the saving of the forwarded interests into the PIT.

Considering the first feature (i.e., the longest prefix match), the adversary can exploit it to generate the MIs that will flood the routers. After identifying a prefix name, for which there is already a forwarding rule in the FIB, the adversary can append a random value to it and transmit in the network, even though it is not satisfiable. Considering the second feature (i.e., the forwarded interests that are saved into the PIT), each entry is kept in the PIT for a certain amount of time, after which it is removed. Thus, the adversary can easily fill the whole PIT capacity.

IFAs are classified into three types according to the type of content requested by the adversary [19]: (i) existing or static content; (ii) dynamically generated content; (iii) nonexistent content.

Concerning the first IFA type, several zombies from multiple locations in the network generate a large number of interests for an existing content, thus exhausting the producer, which starts dropping the interests. The interests will then remain in routers PIT until the expiration time has passed. However, due to the in-network content caching, that prevents the high number of requests from reaching the producer, the impact of such IFA type is quite restricted. In the second IFA type, the adversary sends dynamically generated interests for existing contents, again propagated till the producer(s) and resulting in bandwidth consumption and PIT exhaustion. Moreover, the targeted producer consumes its computational resources due to the content signing (i.e., per-packet operation). Finally, in the third IFA type, the adversary requests unique, nonexistent and unsatisfiable contents that cannot be collapsed by routers and are routed towards the producer. Such interest packets consume the router PIT, where they are stored, until their expiration time has passed. Therefore, a massive number of nonexistent interest packets in the PIT makes the BIs be dropped in the network.

In this paper, we evaluate our ChoKIFA+ solution against all the three IFA types, having routers and legitimate traffic the primary victims of the attack. Using a valid name $prefix$, there are many ways for an adversary to generate unsatisfiable interests: by appending a random value to the name $prefix$ (such interests are propagated towards the producer and are never satisfied), by replacing the `PublisherPublicKeyDigest` field with a random value (no public key would match this value, therefore, will never be satisfied) or by setting the `InterestExclude`

`filter` to exclude all existing content starting with $/prefix$ (the interest can never be satisfied as it concurrently requests and excludes the same content).

4 Related work

In the present section, we describe the existing solutions for IFA mitigation in NDN. Then, we illustrate the role of various active queue management schemes mitigating DDoS attacks in existing IP architecture.

4.1 Solutions mitigating IFA in NDN

In this section, we discuss the approaches and limitations of the state-of-the-art mitigation solutions against IFA. In particular, in order to classify the related work, we refer to the following three IFA mitigation categories based on their functionalities [37]: (i) rate limiting-based countermeasures, in which the mitigation relies on throttling down the overall incoming traffic in an autonomous/collaborative manner; (ii) statistical modeling-based countermeasures, where the mitigation mechanism is based on the statistical information of PIT occupancy; (iii) the other countermeasures, that include the approaches based on the update of routers structures (i.e., forwarding information).

Below, we comprehensively describe the relevant work under each category.

4.1.1 Rate limiting-based countermeasures

Afanasayev et al. [3] proposed four different methods to deal with IFA. The first method introduces a simple limit on the interfaces, based on the physical capacity of the links and resulting in an under-utilization of the network. The second method is an adaptation of the token bucket algorithm [24] providing per-interface fairness. In particular, the algorithm regulates the number of outgoing interests by limiting the assigned tokens to a specific outgoing interface. The major drawback of this method is that tokens are assigned without discriminating between BIs and MIs. Thus, not all MIs are dropped, while some BIs could. The third method is based on the per-interface ratio between the interests sent and the corresponding data packets received, which is also defined as the "satisfaction-based interest acceptance". In this method, tokens are fairly distributed among all the incoming interfaces according to their interest satisfaction rates. The drawback of this method is that the router decision to forward or discard an interest packet relies on a router's local statistics, which is the router's interest satisfaction rate. Due to this reason, the probability of legitimate interests being forwarded declines as the number of hops/routers between the consumer and the producer increases [3]. The last method is a collaborative approach defined as "satisfaction-based push-back". In this case, each router sets an explicit limit value for each incoming interface, and announce this value to all downstream routers. This method has shown to be more effective than the previous ones, but the legitimate stream is still influenced, especially when the path is long. Moreover, it creates unnecessary signaling overhead in the network.

Similar to the works of Afanasyev et al. [3], other detection and reaction solutions have been proposed so far based on an independent or a collaborative approach. For the first case, the attack detection is based on network traffic analysis and/or PIT usage [12,19], while the reaction consists in reducing the incoming/outgoing traffic, independently on each router. Vassilakis et al. [38] also proposed a similar mechanism that relies on the anomalous behavior of the consumers, to detect the attack, and on the reduction in the requesting rate of the traffic from the detected nodes, to block the attack. In addition to the traditional IFA detection criteria, Benmoussa et al. [8] considered also the network congestion to avoid the false positives generated by unintentional IFA phenomenon. Hence, to detect whether there is an intentional or an unintentional IFA, the authors use three parameters: the ISR (i.e., the number of interests issued over the number of data packets received); the incoming interest rate (i.e., the amount of interest packets arriving at a particular interface of an NDN router in a given amount of time); the network congestion, measured considering the number of timed-out interests and of NACK packets. Considering the collaborative component of the above-mentioned mechanisms, this consists in an information exchange among intermediate routers, which inevitably generates a signaling overhead. Moreover, some of those mechanisms [12,13,38] send a "push-back alert" to the downstream interfaces to reduce the data rate. In [13], the authors proposed a collaborative countermeasure known as "interest traceback", which detects an attack when the PIT size increases and produces artificial spoofed data packets for each interest stored in the PIT. Eventually, during the attack, the data packets trace the interests generators. The limitation of the approaches includes an excessive amount of additional traffic in the network, that results in the depletion of bandwidth and performance.

In summary, the main limitations of the rate limiting-based countermeasure are the following ones: (i) lack of differentiation between benign and malicious traffic; (ii) non elimination of the interests already stored in the PIT; (iii) additional overhead due to the massive message exchange among the routers.

4.1.2 Statistical modeling-based countermeasures

In [28], the authors face IFA through a detection scheme focused on Statistical Hypothesis Testing Theory, that relies on the Neyman-Pearson bi-criteria approach by providing a

test that does not depend on router characteristics or measured values. Such framework considers two cases: (i) a scenario in which all traffic parameters are known (in this context, the optimal test is designed and its statistical performance is given); (ii) a linear parametric model proposed to estimate unknown parameters and to design a practical test. However, the evaluation only uses a simple binary tree graph of eight clients and one adversary. It is difficult to analyze the efficacy of the scheme for larger networks or during distributed attacks.

The authors in [30] also leverage the hypothesis testing theory to develop a generalized likelihood ratio test, adapted to evolve IFA attacks, especially, in the context of coupling NDN with IP, which can hardly be addressed by traditional solutions.

In [29], the authors proved the feasibility of IFA in a real NDN deployment. To this aim, the authors proposed a comprehensive set of 18 NFD metrics. For each metric, a microdetector is designed to capture any abnormal variation from the metrics normal behavior. The relevance of the microdetector design was evaluated through its performance against IFA in a testbed. To assess the generality of the approach, this should be tested with other attacks and larger topologies.

In [21], the authors proposed a Theil-Based Countermeasure (TC) to detect the distributions of BIs and MIs in the NDN routers and identify an IFA. Each NDN router records the name of each interest packet received. The router can use the statistical distribution of the names of the interest packets to detect the IFA. When an adversary launches an IFA, the occurrence frequency of the fake names in the MIs will increase significantly. The shift in the value of the Theil entropy can be used to determine the networks situation. In addition, the Theil entropy can divide the interest packets into groups based on a preset rule to evaluate the contribution of intra-group and inter-group differences. Then, under IFA, the TC detects the attack based on the change of intra-group heterogeneity. When an IFA is detected, a trace-back method on MIs can be used to find the adversary's position and avoid further attempts. Despite that solution induces additional signaling overhead, its evaluation uses only a simple binary tree graph with eight clients and two adversaries. The effectiveness of the scheme for larger networks or during distributed attacks is difficult to analyze.

In summary, this category also has some drawbacks: (i) the attack detection is difficult close to attack sources, resulting in late detection and no reaction; (ii) the detection takes place after wasting a huge amount of resources in several regions of the network, and it likely cannot prevent the attack before causing a severe damage; (iii) mitigation mechanisms are also unable to remove malicious traffic from the PIT; (iv) the applied detection algorithms do not distinguish MIs from BIs. Hence, they may also harm legitimate traffic.

4.1.3 Other countermeasures

This category of DoS mitigation includes approaches that change routers' structures, such as PIT and CS. Wang et al. [39] proposed a mitigation mechanism called Disabling PIT Exhaustion (DPE) which diverts all the MIs out of the PIT. In particular, the state information of the MIs is recorded separately in the name of malicious-list instead of PIT. It introduces also a packet marking scheme to help with data packets forwarding without PIT. Although this method prevents the PIT from being overloaded, MIs will still be forwarded, which may lead to the network congestion and poor bandwidth utilization. Other disadvantages include extra packet overhead and processing load.

InterestFence [14] detects IFA based on content servers rather than routers to guarantee accurate detection. All content items with the same prefix within a content server have a Hash-based Security Label (HSL) to claim their existence, and a HSL verification method is securely transmitted to related routers to help filtering and cleaning IFA traffic in transit. InterestFence consists of three key functional entities: InterestFence-enabled router, InterestFence-enabled content server and the communication between them.

In [47], the authors proposed a Charging/Rewarding mechanism based on Hidden Markov Model (HMM) to defend against IFA. The approach is based on the external characteristic parameters of the consumers to establish HMM and then detect the malicious user. In particular, the HMM is established by the Baum-Welch algorithm. In order to limit the malicious users and stimulate the legitimate consumer, the edge router will charge consumers when providing forward service and reward legitimate users.

However, such mechanisms will still be responsible for forwarding MIs, resulting in network congestion and starvation of legitimate clients. In addition, the mechanisms also put additional processing burden on the routers and increase packet overhead.

In this paper, we propose a solution that exploits an Active Queue Management (AQM) mechanism to defend against IFA in NDN. Our solution is the first one that removes the MIs from the PIT as a mitigation strategy; that works both as a detection and reaction approach; that avoids the drawbacks of the first and second categories of the state of the art and that does not add new additional structures to routers, like the third category of the existing solutions.

In Table 1, we extend the review of the existing IFA solutions [34] by providing a detailed comparison of the existing countermeasures along with their functional details.

Table 1 Comparison of IFA countermeasures in NDN

	Type of solution		Detection			Detection mechanism	Mitigating mechanism	Stateless	Remove fake interest	Topology tested
	Reactive	Proactive	Per router	Per interface	Per prefix					
Poseidon [12]	✓	✗	✗	✓	✗	ISR + PIT usage	Push back	✗	✗	AT&T
Traceback [13]	✓	✗	✓	✗	✗	Thresholds	Spoofed data	✗	✗	Rocket fuel topology
Tocken bucket [3]	✓	✗	✓	✗	✗	Token	Trace back	✗	✗	ISP topology
Simple limits [3]	✓	✗	✓	✗	✗	PIT usage	Rate limit	✗	✗	ISP topology
Satisfaction based [3]	✓	✗	✗	✓	✗	ISR based	Rate limit	✗	✗	ISP topology
mTBAD [27]	✓	✗	✗	✓	✗	Thresholds	m-NACK	✗	✗	AT&T
Satisfaction pushback [3]	✓	✗	✗	✓	✗	Thresholds	Push back	✗	✗	AT&T
TDM [40]	✓	✗	✓	✗	✗	Thresholds	Rate limit	✗	✗	AT&T
Congestion-Aware IFA [8]	✓	✗	✗	✓	✗	ISR + Interest rate + congestion	Rate block	✗	✗	AT&T
Hypothesis Testing [28]	✓	✗	✗	✓	✗	Hypothesis Testing	✗	✗	✗	testbed
Reliable Detection [30]	✓	✗	✗	✓	✗	Hypothesis Testing	✗	✗	✗	testbed NDN with IP
Security Monitoring [29]	✓	✗	✗	✓	✗	18 metrics + monitoring plane	✗	✗	✗	testbed
Theil-Based [21]	✓	✗	✗	✓	✓	Theil entropy	Traceback	✗	✗	Small tree
InterestFence [14]	✓	✗	✗	✓	✓	hash-based security label (HSL) filter	✗	✗	✗	Tree
Charge/Reward [47]	✓	✗	✗	✓	✓	HMM	✗	✗	✗	Tree
Decoupling DPE [39]	✗	✓	✗	✗	✓	Malicious list	Decoupling	✗	✗	AT&T
ChoKIFA [7]	✓	✗	✗	✓	✓	AQM technic	Remove fake interest	✓	✓	AT&T
ChoKIFA+ AQM technic Edge router	✓	✗	✗	✓	✓	AQM technic + Edge router	Remove fake interest + rate limit	✓	✓	AT&T

4.2 Mitigating DDoS through active queue management schemes in IP

To mitigate DDoS attacks in IP, the congestion handling techniques, such as AQM, have gathered significant attention from the research community [6,20,23,42]. AQM methods are classified into two categories according to their functionality and to the type of traffic they are able to handle [5]. The first category aims to provide fairness during network congestion, when the incoming traffic consists of only responsive flows (i.e., to which server responds). Random Early Detection (RED) [17], BLUE [16], and Adaptive Virtual Queue (AVQ) [25] approaches belong to the first category. On the other hand, the second category provides fairness when the incoming traffic consists of both responsive and unresponsive flows, such as CHOKe [33], Stochastic Fair Blue (SFB) [15], and Fair Random Early Detection (FRED) [26].

RED tries to acquire a better route queue stability by estimating the level of congestion in the router buffer and drops packets accordingly, by using an exponentially weighted moving average (EWMA) of the queue length. One of the RED limitations is that it is unable to identify unresponsive flows. Thus, it requires a significant parameters tuning to attain optimal results. To address this limitation, several methods, based on the idea of (or function with) RED, have been proposed: CHOKe [33], xCHOKe [10], RECHOKe [42] and FRED [26]. In particular, CHOKe is a stateless technique that tries to handle unresponsive flows by identifying and penalizing them through the drop of their packets.

5 IFA mitigation through an active queue management scheme

In this paper, we take a footstep in the direction of identifying and differentiating MIs from BIs during IFA, in order to selectively stop the malicious traffic. In particular, we exploit an AQM algorithm [33] (i.e., CHOose to Kill malicious Interest, CHOose to keep genuine Interest for IFA - ChoKIFA) to stabilize routers PIT, by dropping the incoming MIs and removing the ones already stored in the PIT, without any global knowledge of the network (i.e., state information).

5.1 Approach overview

The fundamental idea behind ChoKIFA is to exploit the PIT state, which provides adequate statistics regarding the incoming and outgoing interest packets, and use it to identify and drop MIs. When an interest arrives at a router, ChoKIFA randomly draws an interest from the PIT and compares it with the incoming one. If both interests belong to the same traffic flow,[2] then they are both dropped. Otherwise, the randomly drawn interest is left stored in the PIT and the incoming one is stored in the PIT according to a probability that depends on the level of PIT occupancy. The intuition of ChoKIFA is that, during IFA, the router PIT is likely to have more entries filled with MIs.[3] Thus, it is more likely that MIs will be chosen for the comparison with the incoming interests and that the malicious traffic will be stopped.

5.2 Adversary model

Our adversary follows the third IFA type, as illustrated in Sect. 3, by generating a massive amount of MIs, requesting non-existing contents. We assume that the adversary aims to saturate the routers PIT through the rapid generation of a large number of MIs [3,12,39], so that, when the PIT is full, the incoming BIs are dropped. In comparison to the first and the second IFA types, the third one has a bigger impact for the following reasons: (i) the MIs referring to non-existing contents are kept stored in the router PIT; (ii) the sending rate of MIs does not depend on the bandwidth allocated by the R to content packets, or on the capability of the adversary to receive content; (iii) MIs cannot be satisfied by the contents saved in router caches; (iv) if generated in a smart way (e.g., with a random component at the end of each name), MIs never collapse until the interests decay.

Without loss of generality, we assume that the adversary is able to corrupt a set of C (i.e., botnet) and use them to send the attack. This assumption is realistic and justified by the current scenarios of DDoS attacks [18]. The adversary issues MIs referring to a $prefix$, which is registered by P. To make each MI reach P, the adversary attaches a random value to each interest, i.e., $prefix/Rnd$ where Rnd is a random string. Moreover, the 50% of C in the whole network are used in the botnet [3]. Lastly, similar to C, the adversary starts sending MIs at time t. Table 2 summarizes the notations used in the paper.

5.3 ChoKIFA: CHOose to kill interest flooding attack

To be effective against IFA, a mitigation approach has to differentiate between MIs and BIs. To this purpose, ChoKIFA relies on the traffic flow to differentiate and penalize the MIs from BIs.

Unlike in the IP architecture, where the traffic flow is measured through accountable attributes (e.g., source/destination address, interface number, number of packets/bytes sent forward and backward [9]), in NDN, the traffic flow is based

[2] The NDN traffic flow measurement differs from the IP one and we present the comparison between them in Sect. 5.3.

[3] Recall that unsatisfiable interests refer to non-existing contents and saturate the PIT.

Table 2 Summary of the notations used

Notation	Meaning
Adv	Adversary
C	Consumer
P	Producer
R	Set of all routers
MI	Malicious interest
BI	Benign interest
r_i	i-th router, $r \epsilon \lvert R \rvert$
r_i^j	j-th interface of router i
$\delta(r_i^j)$	Interest satisfaction rate
$\delta_{th}(r_i^j)$	Threshold for interest satisfaction rate
$\rho_{avg}(r_i^j)$	Average PIT size
w_ρ	Weight factor for EWMA
$\rho_{size}(r_i^j)$	Actual PIT size
$\rho_{th}^{\min}(r_i^j)$	Minimum threshold for PIT size
$\rho_{th}^{\max}(r_i^j)$	Maximum threshold for PIT size
P_b	Interest drop probability
$P_{\max}$	Maximum drop probability

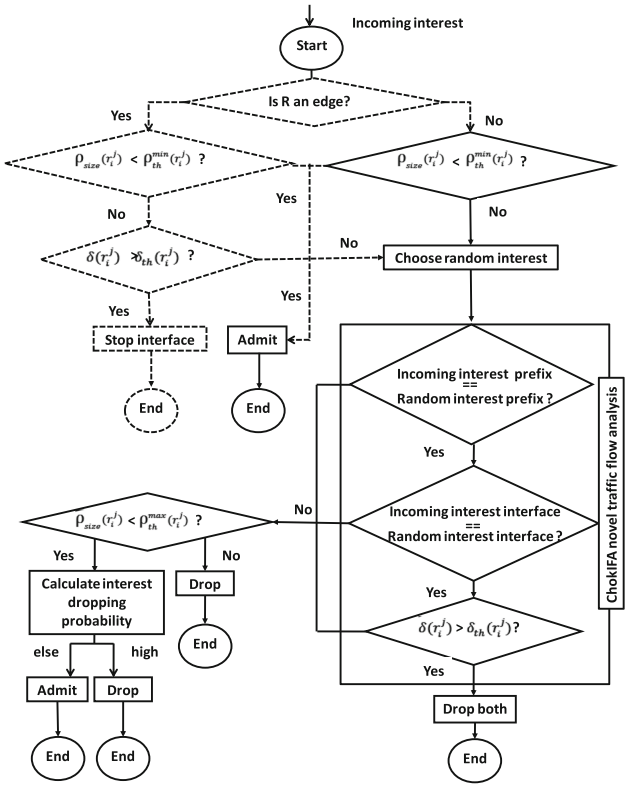

Fig. 3 ChoKIFA+ algorithm flowchart

on a content-oriented communication model [31]. In particular, the traffic flow carries the content name, together with further information used to transport the multiple chunks or segments of the same content. Considering this, we identified the following three attributes to analyze the NDN traffic flow reaching each router: (i) name prefix; (ii) interface; (iii) ISR, i.e., the rate between incoming interests and outgoing content, denoted as $\delta(r_i^j)$ and used to measure routers capability to satisfy interest on a particular interface [2,12]. In particular, if $\delta(r_i^j) > 1$, it means that the number of content packets received by r_i^j is less than the number of interests forwarded from the same interface.

ChoKIFA is a justified stateless algorithm that does not require any specific data synchronization between intermediate routers and traffic analyzers. Compared to the existing rate limiting IFA approaches [3], ChoKIFA performs few operations and it not only restricts the rate of new incoming MIs, but also remove MIs which are already stored in PIT during attack. A detailed flowchart of ChoKIFA is given in Fig. 3, i.e., explicitly presented in solid lines.

To mitigate IFA, ChoKIFA relies on the dynamic computation of the PIT size, denoted as $\rho_{size}(r_i^j)$, and on its evaluation against three thresholds: $\rho_{th}^{\min}(r_i^j)$, $\rho_{th}^{\max}(r_i^j)$ and $\delta_{th}(r_i^j)$ [12]. For each interest arriving at r_i^j, if the PIT size is less than the $\rho_{th}^{\min}(r_i^j)$, the interest gets stored in the router's PIT. During a normal network trend, all interests sent by C are satisfied by P or by a router's cache. Moreover, there are no interests requested for nonexisting contents and there is no massive network traffic going to R. In this scenario, the PIT size almost never reaches $\rho_{th}^{\min}(r_i^j)$. However, when there is either a default delay in the network (e.g., due to congestion or packet loss) or an ongoing IFA, the PIT size can reach $\rho_{th}^{\min}(r_i^j)$. In case of network congestion, the content retrieval is delayed and the PIT entries are kept longer in the PIT, while, in case of IFA, the massive amount of MIs sent by the adversary to saturate the PIT. Thus, the choice of the $\rho_{th}^{\min}(r_i^j)$ threshold is critical, since it has to consider both the legitimate busty traffic and the default delay (more details in Sect. 5.5.1). When the PIT size is greater than $\rho_{th}^{\min}(r_i^j)$ and lower than $\rho_{th}^{\max}(r_i^j)$ (i.e., $\rho_{th}^{\min}(r_i^j) < \rho_{size}(r_i^j) < \rho_{th}^{\max}(r_i^j)$), each new incoming interest is compared with an interest that is randomly selected from the PIT and named as *drop interest candidate*, to verify whether they belong to the same traffic flow, which means: (i) having the same prefix; (ii) coming from the same interface; (iii) checking if the current $\delta(r_i^j)$ exceeds $\delta_{th}(r_i^j)$, in case both previous conditions are verified. If both interests have the same traffic flow, then both are dropped, since most of the PIT entries are likely to be occupied by MIs under IFA. Otherwise, the randomly selected interest is kept stored in the PIT, and the incoming interest is dropped with the probability (P_b), which depends on the average PIT size ($\rho_{avg}(r_i^j)$), as illustrated in Eq. 1 [17].

$$P_b = \frac{P_{\max} * (\rho_{avg}(r_i^j) - \rho_{th}^{\max}(r_i^j))}{(\rho_{th}^{\max}(r_i^j) - \rho_{th}^{\min}(r_i^j))}, \quad (1)$$

$P_{\max}$ denotes the maximum probability.[4] Since the average PIT size varies between $\rho_{th}^{\min}(r_i^j)$ and $\rho_{th}^{\max}(r_i^j)$, the interest dropping probability P_b varies between 0 and $P_{\max}$. This means that an interest is dropped with a probability equal to 1, if it arrives when the average PIT size exceeds $\rho_{th}^{\max}(r_i^j)$, otherwise it is accepted and stored in PIT. In particular, the interest dropping probability is computed by exploiting the mechanism of packet dropping probability of RED [17]. Algorithm 1 illustrates the functionality of the proposed approach at the core routers.

Algorithm 1 Processing incoming packet to core router.

```
input : packet,PITsize,Minth,Maxth
begin
    if packet = Content then
        process packet as Content
    end
    if packet = Interest then
        if (PITsize < Minth) then
            Accepte(Interest)
        else
            Select the old interest (RI) from PIT
            if (Préfixe(Interest) = Préfixe(RI)) & (Interface(Interest)
            = Interface(RI)) & Ratio(Interface(Interest)) ≥ th
            then
                Drop(interest,RI)
            else
                if (PITsize > Maxth) then
                    Drop(Interest)
                else
                    Calculate drooping probability (RED) DPRED
                    if DPRED then
                        Accepte(Interest)
                    else
                        Drop(Interest)
                    end
                end
            end
        end
    end
end
```

5.4 Edge router functionality

The enhanced version of ChoKIFA (i.e., ChoKIFA+) provides an additional security wall on the edges of the network to further improve the network health (i.e., the PIT occupancy) during IFA. In ChoKIFA+, the edge routers detect the malicious behavior of the users close to the adversary, and therefore, they reduce the amount of malicious traffic towards the core routers. In particular, the edge routers detect the adversary location and interface so that they are able to stop the malicious traffic coming from it. The flowchart in Fig. 3 illustrates the functionality of edge routers (i.e., in dotted lines). In particular, an edge router monitors the statistics about the ISR for each consumer, considering the threshold $\delta_{th}(r_i^j)$ to discriminate between legitimate and malicious traffic. For each interest arriving at the edge router, if the current PIT size $\rho_{size}(r_i^j)$ is greater than $\rho_{th}^{\min}(r_i^j)$, the edge router compares the current ISR for that interface (i.e., $\delta_{size}(r_i^j)$) with the threshold $\delta_{th}(r_i^j)$. If $\delta_{size}(r_i^j)$ is greater than $\delta_{th}(r_i^j)$, the consumer is considered malicious and the interface is blocked until the $\delta_{size}(r_i^j)$ ranges under threshold. Algorithm 2 illustrates the procedure followed by the edge router when receiving an interest.

Algorithm 2 Processing incoming packet to edge router.

```
input : packet,PITsize,Minth,Maxth
begin
    if packet = Content then
        process packet as Content
    end
    if packet = Interest then
        if (PITsize < Minth) then
            Accepte(Interest)
        end
        else
            if Ratio(Interface(Interest)) ≥ th then
                stop interface
            end
            else
                process as ChoKIFA
            end
        end
    end
end
```

5.5 Parameters setting

In ChoKIFA, the values of the parameters $\rho_{avg}(r_i^j)$, $\rho_{th}^{\min}(r_i^j)$ and $\rho_{th}^{\max}(r_i^j)$ are essential to make decisions about the desired average PIT size, which directly impacts the interest dropping probability. A proper choice of these parameters ensures the handling of small bursts of benign traffic, which might happen in case of network congestion or content retrieval delay. We will now illustrate the rules applied to choose the best parameters values for both keeping an effective performance of the network and a mitigation of the attack.

5.5.1 Average PIT size calculation

To calculate the average PIT size $\rho_{avg}(r_i^j)$, ChoKIFA uses EWMA, that guarantees that a short-term increase in PIT

[4] We take the value of maximum probability ($P_{\max}$) to be one.

size, due to a burst of BIs, does not result in the significant increase in the average PIT size. Eq. 2 illustrates the calculation of $\rho_{avg}(r_i^j)$, where w_ρ is the weight factor for calculating EWMA and $\rho_{size}(r_i^j)$ is the current/actual PIT size. Then, the average PIT size, used for the interest dropping probability, is updated as follows [17],

$$\rho_{avg}(r_i^j) = (1 - w_\rho) * \rho_{avg}(r_i^j) + w_\rho * \rho_{size}(r_i^j). \quad (2)$$

Note that the calculation of the average PIT size can be made particularly efficient when w_ρ is set to a negative power of two. If w_ρ is too large, then the averaging procedure will not filter out the temporary congestion of PIT [17].

5.5.2 PIT size minimum threshold

The optimal value for $\rho_{th}^{\min}(r_i^j)$ depends on the desired level of the average PIT size and on the default network conditions. In case the default traffic is often congested, the $\rho_{th}^{\min}(r_i^j)$ threshold should be large enough to allow the PIT usage under acceptable levels.

5.5.3 PIT size maximum threshold

The $\rho_{th}^{\max}(r_i^j)$ threshold affects the maximum number of interests that can be stored in the PIT and its value is critical due to the consequences it might generate. If $\rho_{th}^{\max}(r_i^j)$ is too low, the attack might be restricted or ineffective, since the MIs will not be stored in the PIT. However, a low $\rho_{th}^{\max}(r_i^j)$ threshold may also damage the ChoKIFA+ performance detection algorithm, making it report a large number of false positives. This happens, because the ChoKIFA+ detection algorithm requires MIs to be stored in the PIT and to allow the comparison between the incoming malicious traffic with the one already stored in the PIT. On the other side, if $\rho_{th}^{\max}(r_i^j)$ is too high, the adversary could store even a larger amount of MIs and, consequently, delay the detection of the attack and the sensitivity of the algorithm.

5.5.4 Setting $\rho_{th}^{\max}(r_i^j)$ and $\rho_{th}^{\min}(r_i^j)$ to avoid global synchronization

The optimal value for $\rho_{th}^{\max}(r_i^j)$ depends also from the maximum average delay that can be allowed to interest (e.g., round trip time for interest to retrieve data) and from the PIT total size. A useful rule used by ChoKIFA is to set the value of $\rho_{th}^{\max}(r_i^j)$ more than three times the value of $\rho_{th}^{\min}(r_i^j)$ [17], since the mitigation mechanism works efficiently when $\rho_{th}^{\max}(r_i^j)$-$\rho_{th}^{\min}(r_i^j)$ is larger than the typical increase in average PIT size.

Table 3 Parameters for simulation

Parameters	Value
Interest sending rate for C (interests/s)	30
Interest sending rate for Adv (interests/s)	1000
Interest size (kB)	1
Number of C	8
Number of P	1
Number of malicious nodes	4
Number of benign nodes	4
Number of routers	9
Link capacity (Mbps)	10
Link delay (ms)	10
Interest life time (s)	1
R Total PIT size (kB)	600
Min. threshold for PIT size (kB)	1/8 of PIT (75)
Max. threshold for PIT size (kB)	3/4 of PIT (450)
Weight factor (w_ρ) for EWMA	0.001
$P_{\max}$	1
Interest satisfaction ratio threshold	3
Simulation time (s)	100
Simulator version	ndnSIM 2.1
Operating system	Ubuntu 16.04

6 Evaluation

In this section, we evaluate the effectiveness and efficiency of our proposed approach under IFA, by comparing ChoKIFA with ChoKIFA+ and with the state of the art IFA mitigation approaches [3]. In particular, among those we chose to compare with the ones that implement an interest rate limiting based on interface fairness, ISR and limit announcement technique. To this end, we implemented the protocol and performed extensive simulations using the open-source ndnSIM [1] simulator.[5]

6.1 Simulations setup

We ran the simulations on two different network topologies, both for 100 s: a tree topology [13] and a more realistic large-scale ISP-like topology (i.e., AS-7018 [36]). We chose the tree topology, because it provides one of the worsts scenarios to apply a defense against IFA [3], while we chose the larger ISP topology to evaluate the performance of the mitigation approach, when deployed on the real Internet. Table 3 illustrates the other network parameters we used for the simulation setup.

[5] ndnSIM implements the NDN protocol stack on NS-3 simulator.

6.2 Evaluation metrics

To evaluate the impact of IFA and to compare our solution with the state of the art, we adopted the following metrics, which have been widely used in related work [12,35,39]:

- the PIT usage: this indicates the available capacity of the routers to process benign traffic during an attack;
- the percentage of BIs and MIs dropped by the network during IFA: this measures the attack impact and the effectiveness of the countermeasure;
- the ISR of benign users: this is intended to measure the benign traffic received by users during IFA. The lower the ISR, the greater the amount of false positives generated by the mitigation approach when distinguishing between MIs and BIs;
- any additional delay encountered by the clients, when the proposed mitigation approach is active:

6.3 Evaluation on a small-scale topology

As a small-scale topology [3,13], we used the one shown in Fig. 4, that has multiple benign consumers (i.e., C) retrieving the desired content from a producer (i.e., P), that publishes contents under a specific name prefix (i.e., $prefix$). We assume that C sends a BI for a content specified as $prefix/data$, which can be satisfied by P after traversing multiple routers R. On the contrary, the adversary sends requests for non-existing content (i.e., MI), which exhibits a distinct suffix ($/good/rnd$) compared to the one of the existing contents ($/good/data$). Each router has the NDN default features [46] and it is referred as $r_i^j \epsilon |R|$, where j is the interface of i-th router. In addition, each router performs caching and uses *best route* as forwarding strategy.

In this section, we used the tree topology to evaluate: (i) the impact of the three IFA types on the routers PIT occupancy and on the ISR of benign consumers; (ii) the effectiveness of ChoKIFA and ChoKIFA+ in terms of routers PIT occupancy and ISR of benign consumers; (iii) the effectiveness of ChoKIFA+ against the three IFA types in terms of routers PIT occupancy and ISR of benign consumers.

6.3.1 Impact of interest flooding attacks

Figure 5 shows the impact of the three IFA types on the routers PIT occupancy with no active countermeasure. On the y-axis, we plot the average PIT usage (i.e., 600 kB is the maximum PIT size) and the variation of its usage under IFA. At 20 s, the adversary starts the attack by issuing MIs with a rate equal to 1000 interests/s. As shown in Fig. 5, the impact of the third IFA type is more significant than the impact of the other two IFA types. This is motivated by the in-network caching, that provides an intrinsic defense against the first

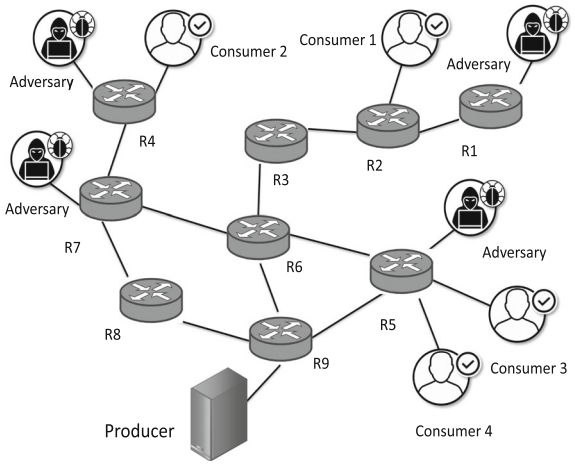

Fig. 4 Small-scale topology used for simulations

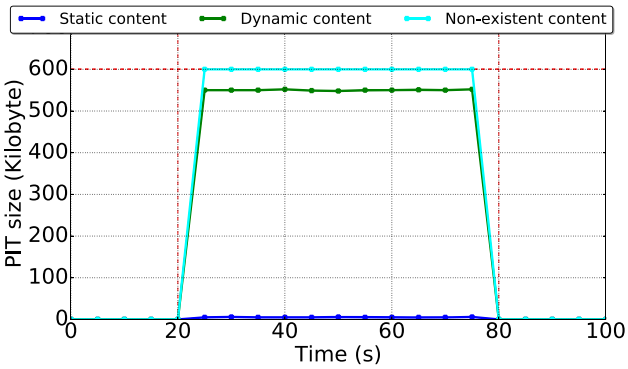

Fig. 5 Impact of the three IFA types on the PIT usage in a small-scale topology

and the second IFA type, while in the third IFA type, the interests are kept stored in the PIT until their lifetime has passed. Before the 20th second, only legitimate consumers send BIs with an interest sending rate equal to 30 interests/s. Since those interests refer to existing contents, they are all satisfied and the PIT average occupancy is equal to zero.

Figure 6 shows the impact of all IFAs types on the consumers ISR with no active countermeasure. This quantifies the quality of service observed by legitimate users, while the network is under IFA. Under the first IFA type, consumers have an ISR between 20 and 30% higher than the ISR under the second IFA type. This result was expected, since requests for static contents are satisfied by all on-path-router's caches. On the contrary, the requests for dynamic contents are routed to the producer. Finally, the third IFA type generates the worst ISR, since the network drops almost 90% of the legitimate traffic. For the rest of simulations, we focus on the IFA where adversary generates unsatisfiable interests, i.e., the third IFA type.

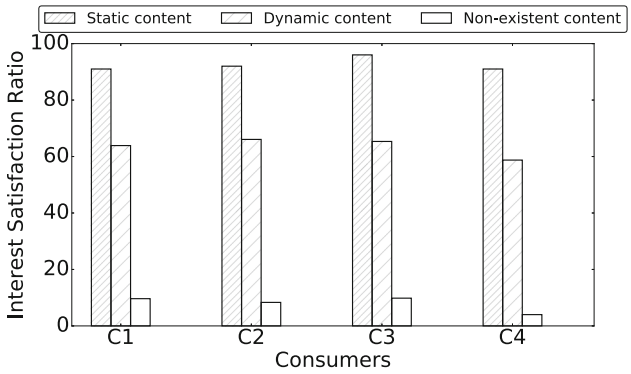

Fig. 6 Impact of the three IFA types on consumers ISR in a small-scale topology

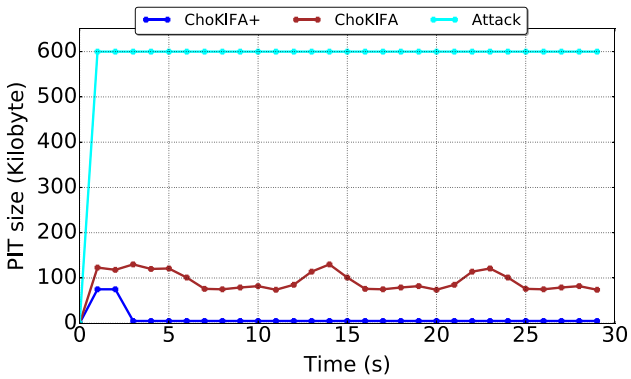

Fig. 7 Effectiveness of ChoKIFA and ChoKIFA+ on the PIT usage in a small-scale topology under the third IFA type

6.3.2 Comparison between ChoKIFA and ChoKIFA+ effectiveness

Figure 7 compares the effectiveness of ChoKIFA and ChoKIFA+ against the third IFA type in terms of PIT average usage. In the simulations, the adversaries start the attack at the same time as the benign users which send requests for existing content from the beginning of the simulation. Following the ChoKIFA algorithm, the PIT size needs to be filled with a certain amount of MIs before the drop of the malicious traffic starts. Thus, in both ChoKIFA and ChoKIFA+, the PIT size reaches a greater value than the minimum threshold, after which our approach starts treating each new incoming interest separately, identifying first its traffic flow (i.e., malicious or benign), and then deciding whether to drop the interest or not. However, with respect to ChoKIFA, ChoKIFA+ keeps the average PIT usage significantly lower because of the security wall introduced at the edge routers. Under IFA, the edge routers readily detect the attack by monitoring the flow of consumers over each interface.

Figure 8 compares the performance of ChoKIFA and ChoKIFA+ under IFA in terms of ISR observed by consumers. The legitimate traffic is slightly affected by the attack as, on average, only the 4% of BIs are dropped.

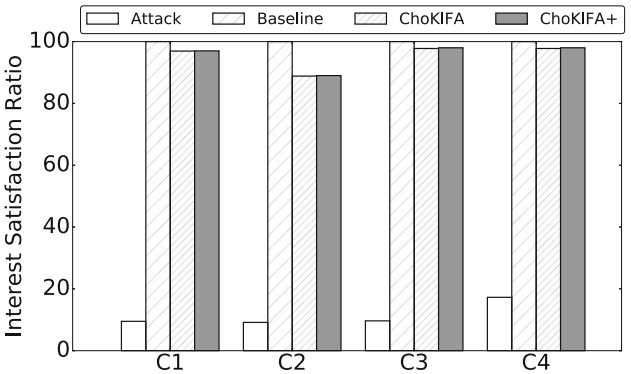

Fig. 8 Effectiveness of ChoKIFA and ChoKIFA+ on consumers ISR in a small-scale topology under the third IFA type

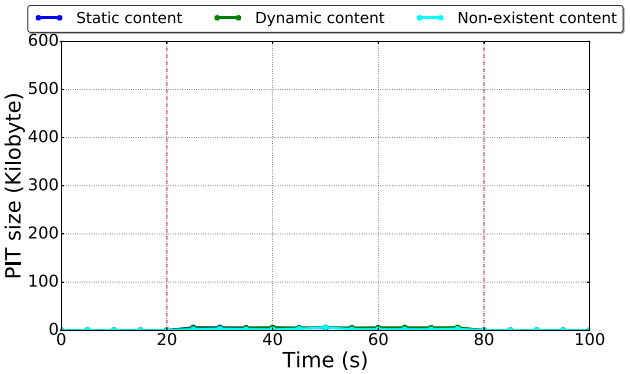

Fig. 9 Effectiveness of ChoKIFA+ ont he PIT usage in a small-scale topology under the three IFA types

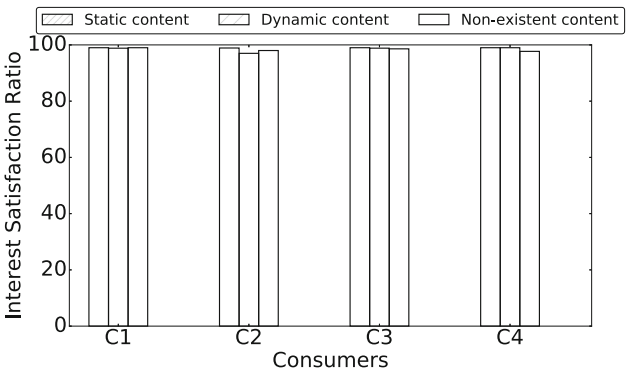

Fig. 10 Effectiveness of ChoKIFA+ on consumers ISR in a small-scale topology under the three IFA types

Finally, Figs. 9 and 10 illustrate the effectiveness of ChoKIFA+ in terms of PIT usage and ISR encountered by the consumers when the three IFA types are active. Thanks to the additional security wall introduced at the edge routers, ChoKIFA+ can readily recognize an attack and guarantee good network performances: an average PIT usage close to zero and a consumer ISR around 98%.

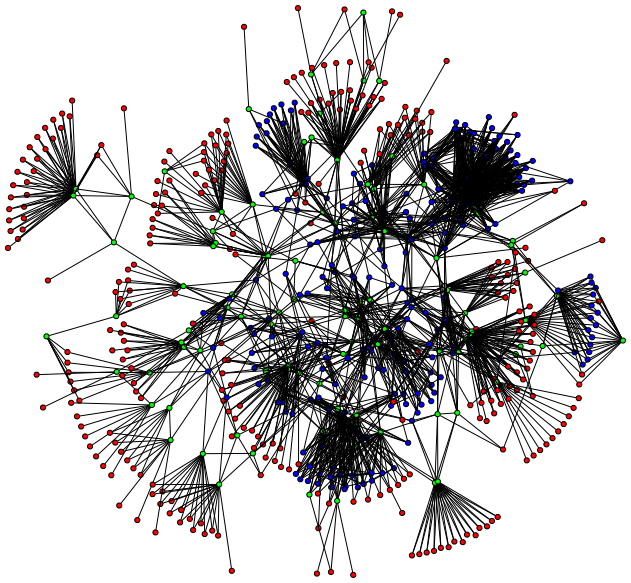

Fig. 11 Large scale topology used for simulations

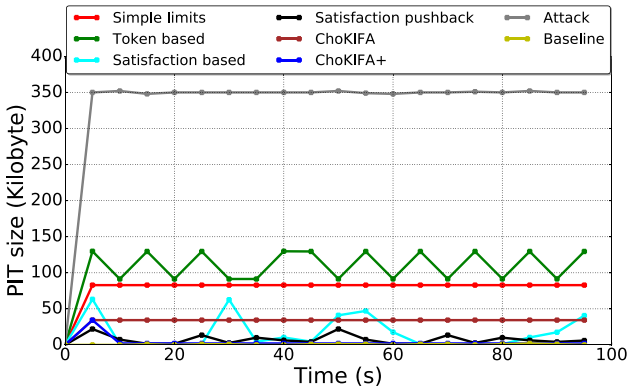

Fig. 12 Effectiveness of ChoKIFA, ChoKIFA+ and the existing solutions on the PIT usage in a large scale topology under the third IFA type

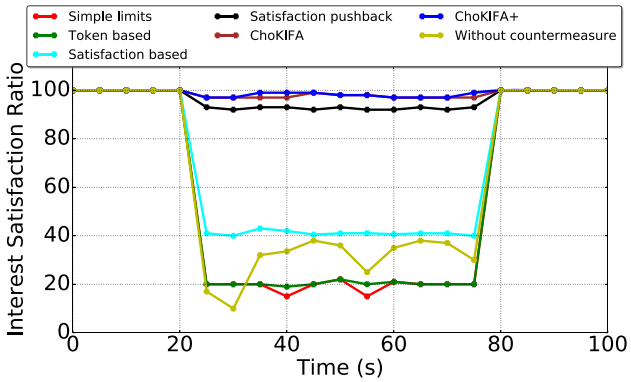

Fig. 13 Effectiveness of ChoKIFA, ChoKIFA+ and the existing solutions on consumers ISR in a large scale topology under the third IFA type

6.4 Evaluation on a large scale topology

As a large scale topology, we considered the AS 7018 topology, measured by the Rocket fuel project [36] and shown in Fig. 11. The topology involves 625 nodes, separated into three categories: clients, gateways, and backbones. 296 nodes are classified as clients due to a degree less than four, while the 108 nodes connected to clients are classified as gateways. The remaining 221 nodes are classified as backbones. A large ISP topology reflects how our mitigation methods would perform when deployed on the real Internet. To study the performance of our proposed mitigation strategy under a range of conditions, we varied the percentage of adversaries in the network and the frequency with which they send MIs.

6.4.1 ChoKIFA+ effectiveness and comparison with the state-of-the-art

Figure 12 shows how ChoKIFA, ChoKIFA+ and four different IFA mitigation approaches perform with respect to the average PIT usage under the third IFA type. As discussed in Sect. 4, the four approaches, proposed in [3] and available online,[6] are: (i) *simple limits*, (ii) *token based*, (iii) *satisfaction based* and (iv) *satisfaction pushback*.

In our simulation, the adversaries start launching the attack at the same time as the benign users start sending requests, i.e., from the beginning. As shown in the figure, ChoKIFA attains slightly higher PIT size than the *satisfaction pushback*. This is because all routers allow the PIT to be filled till the minimum threshold is reached. On the other side, ChoKIFA+ detects and blocks the malicious traffic at the edge routers, and stops the MI from being stored in the PIT.

Figure 13 illustrates the performance of all the approaches in terms of consumers ISR under an attack that starts at second 20 and ends at second 80 of the simulation. As shown in the figure, the rate limiting approaches [3] are not able to maintain an acceptable ISR for benign users in a large topology. In particular, only the *satisfaction pushback* provides a reasonable ISR, but it is overcome by both ChoKIFA and ChoKIFA+.

Figure 14 shows the ISR percentage for legitimate interests generated in the network, under a varying number of adversaries in the network (i.e., from 6% of adversaries to over 50% of adversaries in the network). For any mitigation technique, the ISR ratio for legitimate interests decreases (i.e., aggregated for all benign users) together with the increase in the number of adversaries. However, while the *token based* algorithm faces an ISR worsening already for 3 adversaries and 13 legitimate users, the other approaches have worse performances only for a higher number of adversaries. In particular, ChoKIFA+ outperforms all mitigation

[6] https://github.com/cawka/ndnSIM-ddos-interest-flooding.

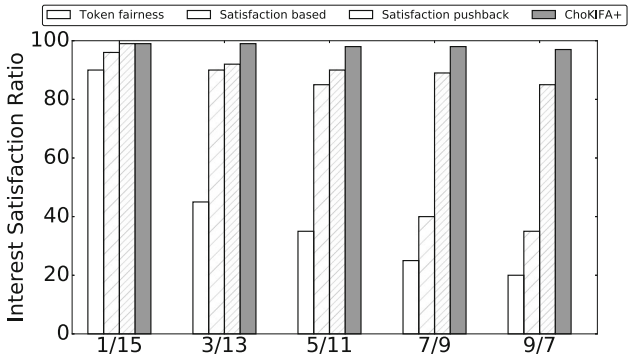

Fig. 14 Consumers ISR with different mitigation approaches and an increasing number of adversaries in a large scale topology

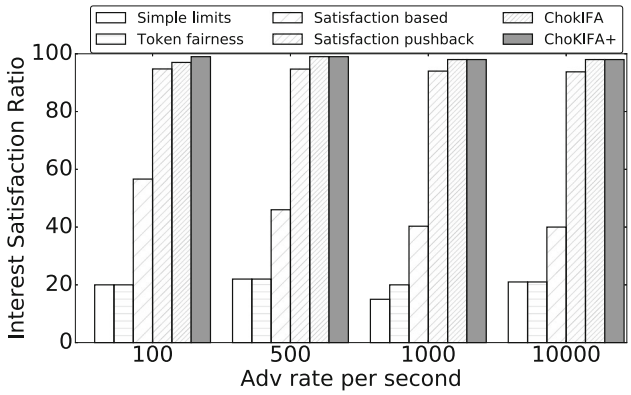

Fig. 15 Consumers ISR with different mitigation approaches and an increasing MI sending rate in a large scale topology

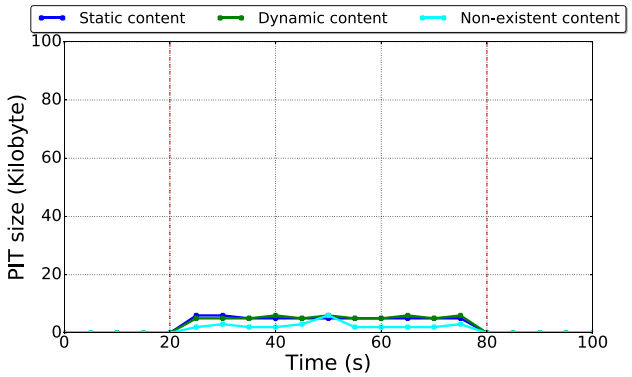

Fig. 16 Effectiveness of ChoKIFA+ on the PIT usage under the three IFA types

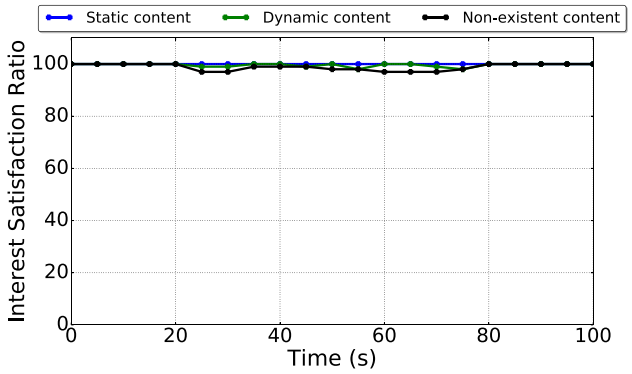

Fig. 17 Effectiveness of ChoKIFA+ on consumers ISR under the three IFA types

algorithms and shows a very minor reduction in ISR ratio (i.e., approximately 3%) even when the adversary percentage is raised more than 50%.

Figure 15 shows the aggregated legitimate ISR ratio under a varying adversary interest sending rate (i.e., from 100 to 10,000 interests/s). As shown in the figure, both ChoKIFA and ChoKIFA+ are almost unaffected, even with a high increase in the adversary interest sending rate, while, among the state-of-the-art approaches, only the *satisfaction pushback* shows satisfactory results.

In Figs. 16 and 17, we show the impact of the ChoKIFA+ mitigation strategy under the three IFA types in a large topology. Considering Fig. 16, ChoKIFA+ proves to be able to keep roughly the 90% of the PIT size available for handling new traffic. Similarly, ChoKIFA+ is able to maintain almost the 97% of the global ISR.

Finally, we also analyzed the performance of ChoKIFA+ in terms of delay encountered between the interest sending and the corresponding data receiving. In particular, Fig. 18 shows the additional delay introduced for each consumer by ChoKIFA+ with respect to the default delay measured in the network, without any attack. As shown in the figure, ChoKIFA+ introduces an extra delay equal to 10.9 µs on average.

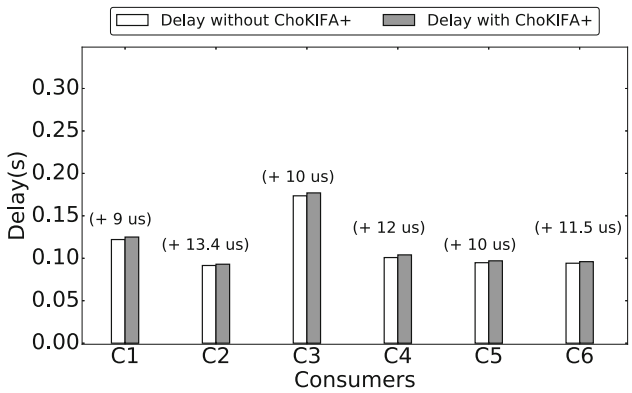

Fig. 18 Additional delay time in μs introduced by ChoKIFA+

7 Conclusion

In this paper, we address the interest flooding-based DDoS over NDN, which is explicitly named as IFA. More specifically, we have found that several proposed countermeasures, that adopt detection and reaction mechanisms based on interest rate limiting, are not highly effective and also damage the legitimate traffic.

In our solution, we exploited an AQM scheme to propose an efficient detection and mitigation mechanism against IFA, which stabilizes the router PIT. The proposed approach penalizes the unresponsive flows generated by adversarial traffic by dropping MIs generated during the IFA. To evaluate the effectiveness of our solution, we implemented the proposed protocol on the open-source ndnSIM simulator and compared it with the state of the art. The results report that our proposed protocol effectively mitigates the adverse effects of IFA and shows significantly less false positives in comparison to the state-of-the-art IFA mitigation approaches.

Funding This work was supported by the European Commission under the Horizon 2020 Programme (H2020), as part of the LOCARD project (Grant Agreement No. 832735).

Compliance with ethical standards

Conflict of interest Abdelmadjid Benarfa declares that he has no conflict of interest. Muhammad Hassan declares that he has no conflict of interest. Eleonora Losiouk declares that she has no conflict of interest. Alberto Compagno declares that he has no conflict of interest. Mohamed bachir Yagoubi declares that he has no conflict of interest. Mauro Conti declares that he has no conflict of interest.

Ethical approval This article does not contain any study with human participants or animals performed by any of the authors.

References

1. Afanasyev, A., Moiseenko, I., Zhang, L.: ndnSIM: NDN simulator for NS-3. Technical Report NDN-0005, NDN. http://named-data.net/techreports.html (2012). Accessed Apr 2018
2. Afanasyev, A., Moiseenko, I., Zhang, L., et al.: ndnsim: Ndn simulator for ns-3. University of California, Los Angeles, Technical Report 4 (2012)
3. Afanasyev, A., Mahadevan, P., Moiseenko, I., Uzun, E., Zhang, L.: Interest flooding attack and countermeasures in named data networking. In: Ifip Networking Conference, pp. 1–9. IEEE (2013)
4. Ahlgren, B., Dannewitz, C., Imbrenda, C., Kutscher, D., Ohlman, B.: A survey of information-centric networking. IEEE Commun. Mag. **50**(7), 26–36 (2012)
5. Bedi, H., Roy, S., Shiva, S.: Mitigating congestion-based denial of service attacks with active queue management. In: IEEE Global Communications Conference (Globecom), pp. 1440–1445. IEEE (2013)
6. Bedi, H., Sankardas, R., Sajjan, S.: Mitigating congestion based dos attacks with an enhanced aqm technique. Comput. Commun. **56**, 60–73 (2015). https://doi.org/10.1016/j.comcom.2014.09.002
7. Benarfa, A., Hassan, M., Compagno, A., Losiouk, E., Yagoubi, M.B., Conti, M.: Chokifa: A new detection and mitigation approach against interest flooding attacks in ndn. In: International Conference on Wired/Wireless Internet Communication, pp. 53–65. Springer (2019)
8. Benmoussa, Ahmed, Tahari, A.K., Lagaa, N., Lakas, A., Ahmad, F., Hussain, R., Kerrache, C.A., Kurugollu, F.: A novel congestion-aware interest flooding attacks detection mechanism in named data networking. In: 28th International Conference on Computer Communication and Networks (ICCCN), pp. 1–6. IEEE (2019)
9. Brownlee, N., Mills, C., Ruth, G.: Traffic flow measurement: architecture (1997)
10. Chhabra, P., Chuig, S., Goel, A., John, A., Kumar, A., Saran, H., Shorey, R.: Xchoke: malicious source control for congestion avoidance at internet gateways. In: Proceedings. 10th IEEE International Conference on Network Protocols, 2002, pp. 186–187. IEEE (2002)
11. Compagno, A., Conti, M., Hassan, M.: An icn-based authentication protocol for a simplified lte architecture. In: Baldi, M., Quaglia, E.A., Tomasin, S. (eds.). Cham: Springer (2018)
12. Compagno, A., Conti, M., Gasti, P., Tsudik, G.: Poseidon: mitigating interest flooding ddos attacks in ndn. In: IEEE 38th Conference on Local Computer Networks (lCN), pp. 630–638. IEEE (2013)
13. Dai, H., Wang, Y., Fan, J., Liu, B.: Mitigate ddos attacks in ndn by interest traceback. In: IEEE Conference on Computer Communications Workshops (Infocom Workshops), pp. 381–386. IEEE (2013)
14. Dong, J., Wang, K., Lyu, Y., Jiao, L., Yin, H.: Interestfence: countering interest flooding attacks by using hash-based security labels. In: International Conference on Algorithms and Architectures for Parallel Processing, pp. 527–537. Springer (2018)
15. Feng, W., Kandlur, D.D., Saha, D., Shin, K.G.: Stochastic fair blue: a queue management algorithm for enforcing fairness. In: Infocom 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE, vol. 3, pp. 1520–1529. IEEE (2001)
16. Feng, W., Shin, K.G., Kandlur, D.D., Saha, D.: The blue active queue management algorithms. IEEE/ACM Trans. Netw. **10**(4), 513–528 (2002)
17. Floyd, S., Jacobson, V.: Random early detection gateways for congestion avoidance. IEEE/ACM Trans. Netw. **1**(4), 397–413 (1993)
18. Gasti, P., Tsudik, G., Uzun, E., Zhang, L.: Dos and ddos in named data networking. In: 2013 22nd International Conference on Computer Communication and Networks (ICCCN), pp. 1–7. https://doi.org/10.1109/ICCCN.2013.6614127 (2013a)
19. Gasti, P., Tsudik, G., Uzun, E., Zhang, L.: Dos and ddos in named data networking. In: 22nd International Conference on Computer Communications and Networks (ICCCN), pp. 1–7. IEEE (2013)
20. Govindaswamy, V.V., Záruba, G., Balasekaran, G: Rechoke: a scheme for detection, control and punishment of malicious flows in ip networks. In: Global Telecommunications Conference, 2007. Globecom'07, pp. 16–21. IEEE (2007)
21. Hou, R., Han, M., Chen, J., Wenbin, H., Tan, X., Luo, J., Ma, M.: Theil-based countermeasure against interest flooding attacks for named data networks. IEEE Netw. **33**(3), 116–121 (2019)
22. Jacobson, V., et al.: Networking named content. In: ACM International Conference on Emerging Networking Experiments and Technologies, pp. 1–12 (2009)
23. Jiang, X., Yang, J., Jin, G., Wei, W.: Red-ft: a scalable random early detection scheme with flow trust against dos attacks. IEEE Commun. Lett. **17**(5), 1032–1035 (2013). https://doi.org/10.1109/LCOMM.2013.022713.122652
24. Kidambi, J., Ghosal, D., Mukherjee, B.: Dynamic token bucket (dtb): a fair bandwidth allocation algorithm for high-speed networks. J. High Speed Netw. **9**(2), 67–87 (2000)
25. Kunniyur, S.S., Srikant, R.: An adaptive virtual queue (avq) algorithm for active queue management. IEEE/ACM Trans. Netw. **12**(2), 286–299 (2004)
26. Lin, D., Morris, R.: Dynamics of random early detection. In: ACM Sigcomm Computer Communication Review, vol. 27, pp. 127–137. ACM (1997)
27. Liu, G., Quan, W., Cheng, N., Wang, K., Zhang, H.: Accuracy or delay? A game in detecting interest flooding attacks. Internet Technol. Lett. **1**(2), 31 (2018)
28. Nguyen, T., Cogranne, R., Doyen, G.: An optimal statistical test for robust detection against interest flooding attacks in ccn. In:

Ifip/IEEE International Symposium on Integrated Network Management (IM), pp. 252–260. IEEE (2015)
29. Nguyen, T., Mai, H.-L., Doyen, G., Cogranne, R., Mallouli, W., Montes, E., de Oca, O.: Festor: a security monitoring plane for named data networking deployment. IEEE Commun. Mag. **56**(11), 88–94 (2018)
30. Nguyen, T., Mai, H.-L., Cogranne, R., Doyen, G., Mallouli, W., Nguyen, L., El Aoun, M., Oca, E.M.D., Festor, O.: Reliable detection of interest flooding attack in real deployment of named data networking. IEEE Trans. Inf. Forensics Secur. **14**(9), 2470–2485 (2019)
31. Oueslati, S., Roberts, J., Sbihi, N.: Flow-aware traffic control for a content-centric network. In: 2012 Proceedings IEEE Infocom, pp. 2417–2425. https://doi.org/10.1109/INFCOM.2012.6195631 (2012)
32. Pan, J., Paul, S., Jain, R.: A survey of the research on future internet architectures. IEEE Commun. Mag. **49**(7), 26–36 (2011)
33. Pan, R., Prabhakar, B., Psounis, K.: Choke-a stateless active queue management scheme for approximating fair bandwidth allocation. In: Infocom 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEEE, vol. 2, pp. 942–951. IEEE (2000)
34. Rai, S., Sharma, K., Dhakal, D.: A survey on detection and mitigation of distributed denial-of-service attack in named data networking. In: Advances in Communication, Cloud, and Big Data, pp. 163–171. Springer (2019)
35. Salah, H., Wulfheide, J., Strufe, T.: Coordination supports security: a new defence mechanism against interest flooding in ndn. In: 2015 IEEE 40th Conference on Local Computer Networks (ICN), pp. 73–81. https://doi.org/10.1109/LCN.2015.7366285 (2015)
36. Spring, N., et al.: Measuring ISP topologies with rocketfuel. IEEE/ACM Trans. Netw. **12**, 2–16 (2004)
37. Tourani, R., Misra, S., Mick, T., Panwar, G.: Security, privacy, and access control in information-centric networking: a survey. IEEE Commun. Surv. Tutor. **20**(1), 566–600 (2017)
38. Vassilakis, V.G., Alohali, B.A., Moscholios, I.D., Logothetis, M.D.: Mitigating distributed denial-of-service attacks in named data networking. In: Proceedings of the 11th Advanced International Conference on Telecommunications (AICT), Brussels, Belgium, pp. 18–23 (2015)
39. Wang, K., Zhou, H., Qin, Y., Chen, J., Zhang, H.: Decoupling malicious interests from pending interest table to mitigate interest flooding attacks. In: Globecom Workshops (gc wkshps), 2013 IEEE, pp. 963–968. IEEE (2013)
40. Wang, K., Zhou, H., Luo, H., Guan, J., Qin, Y., Zhang, H.: Detecting and mitigating interest flooding attacks in content-centric network. Secur. Commun. Netw. **7**(4), 685–699 (2014)
41. Xylomenos, G., Ververidis, C.N., Siris, V.A., Fotiou, N., Tsilopoulos, C., Vasilakos, X., Katsaros, K.V., Polyzos, G.C.: A survey of information-centric networking research. IEEE Commun. Surv. Tutor. **16**(2), 1024–1049 (2014). https://doi.org/10.1109/SURV.2013.070813.00063
42. Zhang, C., Yin, J., Cai, Z., Chen, W.: Rred: robust red algorithm to counter low-rate denial-of-service attacks. IEEE Commun. Lett. **14**(5), 489–491 (2010)
43. Zhang, G., Li, Y., Lin, T.: Caching in information centric networking: a survey. Comput. Netw. **57**(16), 3128–3141 (2013). https://doi.org/10.1016/j.comnet.2013.07.007
44. Zhang, L., et al.: Named data networking. ACM SIGCOMM CCR **44**(3), 66–73 (2014)
45. Zhang, L., Estrin, D., Burke, J., Jacobson, V., Thornton, J.D., Smetters, D.K., Zhang, B., Tsudik, G., Massey, D., Papadopoulos, C., et al.: Named data networking (ndn) project. Relatório Técnico NDN-0001, Xerox Palo Alto Research Center-PARC 157: 158 (2010)
46. Zhang, L., Afanasyev, A., Burke, J., Jacobson, V., Crowley, P., Papadopoulos, C., Wang, L., Zhang, B., et al.: Named data networking. ACM SIGCOMM Computer Communication Review **44**(3), 66–73 (2014)
47. Zhang, X., Li, R.: A charging, rewarding mechanism-based interest flooding attack mitigation strategy in ndn. In: Ifip/IEEE Symposium on Integrated Network and Service Management (IM), pp. 402–407. IEEE (2019)
48. Zhang, Z., Yu, Y., Zhang, H., Newberry, E., Mastorakis, S., Li, Y., Afanasyev, A., Zhang, L.: Revision 2, April 8, An Overview of Security Support in Named Data Networking (2018)