Reliable Detection of Interest Flooding Attack in Real Deployment of Named Data Networking

Tan Nguyen, Hoang-Long Mai, Rémi Cogranne, *Member, IEEE*, Guillaume Doyen, Wissam Mallouli, Luong Nguyen, Moustapha El Aoun, Edgardo Montes de Oca and Olivier Festor.

Copyright ©2018 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org

Accepted version, final version available online on ieeexplore.ieee.org. DOI: TBA

Abstract—Named Data Networking (NDN) is a disruptive yet promising architecture for the future Internet, in which the content diffusion mechanisms are shifted from the conventional host-centric to content-centric ones so that the data delivery can be significantly improved. After a decade of research and development, NDN and the related NDN Forwarding Daemon (NFD) implementations are now mature enough to enable stakeholders, such as telcos, to consider them for a real deployment. Consequently, NDN and IP will likely cohabit, and the Future Internet may be formed of isolated administrative domains, each deploying one of these two network paradigms. The security question of the resulting architecture naturally arises. In this paper, we consider the case of Denial of Service. Even though the Interest Flooding Attack (IFA) has been largely studied and mitigated through NACK packets in pure NDN networks, we demonstrate in this paper through experimental assessments that there are still some ways to mount such an attack, and especially in the context of coupling NDN with IP, that can hardly be addressed by current solutions. Subsequently, we leverage hypothesis testing theory to develop a Generalized Likelihood Ratio Test (GLRT) adapted to evolved IFA attacks. Simulations show the relevance of the proposed model for guaranteeing the prescribed Probability of False Alarm (PFA) and highlights the trade-off between detection power and delay. Finally, we consider a real deployment scenario where NDN is coupled with IP to carry HTTP traffic. We show that the model of IFA attacks is not very accurate in practice and further develops a sequential detector to keep a high detection accuracy. By considering data from the testbed, we show the efficiency of the overall detection method.

Index Terms—Computer networking attack, Interest flooding attack, DDoS detection, Hypothesis testing, Named Data Networking.

Copyright (c) 2018 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org

Accepted version, final version available online on ieeexplore.ieee.org. DOI: TBA

Tan Nguyen, Rémi Cogranne, Guillaume Doyen and Moustapha El Aoun are with transdisciplinary cyber-security team, ICD, FRE 2018 CNRS, Troyes University of Technology, Troyes, France. Rémi Cogranne and Moustapha El Aoun are also with the Lab. of System Modeling and Dependability and Guillaume Doyen is with the Autonomous Network Environment Team.

Hoang-Long Mai, Wissam Mallouli, Luong Nguyen and Edgardo Montes-De-Oca are with Montimage Research Labs, 39 rue Bobillot, 75013, Paris, France

Olivier Festor is with TELECOM Nancy - University of Lorraine, 54600 Villers-les-Nancy, France

I. INTRODUCTION

OVER the last few years, the amount of traffic on the Internet has kept increasing due to a wider range of connected devices, as well as bandwidth-consuming services and the amount of content they generate¹. Such growth has put high pressure on the underlying infrastructure which has been invented decades ago with a different use of the Internet in mind. Recent research efforts to confront this challenge have resulted in several disruptive network architectures. Among them, Information-Centric Networking (ICN) [1], [2], and particularly Named Data Networking (NDN) [3], [4] has been regarded as the most promising Internet architecture for the future. It is advocated that, by moving from the conventional host-centric content diffusion mechanisms to content-centric ones, the data delivery can be more efficiently optimized. Specifically, in such a network, each content object is given a name which can be addressed at the network level by network elements, instead of the hosts' Internet Protocol (IP) address. To improve the delivery of popular data, contents can be delivered in a multicast manner and from any nodes, thanks to stateful NDN routers with a caching capability.

After a decade of research and development, the NDN architecture together with its forwarding daemon implementation is now widely acknowledged as the most mature ICN proposal, and the most promising solution regarding real deployment considerations. Similarly to IPv6 which has been progressively introduced and which cohabits with IPv4, the most credible scenario for an NDN deployment will most likely consist of dedicated domains, each composed of a single protocol stack, altogether interconnected by dedicated gateways. Although the security of each individual protocol stack has been largely studied, the question of their coupling remains open and challenging. In this paper, we consider the case of Denial of Service which, in such a deployment context, can be inefficient in an IP domain due to the stateless nature of the protocol, while being wasteful in an NDN one due to its stateful nature [5], as assessed by the NDN board².

In this paper, we study the statistical detection of Interest Flooding Attack (IFA) in the realistic context of NDN coupled

¹See: Cisco Visual Networking Index (VNI) 2017.

²See: named-data.net/project/faq/

with IP and altogether carrying Hyper Text Transfer Protocol (HTTP) traffic issued by IP hosts. The problem is cast within the framework of hypothesis testing theory. As opposed to machine-learning based approach, hypothesis testing requires an accurate statistical model of the problem it is aimed at detecting. However, this methodology has indisputable advantages including assessment of the test performance and interesting insights on how the parameters affects the detectability. To the best of our knowledge, this approach has only been used for IFA detection in our prior works [6], [7], [32] that the present paper intends to extend. The present paper proposes to make a step toward the practical application of those prior works by using realistic data, obtained from a real deployment in a testbed environment, where web content is retrieved through an NDN island coupled with existing IP networks. The contribution of the present paper is detailed in Section II-D, the main ones are the following:

1. We argue for the persistence of IFA's threat in NDN by revealing an IFA scenario that succeeds even with the existence of NDN protection mechanisms, and assess the attack scenario by providing the results of experimentations conducted on the last NDN implementation available to date. Note that almost none of the previous works, and especially our prior works, considered such protection mechanism.
2. We design an optimal Likelihood Ratio Test (LRT) of IFA in the theoretical case of a perfectly known legitimate traffic. The optimality of this statistical test is ensured no matter what the attack payload is. This test serves as an upper bound of the expected detection accuracy for IFA.
3. We propose a parametric statistical model upon which is designed a practical Generalized LRT (GLRT) for the scenario where the legitimate traffic is unknown.
4. We extend our detector to increase its accuracy by developing a sequential version based on the initial snapshot GLRT. Using a sequential approach is almost necessary because it is shown, in this paper, that in a real implementation the effect of IFA is much less obvious as what has been assumed from simulations; therefore, gathering evidence from several samples is crucial.
5. We assess the overall performance of our detection solution, using both simulations and real experimentations. Simulation results first allow us to evaluate intrinsic properties of our detection approach. The statistical properties of the proposed GLRT are then evaluated with real data collected from the deployed attack scenario. As far as we know, such a real deployment has never been done before in any previous work. This allows us to compare the theoretical findings with empirical results as well as to compare the performance of our proposal to other detection methods.³

The rest of the paper is organized as follows. Section II presents a background on NDN, on IFA and how previous works address it. Section III summarizes our previous work of investigating IFA in native NDN and introduces the newly investigated use-case of NDN coupled to HTTP, as well as the IFA scenario regardless of the *NACK* existence. Section IV

presents, step by step, our design for the IFA detection using the hypothesis testing theory including a detection problem statement, an optimal detection when traffic parameters are known, a generalized detection when traffic parameters are unknown and finally, a sequential detection. In Section V, the proposed detection method is assessed with simulated data to show the relevance of the theoretical findings, followed by an evaluation with real data, to demonstrate the efficiency and accuracy of the detection method. Finally, Section VI concludes the paper and discusses future work.

II. RELATED WORK

In this section, we provide an NDN background and present the principle of IFA. Then, we survey a set of remarkable proposals which aim at detecting and mitigating this attack. Also, we pay particular attention to an NDN protocol enhancement whose purpose is to solve the IFA intrinsically. Finally, we motivate our proposal with regards to all the previously identified competitors in this area.

A. Named Data Networking

ICN is a networking paradigm which is based on data objects. The key concept in ICN is that it names each data object in the network, instead of using IP addresses for naming hosts and nodes. ICN also deploys in-network caching to enhance the delivery of popular data objects. Besides, a node in ICN does not have to connect to one specific server to get data. Alternately, this node sends a request with the data name and the network will return the corresponding object, either from a cache or from the original provider. Based on these concepts, many ICN architectures have been introduced, including Data-Oriented Network Architecture (DONA) [8], Publish-Subscribe Internet Technology (PURSUIT) [9], Network of Information (NetInf) [10] and NDN [3] which is the most popular and acknowledged in the research community. Beyond its novel architectural principles, NDN has been fully implemented. Tools such as the *ndnSIM* [11] simulator and the NDN Forwarding Daemon (NFD)⁴ have been developed that serve as reference implementations.

In NDN, communications are based on requests for hierarchical content names and are performed by *Interest* and *Data* packets. A user sends an *Interest* to request a content and receives a *Data* in return. An NDN router includes three main components. First is the Forwarding Information Base (FIB) which contains routing information for *Interest*. Secondly, the Content Store (CS) is essentially a local cache which stores recently requested *Data* to improve performances. Finally, the Pending Interest Table (PIT) contains entries for each forwarded *Interest*, and uses them as reverse-path routing information for *Data*. A PIT entry contains an NDN name and multiple incoming interfaces.

Fig. 1a illustrates the *Interest* forwarding. When an NDN router receives an *Interest*, it checks the CS first. If a cached copy exists, the router sends this copy back to the incoming interface. If a cached copy does not exist, but a PIT entry

³The source codes and data used in the paper will be made freely available for reproducible research.

⁴See: named-data.net/doc/NFD/current.

Fig. 1. NDN forwarding process for (a) *Interest* and (b) *Data* packets

for this content name has already been created, the *Interest*'s incoming interface is added to this entry and the *Interest* is dropped. Otherwise, a new entry is created, and the *Interest* is forwarded using the routing information in the FIB. If no matching route is found, the *Interest* can be discarded or broadcast, depending on the routing policy of the router.

Similarly, Fig. 1b presents the *Data* forwarding process. When an NDN router receives a *Data* packet, it checks the PIT. If a matching PIT entry is found, it caches the *Data* before forwarding it to all the corresponding interfaces in the PIT entry, and then this entry is removed. If there is no matching PIT entry, it means that the NDN router did not request this *Data*, hence the *Data* drop. One can note that there are several reasons why a *Data* may not match any PIT entry, among which: (1) received *Data* may violate specific requirements stated in the *Interest* (for instance users can specify that *Data* must be fresh or exclude some versions) (2) PIT entry may have been removed either because it has already been resolved or due to too long a response delay and (3) the *Data* is received from an interface which is different from the one over which *Interest* has been forwarded. This last case has been studied in [12], leveraging the fact that a router cannot verify all signatures [13], to implement an attack in which an attacker can poison responses and consequently related in-network caches, by sending unsolicited corrupted data for any random (popular) content name.

The whole process ensures that one *Interest* results in only one *Data* packet.

B. Interest Flooding Attack

New network components and protocols come with novel potential attacks and NDN is not an exception, even though it is based on the paradigm of security-by-design [14], [15]. The most important threats identified in NDN are related to privacy [16], due to the in-network data caches, content forgery or pollution [13] to fool their popularity, and PIT overloading [17]. Detailed reviews on NDN security can be found in [18], [19].

In this paper, we focus on an attack that aims at overloading the PIT and referred to as *Interest flooding attack* (IFA).

Roughly speaking, it essentially is a variation of the Denial of Service (DoS) attack in NDN [20]. The principle of IFA consists in sending a lot of malicious *Interest* packets for non-existent contents. Such *Interests* cannot be resolved by any *Data*. Hence, the corresponding PIT entry cannot be removed. When the PIT is overloaded, new *Interest* packets cannot be handled because there is no more room to create a PIT entry and they are thus dropped. This attack can have serious consequences on the network. Given that *Interests* for non-existing content are extremely easy to generate, they can cause large-scale damages on the network infrastructure.

1) *Detection and Mitigation*: Several solutions for detecting and mitigating IFA have been proposed to date, see for instance [18] for a recent review. In [21], Dai et al. present their *Interest* trace back mitigation strategy. When the PIT's size exceeds a threshold, a spoofed *Data* is created by the NDN router to respond to a long-unsatisfied *Interest*. These *Data* are eventually forwarded back to the source of the attack by tracing PIT entries. At the same time, NDN routers also limit the incoming packet rate of interfaces to which they send spoofed *Data*.

In [22], Tang et al. identify the compromised name prefixes used to launch IFA and then announce these malicious prefixes to neighbors. There are two phases in this identification process: rough detection and accurate detection. In the first phase, malicious interfaces are detected by computing a *satisfaction ratio*, a ratio between the number of outgoing *Data* and incoming *Interest* on an interface. When this ratio exceeds a threshold, the interface is considered under attack. The threshold of this phase is preconfigured for all cases. In the accurate detection phase, expired *Interests* on the reported interface are recorded. The prefix that has the largest expired ratio is considered hostile.

Having the same idea of using statistics to identify harmful interfaces, the Poseidon approach [23] maintains two measurements: the satisfaction ratio and the PIT space used up by *Interest* from each interface. Once an alarm occurs, an NDN router issues an alert message to its neighbors on the malicious interface. When an NDN router receives an alert, it also triggers the same countermeasure but with a lower threshold, to better identify the compromised interface.

Among previously mentioned detection and mitigation proposals, the satisfaction-based push back [24] is the most notable one. This proposal is similar to Poseidon: routers exchange announcements to neighbors and adjust their reactions based on these messages. Although this solution monitors the satisfaction ratio, it does not have a separate detection phase. The ratio is used to calculate periodically *Interest* limits in announcements between routers.

A fundamentally different approach has been proposed in [25] which questions the necessity of the stateful nature of ICN. Given IFA's attack mechanism, the authors show that if the NDN protocol is modified into a stateless fashion, the IFA attack becomes much less efficient. However, this solution requires redesigning the whole NDN concept and is only evaluated with IFA designed for an NDN stateful protocol.

It is also noteworthy that several prior works, such as [26], [27], proposed to study statistical methods for IFA detection.

In [26] the author proposed to use the well-known Gini coefficient to measure a side effect of IFA: the discrepancy in the range of requested content. The authors in [27] use a model of the IFA attack in the time-frequency domain using wavelets transform. Based on a statistical model of lower frequencies sub-band, a simple statistical method is proposed.

2) *NACK packet*: Though all those works constitute interesting solutions for IFA detection and mitigation, a recent update in NDN implementation prevents *Interests* to remain awhile in the PIT: instead of addressing a standalone solution against IFA, the authors of [28] proposed to extend the NDN forwarding mechanism by introducing the *NACK* packet. This mitigation is noteworthy since it is integrated into the NDN reference implementation since its version 0.5.1 and potentially deprecates all previous detection proposals in this area. Indeed, as noted in [29], the implementation of *NACK* prevents all attacks scenario in which it is assumed that an *Interest* for a non-existing content will stay a very long time, if not forever, in the PIT since a *NACK* packet for such *Interest* will be issued. When an NDN router can neither satisfy nor forward an *Interest*, a *NACK* is sent to downstream routers with an error code. There are three *NACK* error codes implemented, namely (1) *Duplicate*, (2) *Congestion*, and (3) *No Route*. The first code indicates that the router is still waiting for the *Data* of an identical *Interest* packet. The second one implies that the *Interest* cannot be forwarded due to congestion occurring on the outgoing link and the last one means that the router does not have any eligible route in FIB to forward the *Interest*. As such, downstream routers can adapt their sending rates to avoid overloading the upstream. Also, it helps downstream routers determine the cause of *NACK* to decide further forwarding strategies. Besides, *NACK* prevents *Interest* messages carrying names of non-existing content from being forwarded further to occupy other routers' PIT.

Although *NACK* is an effective mechanism to mitigate IFA, there are still several limitations. First, the *NACK* mechanism is based on the *Interest* prefix to mitigate the overload. If an attacker can send *Interests* with different prefixes, the router must issue a *NACK* for each prefix. Such a reaction can be burdensome while it could be more efficient to limit the sending rate of the face under attack. Secondly, using *NACK* means that a router's forwarding strategy depends on its upstream. A malicious upstream can hence add a delay to its response to postpone the sending of *NACK* to downstream, thus enabling a vulnerability for IFA. Finally, dealing with a non-existing name may be time-consuming, especially when several routes are available because the router will try all available routes before giving up and sending a *NACK* to downstream [28].

C. Common Limitations of Prior Works

Though the topic of security in ICN and NDN in particular has been widely studied, all prior works share three fundamental limitations. First of all, they were all designed and assessed in a purely simulated setup. While simulations may be very useful, real-life experimentation is crucial to evaluate the efficiency of a solution in practice. The real deployment

of an NDN network, coupled with the legacy Internet based on IP, is extremely time consuming but worth considering; this explains why almost all prior works on IFA detection and mitigation, see for instance [25]–[27], and in general on solutions for NDN security issues, see for instance [30], [31] did not make this final step. But, as shown in the present paper, the efficiency of a detection method measured in a completely simulated environment may hardly transfer into real life. In fact, it is somewhat superficial to design a detection tool under several assumptions on network traffic and to assess the proposed detection scheme under exactly the same conditions (simulated). In this paper we observed, for instance, that in real environments the traffic under an IFA attack greatly differs from the widely used and simplistic model of mere increase in loss-packet rate and that this may deprecate almost all prior works on IFA detection and mitigation.

Another common limitation of almost all prior works lies in their unknown performance. Indeed, even statistical methods, such as those proposed in [26], [27] are only evaluated empirically and, again, in simulated environments. Empirical evaluations, on their own, do not allow to guarantee any reliability of the results whose properties are unknown when practical setup is changed. The main limitation of detectors whose statistical performance is not analytically established is that they can meet a prescribed false alarm rate while, given the amount of network components, this is unacceptable for practical application.

Eventually, for a disruptive architecture like NDN, it is highly unrealistic to assume that it will completely replace IP networks all of a sudden. Indeed, NDN will likely coexist with current network architectures before being deployed at large scale. As such, dedicated gateways enabling the seamless translation of traffic crossing different protocol stacks will be implemented and deployed. They will undoubtedly alter the network traffic behavior and thus need to be considered in the scenario an attacker plans to implement. This phenomenon is particularly relevant given the stateless nature of the IP protocol as opposed to NDN which is of a stateful one. However, most previous works on IFA validate their solutions in pure NDN networks, thus questioning their performance in a real deployment.

D. Contribution of the Present Paper

The present paper addresses the three fundamental limitations of prior art stated in the previous Section II-C. First of all, by using the statistical hypothesis testing theory, we overcome the drawback of detection approaches designed and evaluated over solely empirical evaluations. It is thus proposed to design a detector with a well-defined threshold that is calculated according to the user desired false-alarm rate that works independently on each router. We show that the proposed method allows the guaranteeing *in practice* of a prescribed probability of false-alarm (PFA) as well as the accurate calculation of ensuing statistical power. Second, as described in Section II-C, shifting from simulated to real data traffic is not straightforward. We have implemented the proposed detection method into a real testbed and found

that the simple model for IFA network traffic is not accurate. This requires modification on the model used to build and extract IFA attack footprints. Eventually, for maintaining high performance, it is proposed to cast IFA detection within a sequential framework, meaning that the detection is not only carried out for each sample independently but also uses previous results to improve the detection accuracy.

Finally, our work addresses a credible deployment scenario in which NDN is integrated into the current Internet through dedicated islands interconnected with application layer gateways (i.e. HTTP to NDN and inversely in our paper). This scenario is from our perspective relevant since (1) it stands for one of the most credible options to date for an NDN deployment (2) it is easily achievable by an attacker operating some IP hosts (no need to perpetrate some intrusive actions to directly act on the NDN island) and (3) it exhibits a perfect normal appearance regarding an IP network, due to its stateless nature, while being wasteful for the sole NDN island. One can also note that the presence of the application layer gateways change the way the packets are generated and forwarded. As such, all attacks, as well as related detection solutions, must be reinvestigated as compared to IFA attacks perpetrated over pure NDN networks.

The present paper is based on our prior works [6], [7], [32] that study the IFA detection using hypothesis testing theory. The proposed statistical test is thus similar to those from our prior work. However, those prior works only cover the first aforementioned contribution and is only evaluated using simulated traffic. The present paper extends, thus, our prior papers [6], [7], [32] into the following directions. Firstly, by leveraging the implementation described in [32] to generate real traffic in NDN islands, it applies and evaluates the proposed statistical method for IFA detection into conditions which are as close as possible to what could happen in a real operated network. Secondly, this work (1) integrates application layer gateways that affect NDN traffic and (2) circumvents *NACK* packets that intrinsically mitigate IFA effects. By doing so, it is also shown that the network traffic corrupted in the context of an IFA attack differs greatly from what has been assumed so far in pure NDN networks and this requires modification of the detection method.

III. REVISED INTEREST FLOODING ATTACK SCENARIOS

In this section, we take into account the existence of *NACK* packets and investigate the feasibility of IFA scenarios in the literature within the context of (1) native NDN and (2) NDN coupled with the current IP network.

A. Circumventing *NACK* for IFA

To inspect IFA's impact in a native NDN network with *NACK*, we implemented a basic test environment with a single NFD node and monitored its PIT size's evolution. As the PIT gets overloaded, the NFD process crashes. Currently, no protection scheme can prevent such a phenomenon. Therefore, we consider this event, called the PIT collapse point, as an indicator of a successful IFA in native NDN. In this subsection, we reveal credible scenarios to collapse the PIT, and feature factors impacting this phenomenon.

Fig. 2. NFD collapse point according to *Interest* rate and lifetime

1) *Attack Scenarios*: Among all the setups we have tested to stress the router's PIT, there are three scenarios leading to an abnormal PIT increase or the node collapse point, described as follows:

- **Congesting the link between routers and the provider**: the attacker sends a large number of *Interests* in a short time to congest the link between routers and the provider. Consequently, the provider cannot send *NACKs* to notify the router. Therefore, at the time of congestion, the router is under attack without the presence of *NACK*.

- **Accumulating PIT entries with No Data *NACK***: this scenario exploits the vulnerability design of the *No Data NACK*⁵, which allows the PIT to keep an *Interest* until its lifetime expires even if it received a *NACK* [28]. The PIT entry is removed only when the router has no available face in FIB to send the *Interest*.

- **Delaying the response with a malicious provider**: an attacker-controlled provider will delay the response of *Interests*. The delay should be relatively long to occupy routers as long as possible and must be lower than the *Interest* lifetime so that *NACK* is not sent. Consequently, the downstream will not receive any *NACK* packet while its PIT accumulates entries.

2) *Impacting factors of the PIT overload phenomenon*: Based on the scenarios presented, we identified the factors that augment the attack impact. We selected the second scenario because it stands for the straightest case for an attacker to exploit vulnerabilities in the current NDN design and its NFD implementation. The results show that the attack effectiveness is impacted by the following factors:

- **Attack power**: Given by *Interest*'s lifetime and the attack rate in *Interest* per second, these factors facilitate the reach of the PIT collapse point [33]. A collaborated malicious user can flood NFD with large *Interest* lifetime values to multiply the IFA impact and currently there is no protection in NFD to prevent this phenomenon. The results of these experiments are depicted in Fig. 2 which presents an expected behavior, with a constant limit in terms of the number of *Interest* featured by an attack rate inversely proportional to the *Interest* lifetime. These results also assess the potential vulnerability of NDN to the accumulation of PIT entries in case of *No Data NACK*. It mainly shows that an attacker can perform flooding attacks with a small *Interest* rate by merely

⁵At the time of our previous work, the *NACK* is implemented with *No Data* code, not with *No Route* code as currently

Fig. 3. NFD collapse point according to the number of prefix levels

extending the *Interest* lifetime.

- **Number of prefix levels:** The naming convention of NDN follows a hierarchical scheme, which is similar to URI and currently, there is no limit to the number of prefix levels. To measure the impact of this factor, we have created prefixes with various levels with a constant length of 522 characters. The result of Fig. 3 shows that an increase in the number of prefix levels drastically reduces the PIT collapse point; even worse, this phenomenon seems non-linear, and can affect the PIT size, in terms of the number of *Interests*, by three orders of magnitude when the number of prefixes levels increases by 100. To the best of our knowledge, the importance of this factor's impact has not been identified to date. Thus, it introduces an easy-to-exploit flaw that the attacker can use to perform an IFA with limited resources.

- **Length of *Interests*:** Since the implementation of the PIT in NFD is designed as a data structure hosted directly in the NFD process memory, the length of *Interests* names exhibits a clear impact on the PIT collapse point. Hence, the more complex the *Interest* name, the more memory space is required.

- **Memory allocation:** The memory capacity allocated to the NFD process has an important impact on the PIT capacity. Our experiments show that the PIT collapse point is proportional to the amount of allocated memory.

B. Implementing IFA over an NDN Island Deployed in the Current Internet

1) *On the Coupling of NDN with IP:* As a credible deployment scenario for NDN, we investigate a use-case where an Internet Service Provider (ISP) couples an NDN network to the existing IP network to provide the HTTP service to users, as illustrated in Fig.4. Addressing web services is a relevant step toward the integration of NDN into the existing networks because it is among the most popular on the Internet. In this scenario, an NDN island is deployed inside the ISP's core network to leverage the benefits of its caching system and low-latency data delivery for a substantial part of traffic. Thanks to recent virtualization techniques (e.g., Network Function Virtualization), NDN routers can be deployed without the need for dedicated hardware or causing interruptions to existing ISP networks.

As the current Internet and web users do not implement NDN, their IP traffic is forwarded to dedicated gateways [34]

Fig. 4. Context of NDN coupled with HTTP application

which translate HTTP traffic to NDN traffic and vice versa. Two types of gateways are necessary to translate HTTP traffic over IP into NDN packets: *ingress gateway (iGW)* and *egress gateway (eGW)*. The operation of *iGW* and *eGW* is briefly demonstrated in Fig. 5. When an HTTP request arrives at *iGW* (arrow 1), it is translated into *Interests* and injected in the NDN network. *Interests* that cannot be satisfied by NDN routers' cache will reach the *eGW*. The *eGW* checks *Interest*'s name for fragmentation and retrieves remaining chunks of the HTTP request if needed (arrows 3 and 4). After that, the *eGW* reconstructs the original HTTP request and sends it to the corresponding HTTP server (arrow 5). When an HTTP response arrives (arrow 6), the *eGW* converts it into *Data* packets and sends them into the NDN network, reaching the *iGW* (arrow 7). The first *Data* is considered as a response to the *Interest* from the arrow (2). The *eGW* also includes information about fragmentation in the *Data* name, so that the *iGW* can retrieve remaining chunks of the HTTP response (arrows 8, 9). Afterward, the *iGW* reconstructs the HTTP response and delivers it to the client. All of these operations and the existence of the NDN island is entirely transparent to network users. Hence the users can still experience the benefits from NDN without any adaptation effort from their side.

2) *Proposed Attack Scenario:* Fig. 5 reveals an exploitable flaw to corrupt the NDN network: after a long enough delay between the HTTP request and the corresponding HTTP response (arrows 5 and 6), the first *Interest* (arrow 2) expires. Hence, the first *Data* (arrow 7) is considered unsolicited and rejected by *iGW*. As such, only the third attack scenario identified in section III-A1 enables an attacker to implement an IFA in the operational context presented above. More precisely, by leveraging a botnet or an equivalent means, the attacker can own the control of multiple web users and a malicious web server on the Internet (see Fig.4). Attacker-controlled users will browse for the website hosted on the malicious server. With the existence of *NACK* packets, the attack relies on intentionally adding a large delay to the response from the malicious server so that *Interests* exchanged in the NDN core network will linger in the router as long as possible. Moreover, by using a malicious server and delaying its response, the attacker can bypass the protection mechanism against IFA of *NACK* packet. Furthermore, since malicious

Fig. 5. Interaction between ingress and egress gateway

requests occupy NDN routers during the attack, legitimate users will suffer longer delays when accessing a website. It is also important to note that the use of NDN/HTTP gateways may slightly increase the delay to respond legitimate *Interests*. However, as noted in [34], the additional delay due to NDN/HTTP translation seems negligible. Besides, this additional delay affects all *Interests* (both legitimated or IFA generated). Eventually, it is worth noting, as a possible future work, that while NDN/HTTP gateways may allow novel attack vectors, those are outside the scope of the present paper which focuses on IFA attack.

IV. INTEREST FLOODING ATTACK DETECTION

This section presents our proposal to address the case of IFA detection. It first focuses on the local instantaneous detection for a given interface of a given router, assuming that the packet loss rate is known a priori. Then, we extend the local detection to address the case where the packet loss rate is unknown. To further enhance the detection accuracy, a sequential detection is presented. For the sake of clarity, the notation is simplified by omitting the index of the interface and the router.

A. Definitions

In the following, the number of incoming *Interest* packets and outgoing *Data* packets at an instant t , denoted as i_t and d_t respectively, are measured for each router's face. Ideally, each incoming *Interest* at a face should result in one outgoing *Data* packet. However, in any networks, part of the packets could be lost (even under regular legitimate uses, due to either transmission error, congestion or faulty hardware or cabling). Hence, let

$$\ell_t = 1 - \frac{d_t}{i_t} = \frac{i_t - d_t}{i_t} \quad (1)$$

be the measured packet-loss (unresolved *Interests*) rate at the instant t . It is worth noting that the ratio ℓ_t , see (1), represents the packet loss rate under the assumption the round-trip-time is small compared to the sampling period. Additionally, the loss packet rate ratio is negligibly affected by interests resolved over the next sampling period since those are compensated by the number of data packets associated with interests from the previous sampling period. This phenomenon is also taken into account in the present work which assumes that responses to

Interests are a random process.

Following the model proposed in [23], [24], it is assumed that, at the instant t , all *Interests* have the same probability of not being resolved, denoted as p_t . Under normal situation, such probability should correspond to the expectation of measured packet-loss rate, i.e., $\mathbb{E}(\ell_t) = p_t$. Therefore, d_t should follow a binomial distribution $\mathcal{B}(i_t, 1 - p_t)$ with expectation $\mathbb{E}(d_t) = i_t(1 - p_t)$. By contrast, when IFA occurs, a significant number of *Interests* are sent to the pirate server, resulting in an abrupt increase of the packet-loss rate ℓ_t .

To model the impact of IFA, let us denote N_a the number of malicious *Interests* sent, per unit of time, during the IFA by attacker-controlled hosts besides the legitimate *Interests*, denoted as i_t^* . Hence, the IFA can be characterized by an increase in the number of incoming *Interests*:

$$i_t = i_t^* + N_a. \quad (2)$$

One should note that it is impractical to distinguish legitimate *Interests* i_t^* from the whole flow of *Interest* packets i_t . Moreover, because the N_a additional malicious *Interests* cannot be responded to, the expectation of overall packet-loss rate is increased as follows:

$$a = \mathbb{E}(\ell_t) - p_t = \frac{(1 - p_t)N_a}{i_t^* + N_a}. \quad (3)$$

The above relation comes from the fact that, whether an IFA is currently happening or not, the expected number of *Data* packets received at a given face remains the same:

$$\begin{aligned} (1 - p_t)i_t^* &= \mathbb{E}(d_t) = (1 - p_t - a)(i_t^* + N_a), \\ \Leftrightarrow \frac{(1 - p_t)N_a}{i_t^* + N_a} &= a. \end{aligned}$$

B. Detection Problem Statement

To facilitate the problem description, let us assume that the expected packet-loss rate p_t is known. According to the characterization in Section IV-A, the IFA detection problem consists in choosing between two hypotheses: $\mathcal{H}_0$: “the number of incoming *Interests* i_t and outgoing *Data* packets d_t are consistent with what is expected from p_t ” and $\mathcal{H}_1$: “ d_t is significantly lower than what is expected from i_t and p_t ”. Those two can be written formally as a choice between the following statistical hypotheses:

$$\begin{cases} \mathcal{H}_0 : d_t \sim \mathcal{B}(i_t, 1 - p_t), \\ \mathcal{H}_1 : d_t \sim \mathcal{B}(i_t - N_a, 1 - p_t), N_a > 0. \end{cases} \quad (4)$$

Formally, a statistical test is a mapping $\delta : \mathbb{R} \mapsto \{\mathcal{H}_0; \mathcal{H}_1\}$, i.e., hypothesis $\mathcal{H}_i$, $i \in \{0, 1\}$ is accepted if $\delta(x) = \mathcal{H}_i$ (see [35] for a thorough introduction to hypothesis testing). We focus on the Neyman-Pearson bi-criteria approach that simultaneously aims at guaranteeing a prescribed Probability of False-Alarm (PFA) while maximizing the power function (or correct detection probability). Let $\mathbb{P}_i(E)$, $i \in \{0, 1\}$ be the probability of the event E under the hypothesis $\mathcal{H}_i$. For a prescribed PFA α_0 , the Neyman-Pearson approach aims at finding a test δ whose PFA is upper bounded by α_0 . Hence, let:

$$\mathcal{K}_{\alpha_0} = \{\delta : \mathbb{P}_0[\delta(\ell_t) = \mathcal{H}_1] \leq \alpha_0\}, \quad (5)$$

be the class of all tests whose PFA are upper bounded by α_0 . The goal is to find in $\mathcal{K}_{\alpha_0}$ a test that maximizes the power function, or correct detection probability, formally defined by:

$$\beta_\delta = \mathbb{P}_1[\delta(\ell_t) = \mathcal{H}_1]. \quad (6)$$

The hypotheses formulated in (4) highlight the main difficulties of the present testing problem. First, we note that the IFA implies a change in both expectation and variance of the measured loss-packet rate ℓ_t . Secondly, the parameters of attack payload N_a or a (see (2) and (3)) are unknown. Ideally, the test δ should be Uniformly Most Powerful (UMP), that is it should maximize the power function β_δ , regardless the attack payload a [35, Chap.3]. Unfortunately, UMP test scarcely ever exists. Thirdly, the greatest difficulty is that the expected loss-packet rate is unknown in practice. It therefore has to be estimated from data and this problem is studied in detail in Section IV-D.

C. Optimal Likelihood Ratio Test for Known Loss Rate

This section presents the theoretical optimal Likelihood Ratio Test (LRT) and assesses its statistical performance. Since the binomial law belongs to the exponential distribution family, there exists a UMP test that is given by the following decision rule, see [35, Corollary 3.4.1]:

$$\delta^*(d_t) = \begin{cases} \mathcal{H}_0 & \text{if } d_t \geq h(i_t; p_t), \\ \mathcal{H}_1 & \text{if } d_t < h(i_t; p_t), \end{cases} \quad (7)$$

where the threshold $h(i_t; p_t)$ depends on both the number of *Interests* packets i_t and expected pack-loss rate p_t , so that $\delta^* \in \mathcal{K}_{\alpha_0}$ (5). However, evaluating the statistical properties of a test such as (7) is hardly possible. Besides, the decision threshold must always be recomputed because it depends on parameters i_t and p_t that change at each instant t . In short, though a test such as (7) is simple to build, it is hardly usable in practice.

Therefore, it is possible to simplify the IFA detection problem by applying the central limit theorem (CLT) [35, Theorem 11.2.5], assuming that the number of *Interests* sent i_t is large, which is a usual case for a router face. Hence, d_t under $\mathcal{H}_0$ can be modeled as:

$$d_t \rightsquigarrow \mathcal{N}(i_t(1-p_t), i_t p_t(1-p_t)), \quad (8)$$

where $\rightsquigarrow$ represents the convergence in distribution as i_t tends to infinity. Let us define the residual packet-loss rate r_t as the difference between observed and expected loss rates: $r_t = \ell_t - p_t$. Under $\mathcal{H}_0$, the residual one will be:

$$r_t = \left(1 - \frac{d_t}{i_t}\right) - p_t \rightsquigarrow \mathcal{N}\left(0, \frac{p_t(1-p_t)}{i_t}\right). \quad (9)$$

On the opposite, when an IFA happens, the residual tends to:

$$r_t \rightsquigarrow \mathcal{N}\left(a, \frac{p_t(1-p_t)}{i_t} - \frac{N_a p_t(1-p_t)}{i_t^2}\right). \quad (10)$$

Let us denote σ_t^2 the variance under $\mathcal{H}_0$ and σ_a^2 the decrease of variance due to the IFA:

$$\sigma_t^2 = \frac{p_t(1-p_t)}{i_t}, \quad \sigma_a^2 = \frac{N_a p_t(1-p_t)}{i_t^2} = \frac{a p_t}{i_t}. \quad (11)$$

Interestingly, the two terms from variances in Eq. (9)-(10) find their origin in distinct phenomena. On the one hand, σ_t^2 is a directly related random aspect of packet resolution while, on the other hand, the term σ_a^2 is due to the fact that IFA attack *Interests* surely generate no *Data* back which thus “reduces the randomness” of measured packet-loss rate.

One can note that the decrease of variance is due to the increase of i_t during the attack, i.e., from $i_t = i_t^*$ in (9) to $i_t = i_t^* + N_a$ in (10) while the number of received *Data* packets does not change. It follows from Eq. (9) - Eq. (10), that the testing problem (4) can be reformulated as:

$$r_t \sim \begin{cases} \mathcal{N}(0, \sigma_t^2) & \text{under } \mathcal{H}_0, \\ \mathcal{N}(a, \sigma_t^2 - \sigma_a^2) & \text{under } \mathcal{H}_1. \end{cases} \quad (12)$$

Equation (12) shows that parameter a completely characterizes the impact of IFA on packet-loss rate, hence it is used in the remaining of this paper to quantify the attack payload.

Based on the previous equations (8) – (12), it is now possible to define an optimal test for IFA detection as stated in the following theorem:

Theorem 1. *With r_t defined as in (9), the following test:*

$$\delta^*(r_t) = \begin{cases} \mathcal{H}_0 & \text{if } r_t \leq \tau^*, \\ \mathcal{H}_1 & \text{if } r_t > \tau^*, \end{cases} \quad (13)$$

is an Asymptotically Uniformly Most Powerful (AUMP) test for the testing problem (4).

Proof. The proof provided in Appendix A demonstrates the asymptotic optimality of the test (13). $\square$

As previously discussed, the application of the CLT (9) allows establishing the statistical properties of the optimal UMP test, presented by the following proposition:

Proposition 1. *Assuming that the number of *Interests* i_t tends to infinity, for any prescribed PFA α_0 , the decision threshold, τ^* , given by:*

$$\tau^*(\alpha_0) = \sigma_t \Phi^{-1}(1 - \alpha_0), \quad (14)$$

guarantees that the test δ^ (13) is in $\mathcal{K}_{\alpha_0}$. Here Φ and Φ^{-1} are the standard normal cumulative distribution function and its inverse function, respectively. Using the decision threshold given in (14), the power function of the UMP test δ^* (13) is given by:*

$$\beta_{\delta^*}(a) = 1 - \Phi\left(\frac{\tau^*(\alpha_0) - a}{\sqrt{\sigma_t^2 - \sigma_a^2}}\right). \quad (15)$$

The assessment of the statistical performance of AUMP test $\delta^*(r_t)$ serves as an upper bound on the detection performance one can expect from any practical IFA detection method. Another interesting aspect of the proposed asymptotic approach is that it is possible to set a threshold satisfying a prescribed PFA. This threshold only depends on the desired PFA α_0 , i_t and p_t which are all known. This approach simplifies the problem of dealing with a binomial distribution whose cumulative distribution function is difficult to compute. However, a notable consequence of the underlying binomial distribution is that the residual packet-loss rate r_t has both its

expectation and its variance impacted by the IFA. Hence the power function of proposed AUMP test not only depends on the attack payload a but also on the impact on the variance of r_t through the denominator $\sqrt{\sigma_t^2 - \sigma_a^2}$.

D. Generalized Likelihood Ratio Test

This section addresses the case where the expected packet-loss rate p_t is unknown. In such a situation, a usual approach consists in designing a Generalized LRT by substituting the unknown parameter (the expected packet-loss rate p_t in our case) with its estimation using the Maximum Likelihood.

1) *Packet-loss Rate Model*: As a first step, it is proposed to gather the N last measurements of packets-loss rate $\ell_t = (\ell_{t-N+1}, \dots, \ell_t)$. Since the fluctuation of the packet-loss rate is limited and smooth [36], [37], its expectation can be modeled by a polynomial:

$$\mathbf{p}_t = \mathbf{H}\mathbf{x}_t, \quad (16)$$

where $\mathbf{H}$ is a matrix of size $N \times q$ whose elements $h_{n,j} = n^{j-1}$, $n \in \{1, \dots, N\}$, $j \in \{0, \dots, q-1\}$ and $\mathbf{x} = (x_0, \dots, x_{q-1})$ is the vector of the q coefficients of the polynomial. Such a model has been widely used in signal processing, see [38]–[41] for applications in Internet traffic modeling and image processing.

Assuming that packet-loss rate measurements ℓ_t are statistically independent, it follows from previous asymptotic distribution (8) that under the hypothesis $\mathcal{H}_0$, the observations can be modeled as:

$$\ell_t \rightsquigarrow \mathcal{N}(\mathbf{H}\mathbf{x}_t, \mathbf{\Sigma}_0), \quad (17)$$

where $\mathbf{\Sigma}_0$ is a diagonal covariance matrix whose elements are given by $\frac{p_u(1-p_u)}{i_u}$, $u \in \{t-N+1, \dots, t\}$.

When an attack is started at the instant t^{th} , the packet-loss rate drops for the very last samples (before affecting all inspected samples). Hence under hypothesis $\mathcal{H}_1$, as i_T tends to infinity, the packet-loss can be modeled as:

$$\ell_t \rightsquigarrow \mathcal{N}(\mathbf{H}\mathbf{x}_t - a\mathbf{v}_a, \mathbf{\Sigma}_0 - \mathbf{\Sigma}_a), \quad (18)$$

where $\mathbf{\Sigma}_a$, as in Eq. (11)–(12), represents the variance decrease due to the IFA generated *Interests*; $\mathbf{v}_a$ represents the number of samples whose loss-packet rate is changed due to the attack and is to be set according to the user's desire since, as shown in Eqs.(23)–(26) and in Section V-A there is a tradeoff between quick and accurate detection; the user can set, for instance, $\mathbf{v}_a = (0, 0, \dots, 0, 1)^T$, implying that only the very last sample is affected by the IFA attack, for the quickest detection and can set $\mathbf{v}_a = (0, \dots, 0, 0, 1, 1, 1, 1)^T$ to detect IFA footprint over the 5 last samples with higher accuracy at a cost of delayed detection.

Under the Gaussian distribution model, it is well known that the maximum likelihood estimation of packet loss rate $\mathbf{p}$ is equivalent to the least square estimation:

$$\tilde{\mathbf{p}}_t = \mathbf{H}(\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T\ell_t. \quad (19)$$

The ensuing residuals $\mathbf{r}$ are defined, as in Equation (9), as:

$$\tilde{\mathbf{r}}_t = \ell_t - \tilde{\mathbf{p}}_t = \mathbf{H}^\perp \ell_t, \quad (20)$$

where $\mathbf{H}^\perp = \mathbf{I}_N - \mathbf{H}(\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T$, with $\mathbf{I}_N$ the identity matrix of size N , represents the projection onto the orthogonal complement of the subspace spanned by the columns of $\mathbf{H}$.

2) *Generalized Likelihood Ratio Test for Unknown Loss Rate*: Following the model of the packet-loss rate under each hypothesis (17) - (18) and the definition of residuals $\tilde{\mathbf{r}}_t$ (20), the IFA detection problem for unknown loss rate can be formulated as a choice between the following hypotheses:

$$\begin{cases} \mathcal{H}_0: \tilde{\mathbf{r}}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{H}^\perp \mathbf{\Sigma}_0 \mathbf{H}^{\perp T}), \\ \mathcal{H}_1: \tilde{\mathbf{r}}_t \sim \mathcal{N}(a\tilde{\mathbf{v}}_a, \mathbf{H}^\perp \mathbf{\Sigma}_0 \mathbf{H}^{\perp T} - \mathbf{H}^\perp \mathbf{\Sigma}_a \mathbf{H}^{\perp T}), \end{cases} \quad (21)$$

where $\tilde{\mathbf{v}}_a = \mathbf{H}^\perp \mathbf{v}_a$ is the obtained IFA footprint after estimating and removing the expected loss-packet rate (20). Here, it can be noted that, as previously discussed in Section IV-C, the IFA impacts both the expectation and the covariance of the residuals.

As previously explained, by replacing in the LRT the estimated packet-loss (19), it is proposed to design a GLRT as follows:

$$\tilde{\delta}(\tilde{\mathbf{r}}_t) = \begin{cases} \mathcal{H}_0 & \text{if } \tilde{\mathbf{v}}_a^T \tilde{\mathbf{r}}_t \leq \tilde{\tau}, \\ \mathcal{H}_1 & \text{if } \tilde{\mathbf{v}}_a^T \tilde{\mathbf{r}}_t > \tilde{\tau}. \end{cases} \quad (22)$$

From the distribution of the residuals $\tilde{\mathbf{r}}$ (21), it is straightforward that:

$$\tilde{\mathbf{v}}_a^T \tilde{\mathbf{r}}_t \rightsquigarrow \begin{cases} \mathcal{N}(\mathbf{0}, s_0^2) & \text{under } \mathcal{H}_0, \\ \mathcal{N}(a\|\tilde{\mathbf{v}}_a\|_2^2, s_0^2 - s_a^2) & \text{under } \mathcal{H}_1. \end{cases} \quad (23)$$

where the GLR variance s_0^2 under $\mathcal{H}_0$ and the decrease of variance s_a^2 under $\mathcal{H}_1$ are given by:

$$s_0^2 = \mathbf{v}_a^T \mathbf{H}^\perp \mathbf{\Sigma}_0 \mathbf{H}^{\perp T} \mathbf{v}_a, s_a^2 = \mathbf{v}_a^T \mathbf{H}^\perp \mathbf{\Sigma}_a \mathbf{H}^{\perp T} \mathbf{v}_a. \quad (24)$$

Similarly to Proposition 1, based on the distribution of the GLR (23), we can establish the decision threshold and the power function of the proposed GLRT:

Proposition 2. *Assuming that the number of incoming Interest i_t tends to infinity, for any prescribed PFA α_0 , the decision threshold $\tilde{\tau}$ given by:*

$$\tilde{\tau} = \Phi^{-1}(1 - \alpha_0) s_0, \quad (25)$$

guarantees that the test $\tilde{\delta}$ (22) is in $\mathcal{K}_{\alpha_0}$. Using the decision threshold given in (25), the power function of the UMP test (22) is given by:

$$\beta_{\delta^*}(a) = 1 - \Phi\left(\frac{s_0\Phi^{-1}(1 - \alpha_0) - a\|\tilde{\mathbf{v}}_a\|_2^2}{\sqrt{s_0^2 - s_a^2}}\right). \quad (26)$$

From the power function (26), one can note that the loss of optimality of the proposed GLRT is mainly caused by the factor $\|\tilde{\mathbf{v}}_a\|_2^2$. This is explained by the estimation of the unknown and dynamic packet-loss rate. When the IFA starts, the packet-loss rate changes suddenly. However, a non-negligible proportion of such change is modeled as a regular change of the dynamic legitimate traffic.

E. Sequential Detection

The tests presented in Sections IV-C and IV-D2 are devoted to the analysis of a single router interface at a specific time t . However, in practice it seems natural to increase the accuracy of the proposed IFA detection by gathering consecutive samples. Indeed, the IFA footprint may be small enough to make the detection difficult at once, and collecting evidence over time may significantly ease the detection. Besides, as it will be presented in Section V, the traffic under IFA attack is rather different from the expected model of a “simple” increase in the loss-packet rate: it increases the loss-packet rate but also makes it very unstable, changing abruptly from very high to low values. Therefore, a test based on a single measurement may be quite inefficient in practice and, throughout a sequential approach, it is required to collect evidence of an IFA attack over several samples to overcome this behavior of loss-packet rate. Another important reason to analyze the data sequentially is to ease extending the proposed method over all interfaces from all routers to make a global monitoring system that can detect IFA. However, it is very unlikely that, when an IFA starts, its detection will be efficient for all interfaces and all networking devices at the same time. Hence, keeping a record of previous results is crucial. Consequently, this section presents an extension of the previously proposed “snapshot” GLR test (22) by taking into account previous observations within a sequential framework.

In the literature, the problem of change-point detection⁶ has been extensively studied. In brief, the sequential analysis framework not only aims at detecting a specific event with highest accuracy, regarding PFA and missed-detection probability but also introduces the delay as the third criterion of detection performance. More formally, a change-point detection scheme is defined by a stopping rule $S(\ell_1, \dots, \ell_t) \mapsto \{0, 1\}$ such that the IFA is detected at first time S_t for which $S(\ell_1, \dots, \ell_t) = 1$. Here, as in all that precedes, the values $\ell_1, \dots, \ell_t$ represent the loss packet rate, see Eq. (1), over which the IFA detection is carried out. Let us denote ν the IFA starting time and $S_t \geq \nu$ the instant when the attack is correctly detected. Thus, the detection delay is defined as $DD = S_t - \nu$.

Several methods have been proposed in the literature for change-point detection, among which two have been studied in the present paper. We first implemented a sliding window version of Wald Sequential probability Ratio Test (SPRT) as proposed in [44], [45]. However, we have observed empirically that the well-known CUMulative SUM (CUSUM), initially proposed in [46], has better detection accuracy in our cases. It is thus proposed to use it in the present work. More generally, the CUSUM has been shown to be optimal in several cases according to the so-called Lorden’s criterion [47] that consists in minimizing the average worst case detection delay, formally defined as:

$$\sup_{\nu \in \mathbb{N}} \mathbb{E} [S_t - \nu | S_t \geq \nu] \quad (27)$$

for a given worst case average Run Length To False Alarm (RL2FA), defined as:

$$\inf_{\nu \in \mathbb{N}} \mathbb{E} [S_t | S_t < \nu] \quad (28)$$

For a given interface at which observations $\ell_1, \dots, \ell_t$ are collected, see Eq. (1), the CUSUM C_t is defined as:

$$C_t = \max(0; C_{t-1} + \mathbf{v}_a^T \tilde{\mathbf{r}}_t - \kappa) \quad (29)$$

where $\mathbf{v}_a^T \tilde{\mathbf{r}}_t$, as defined in Eq. (23), corresponds to the Generalized Likelihood Ratio between hypotheses $\mathcal{H}_0$ and $\mathcal{H}_1$ computed with observation ℓ_t and with κ , a constant that has to be set. The main idea behind the CUSUM is to compute sequential LRs and reset it to zero whenever it goes below zero, given that observations are independent, and the change-point has not occurred yet. The constant κ can be interpreted as the sensitivity of the CUSUM. A large value for κ makes the reset of C_t to 0 more frequent but may delay the detection. On the other hand, a small κ allows a faster detection at the price of a less frequently reset CUSUM, hence smaller average RL2FA.

V. NUMERICAL RESULTS

In this section, we assess our proposed detection with both simulated data and real data. The simulated data is necessary to validate the intrinsic statistical properties of our detection since, in the simulation environment, we can entirely control the attack power regardless of the NDN network conditions. In a second step, the proposed detection is assessed with data from the real deployment of NDN coupled with IP. Each subsection begins with a description of the deployed topology, utilized tools, experiment setup, followed by the evaluations on the PFA guarantee and the detection power of our detection.

A. Assessment of the Statistical Properties

1) *Simulation tool and topology*: To generate simulated data, we use ndnSIM [11] – an open-source NDN simulation provided by the NDN project. Indeed, ndnSIM faithfully implements the components of an NDN network, allowing us to consider every aspect of the network [11]. Simulated data is processed offline using the *Matlab* numerical computation software. To compare the performance of our approach to existing ones, we reuse one of the topologies from [24], depicted in Fig.6 - a binary tree with eight hosts and one content provider. The topology represents one of the worst cases for IFA detection: indeed, all *Interests* (both legitimate and IFA generated ones) are forwarded to the sole content provider which covers the IFA traffic under the legitimate one. On the opposite, the detection is much easier if the pirated server only receives IFA generated *Interests* resulting in a much more important loss-packet rate under the attack.

2) *Simulation setup*: In all of our simulations, the actual number of *Interests* sent is generated from a Poisson distribution whose mean value is drawn from a uniform random variable. The actual packet-loss rate follows an auto-regressive (AR) model. Such a model has been extensively used to model both the evolution of users’ requests, and packet-loss rate in

⁶In the literature, see [42] and [43], the term “change-point” usually refers to the problem in which samples’ distribution changes at an unknown time ν .

Fig. 6. Topology for data simulation in ndnSIM

computer networks [37], [48]. More precisely, the packet-loss rate is given by $p_t = p_{t-1} + u$ with $p_0 = 0.05$ and u drawn from a uniform distribution with zero mean. Note that for a realistic behavior, the sign of u is flipped if $p_t < 0$ or if $p_t > 0.25$, the latter being quite a high value in practice. Several parameter values have been tested with similar results trends.

For the proposed GLRT, a set of $N = 50$ samples is used, and the polynomial's degree is $q - 1 = 4$, hence the 50×5 size of matrix $\mathbf{H}$. Unless explicitly stated otherwise, in all experiments we focus on the quickest detection, i.e., the calculation of the GLR $\tilde{\mathbf{v}}_a^T \tilde{\mathbf{r}}_t$, see Eq. (22), is carried out assuming that only the last sample can be affected by IFA generated *Interests*. Thus, the IFA footprint is characterized, by default, by $\mathbf{v}_a = (0, 0, \dots, 0, 1)^T$, leading to a footprint after packet-loss rate estimation $\tilde{\mathbf{v}}_a$ with $\|\tilde{\mathbf{v}}_a\|_2^2 \approx 0.6$.

3) *Results' analysis*: Fig. 7 shows a comparison between the theoretical PFA and detection power, given in Proposition 2, and the empirical ones. Even for a threshold that corresponds to $\alpha_0 = 10^{-3}$, the empirical results match well the theoretically established ones. This observation is important since it guarantees a prescribed PFA in a practical situation. This also shows the sharpness of the theoretical findings and relevance of the proposed model.

Fig. 8 then compares the theoretical and empirical power as a function of the IFA payload a for both optimal LRT and proposed GLRT. We also note that the power is computed with two prescribed PFA $\alpha_0 = 0.01$ and $\alpha_0 = 0.1$. The figure shows the relevance of the theoretical findings since empirical power functions match the theoretical ones. However, for low

Fig. 7. Comparison between proposed GLRT theoretical PFA and detection power, see Eq. (22), and empirical ones. The PFA and power are plotted as a function of the decision threshold $\tilde{\tau}$. Fig. 8. Comparison between empirical detection power and theoretical power function for both optimal LRT and the proposed GLRT. The power function is plotted as a function of the strength of the anomaly $a \in [0, 0.02]$.

PFA as $\alpha_0 = 0.01$, the number of required samples is very large, hence the lightly less accurate empirical results.

Fig. 9 takes one step beyond by showing a comparison between the theoretical and the empirical power of the proposed GLRT for three numbers of samples corrupted by IFA generated *Interests*, denoted $M = \|\mathbf{v}_a\|$, 1, 3 and 7. As one would expect, the power increases with the number of corrupted samples. This result emphasizes that the proposed method can be adapted to focus on the quickest detection, thus aiming at detecting only if the last sample is corrupted at the cost of lower detection accuracy. On the other hand, it is also possible to increase the detection delay, consequently focusing on the detection of several last samples corrupted by the IFA, to ensure a higher detection accuracy. Finally, for being comprehensive, Fig. 9 also proposes a comparison with the detector proposed in [24], which is based on a fixed threshold for loss-packet rate. Obviously, such an approach cannot deal with the non-stationary behavior of users and, so, Fig. 9 shows that such a detector performs significantly worse than the one proposed in the present paper.

B. Performance Evaluation under Realistic Conditions

1) *Use-case topology*: To validate the proposed method with real data, we deployed the NDN topology depicted in Fig. 10. The network consists of four nodes with NFD installed (v0.5.1 with NACK implemented). The *iGW* and *eGW* connect, respectively, to web users and the Internet. Users' traffic is generated by a web user emulator described below. On the other side of the network, we deploy a malicious server which connects to the *eGW* and runs an Apache HTTP server to collaborate with the attacker to perform IFA.

2) *Testbed deployment*: The use-case topology is deployed in OpenStack⁷ - a cloud operating system that helps control large pools of computation, storage, and networking resources throughout several physical hardware devices, enabling scalability for large-scale experiments. For each node in the topology, a virtual machine (VM) is created following the template configuration of OpenStack and hosts Ubuntu as the operating system where corresponding applications (e.g., emulator) are

⁷See: www.openstack.org.

Fig. 9. Receiver Operational Characteristic (ROC) curves for the proposed GLRT with different numbers of corrupted samples.

installed. These VMs connect to a virtual network provided by OpenStack Neutron, the OpenStack Networking Service. To collect the data, we use the Montimage Monitoring Tool⁸ (MMT) probe. A plugin for this probe has been developed to extract NDN protocol fields' values as well as to perform the proposed statistical tests.

3) *Web user emulator*: We developed a web user emulator based on Jaunt API⁹ - a Java library for web-scraping that allows retrieving objects in a web page such as images, CSS and javascript. The emulator first randomly selects a website from a given list of popular sites, based on a Zipf distribution - a well-known distribution to model contents popularity [49]. Next, it retrieves the website's object list, then checks for objects' existence on the Internet. If an object exists and does not require HTTPS connection, the emulator passes the HTTP request for that object toward the iGW to retrieve it. After loading the whole object list, the emulator waits for a while before selecting a link randomly on the website to browse. This behavior emulates web users' action of reading and clicking when navigating through a web page. Based on prior works that model users' requests time [50], the waiting time is drawn from the exponential distribution. The same process of loading and waiting is repeated for the secondly selected link. The emulator then selects another website in the given list and repeats the whole process.

4) *Experiment scenario*: Each VM can run several emulators at a time. One can change the amount of traffic generated by modifying the number of emulators and the number of threads given to each emulator (the more threads are given, the shorter the time to retrieve the object list). Each experiment lasts for 30 minutes, including 15 minutes under $\mathcal{H}_0$ followed by 15 minutes under $\mathcal{H}_1$ traffic. In all experiments, 30 emulators, each with 2 threads, are launched to generate legitimate traffic. Their average browsing time is set to 10s. Such a configuration for legitimate traffic generates, on average, approximately 80 HTTP requests/s. The list given to the emulator includes 90 websites chosen among the most popular ones. Using more websites would be unnecessary because websites located in the tail of the Zipf distribution have very low probability of being selected. The probe generates

Fig. 10. Use-case topology and testbed architecture

a sample after each 4s-interval. Each sample consists of the number of incoming *Interests* and outgoing *Data* for each interface of all routers. This setup for legitimate traffic and probe helps reduce the noise in packet loss rate, minimize unnecessary management effort while achieving a relatively high i_t^* , of about 320 legitimate *Interests* per second, for allowing the application of the CLT with sufficient accuracy. Besides, this sample period of 4 seconds constitutes a good trade-off between reactivity and accuracy since it is small enough for reaction and large enough to ensure that the number of *Interests* whose corresponding *Data* is received over the next sample period represents a negligible fraction of packets.

On the attacker's side, bad (or malicious) emulators will request only objects from the malicious server's site. We vary the attack power by changing the number of malicious emulators and number of threads given to each emulator. Because bad emulators want to send as many HTTP requests as possible, they do not take time to browse the site. Hence, we set up a very small value to their average browsing time (1ms). In the following sections, the attack power will be presented in terms of bad HTTP requests/s for easy understanding. Each attack setup is run for 10 times in order to increase the amount of data and, hence, reducing the statistical spread. The malicious server's delay is configured to $5s \pm 100ms$. Longer delay values have been tested, but they are so long that the emulator cannot establish a connection to the malicious server. The hosted website also contains a lot of objects, so that sessions to the malicious server will last longer, prolonging their loads on the NDN network.

5) *Characterization of attack phenomenon*: Under the effect of an IFA, one can expect that the NDN network is severely occupied. As a result, users will experience longer delays when loading and browsing a website. This phenomenon is visualized in Fig.11 which shows the delay's density functions of different scenarios. Note that these probability density functions have been estimated using the Parzen estimator [51]. One can remark that when there is no attack (black line), the distribution possesses a high density in the small value zone and consists of a thin tail, meaning that the delay value is

⁸See: www.doctor-project.org.

⁹See: jaunt-api.com

Fig. 11. Estimated density of delay under different attack setups

small and does not vary widely. On the other hand, when an IFA is occurring (green, blue, red lines), the density decreases around smaller values while the tail becomes thicker with a larger attack payload, showing a shift to higher response delay with a larger spread of latency. Such dispersion is enlarged with more powerful attacks.

As described in the usual behavior of the emulator, a website is randomly selected based on a Zipf distribution. As a result, those located at the tail of content popularity are less likely to be picked and, hence, will contribute less to the density presented above. To provide an overall view of the IFA effect on the whole website list with various delays, we record the latency 20 times for all websites under different attack scenarios. No click is emulated since randomly selected links may add more delay to the record, causing inconsistency between different sites. Websites that are too long to retrieve (> 300 s) are considered unreachable. Fig.12 illustrates the average delay under attack as a function of average delay under $\mathcal{H}_0$. This figure helps visualize the severity of the increased delay that users suffer when IFA happens. Each point represents the measurement of an individual website. For readability, the black dash line shows the equation $y = x$, that is a constant delay under $\mathcal{H}_0$ and $\mathcal{H}_1$, while solid lines show the results of an affine regression of these measurements, using least mean squared error, for each attack scenario. One can remark that the attack is still successful in terms of increasing the delay of legitimate clients even when it has a relatively small power (green line, equivalent to 6 bad HTTP requests/s). This trend, however, is not significant since it is quite close to

Fig. 12. Attack effect on increasing the delay of individual website

Fig. 13. Theoretical and empirical loss rate

the neutral line. When the attack power increases, the delay gets worse, as indicated by the increase of the trends' slope.

Also, we observe a difference between the theoretical and the empirical packet loss rate, depicted in Fig.13. In theory, when the attack starts, the packet loss rate is expected to increase and then remain stable at a specific value during the attack period. However, the empirical data shows that the packet loss rate repeatedly increases and drops during the attack. These changes occur quite abruptly and seem not to follow any particular pattern. This phenomenon can be explained by the fact that it has often been assumed that IFA creates an increase, in the packet-loss rate, while the proposed attack scenario, taking into account *NACK* packets, only creates a delay but the corresponding data eventually arrives much later. Thus, the loss packet rate during the attack changes quite abruptly depending on the exact number of attack *Interest* packets sent over each sampling period and the exact additional server delay, which are both stochastic processes. Note that for readability, Fig.13 presents a rather obvious attack scenario in which the loss packet rate is multiplied by a time factor after the IFA starts.

6) *Guarantee of false alarm probability and detector configuration selection*: One of the important properties of a statistical test is its guarantee of a prescribed PFA. To validate this aspect, we evaluated the empirical packet-loss rate under all the data collected under $\mathcal{H}_0$, that is in more than 48000 samples. For being comprehensive, several parameters of the proposed method have been selected, namely the window length N , model degree q and number M of corrupted samples it is aimed at detecting. Fig.14 depicts the empirically measured PFA as a function of the detection threshold τ as well as a comparison with the theoretical PFA given in (25). This figure shows that empirical results match theoretical ones in the range of $PFA > 5.10^{-3}$, implying the accuracy of our model under $\mathcal{H}_0$. Moreover, results from different configurations are close, showing the ability to guarantee PFA under various setups.

7) *Sequential detector performance*: To emphasize the advantages of gathering consecutive samples, Fig. 15 and 16 compare the performance of the snapshot test in Eq. (22) with the proposed sequential detection method based on the CUSUM in Eq. (29) on two different aspects. More specifically, Fig. 15 illustrates the detection power under a maximum detection delay constraint $\mathbb{P}[S_t - \nu \leq M_{\max}]$ as a function of

Fig. 14. PFA as a function of the detection threshold τ for different set of parameters for the detector, window size N , polynomial degree q and number of corrupted samples $\|\mathbf{v}_a\|$.

the average RL2FA, see (28). The maximal detection delay is set to $M_{\max} = 10$ seconds. On the other hand, Fig. 16 depicts the average detection delay, i.e., $\mathbb{E}[S_t - \nu | S_t \geq \nu]$, as a function of the average RL2FA.

The comparison is made on both figures for two different attack payloads, 43 and 50 bad requests/s. For the CUSUM test, the constant κ is set to 0.005 which corresponds roughly to the 1% largest values of detection statistics $\tilde{\mathbf{v}}_a^T \tilde{\mathbf{r}}$ under $\mathcal{H}_0$ (see Fig. 7). For the snapshot test, the parameter detection delay M is set to 10s. To be comprehensive, these figures also offer a comparison with the proposed detector in [24] by replacing the proposed GLRT statistics $\mathbf{v}_a^T \tilde{\mathbf{r}}_t$ with the test [24] in the CUSUM equation (29).

Fig. 15 and 16 clearly show that the gain obtained by gathering consecutive samples is huge. Indeed, for an average RL2FA of 300 seconds (5 minutes), one can note that the probability of detecting an IFA after 10 seconds increased from roughly 5% for the snapshot test to more than 90% using the CUSUM procedure. Similarly, Fig. 16 shows that, for the same average RL2FA, the average detection delay is decreased by a factor of about 8 from about 40 seconds for the snapshot test to 5 seconds for the CUSUM test.

Fig. 15. Power of sequential detection method (probability of detection with maximal constraint delay) as a function of average RL2FA.

Fig. 16. Average Detection Delay as a function of average RL2FA.

VI. CONCLUSION

In this paper, we address the problem of IFA in the coupled NDN-IP network. As compared to existing studies and given the maturity level of the NDN paradigm, the paper focuses on a credible deployment scenario in which IP and NDN coexist in isolated domains interconnected by dedicated gateways. In such a context, we first demonstrate that despite the implementation of the *NACK* packet, IFA is still possible. Moreover, we propose a practical attack scenario leveraging HTTP traffic. Results have shown that the proposed IFA scenario succeeds in degrading the user's experience (i.e., web loading delay). To tackle this threat, we first design a GLRT detector and evaluate its intrinsic performance in a simulation environment. The results show the relevance of the proposed model with the close match of empirical and theoretical results, and also the ability to guarantee a prescribed PFA and to establish the trade-off between detection power and delay. To address the IFA in a real context of NDN network coupled with IP to carry web traffic, the detector is extended to increase its accuracy by developing a sequential version based on the initial snapshot GLRT. The results demonstrate the good capability of our approach to operate in a real context and the significant gain obtained by the sequential detector, regarding the average detection delay and the probability of true alarm with the constraint of maximum detection delay.

As a prospective future work, we have started addressing other types of attacks in NDN (e.g., content poisoning [52]) by proposing a monitoring plane that can consider a wide range of metrics to detect anomalous behavior using the approaches presented here. First promising results have been obtained in this direction [53], [54]. Other perspectives are to study the scalability of the proposed detection in large networks and to design a collaborative detection method which gathers information from several routers to enhance the detection accuracy.

ACKNOWLEDGEMENTS

This work has been funded in part by the French National Research Agency (ANR), DOCTOR project <ANR-14-CE28-000>, by the French Systematic ICT cluster and by the CRCA and FEDER Transdisciplinary CyberSec Platform <D201304601> ; www.cybersec.utt.fr.

APPENDIX A

ASYMPTOTIC OPTIMALITY OF PROPOSED TEST $\delta^*(r_t)$

This appendix proves that the proposed test $\delta^*(r_t)$, defined in equation (13) is Asymptotically Uniformly Most Powerful (AUMP) for the testing problem (12) into two steps. First it shows that the test $\delta^*(r_t)$ is UMP for the testing problem (12) and, then a proof that this test is AUMP is given.

Let us recall that from Eqs (8)–(11), the testing problem (12) is defined by the following hypotheses:

$$r_t \sim \begin{cases} \mathcal{N}(0, \sigma_t^2) & \text{under } \mathcal{H}_0, \\ \mathcal{N}(a, \sigma_t^2 - \sigma_a^2) & \text{under } \mathcal{H}_1. \end{cases}$$

with $\sigma_t^2 = \frac{p_t(1-p_t)}{i_t}$, $\sigma_a^2 = \frac{ap_t}{i_t}$ and $\frac{(1-p_t)N_a}{i_t + N_a} = a$. The Likelihood Ratio $\Lambda(x)$ is given by:

$$\Lambda(x) = \frac{\sqrt{\sigma_t^2}}{\sqrt{\sigma_t^2 - \sigma_a^2}} \exp\left(\frac{x^2}{2\sigma_t^2} - \frac{(x-a)^2}{2\sigma_t^2 - 2\sigma_a^2}\right). \quad (30)$$

It immediately follows that :

$$\frac{\partial \Lambda(x)}{\partial x} = \left(\frac{x}{\sigma_t^2} - \frac{x-a}{\sigma_t^2 - \sigma_a^2}\right) \Lambda(x). \quad (31)$$

Therefore, searching the values on x for which the first term of Eq.(31) is positive, one finds straightforwardly:

$$x < \frac{a\sigma_t^2}{\sigma_a^2} \quad (32)$$

Replacing σ_t , σ_a , and a by their definition in Eq.(32), a short algebra gives:

$$x < 1 - p_t \quad (33)$$

Because $r_t = (1 - d_t/i_t) - p_t$, see Eq.(9), Eq.(33), one eventually has :

$$1 - \frac{d_t}{i_t} - p_t < (1 - p_t) \Rightarrow -\frac{d_t}{i_t} < 0 \quad (34)$$

which is always true since i_t and d_t respectively represent the received number of *Interest* and *Data* packets respectively and therefore are positive numbers. We note that in the case where $i_t = 0 = d_t$, the LR $\Lambda(r_t)$ is “degenerated” because $\sigma_t = 0$ and $\sigma_a = \infty$. It thus follows that the hypotheses $\mathcal{H}_0$ and $\mathcal{H}_1$ (12) admit a monotone LR with respect to r_t and therefore, it follows from the Karlin–Rubin Theorem [35, Theorem 3.4.1] that the test $\delta^*(r_t)$, see (13), is UMP.

According to the definition of convergence in distribution, see [35, Definition 11.2.1] and in virtue of both the Portmanteau [35, Theorem 11.2.1] and the continuous mapping theorem [35, Theorem 11.2.13] it follows that the power function of the test $\delta^*(r_t)$ converges to the power function of the MP test $\delta^*(d_t)$, defined in (7), and therefore that the proposed test $\delta^*(r_t)$ is Asymptotically UMP for the testing problem (4).

□

REFERENCES

- [1] G. M. De Brito, P. B. Velloso, and I. M. Moraes, *Information Centric Networks: A New Paradigm for the Internet*. John Wiley & Sons, 2013.
- [2] G. Xylomenos, C. N. Ververidis, V. Siris, N. Fotiou, C. Tsilopoulos, X. Vasilakos, K. V. Katsaros, G. C. Polyzos *et al.*, “A survey of information-centric networking research,” *Communications Surveys & Tutorials, IEEE*, vol. 16, no. 2, pp. 1024–1049, 2014.
- [3] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, “Networking named content,” in *Proceedings of the 5th international conference on Emerging networking experiments and technologies*. ACM, 2009, pp. 1–12.
- [4] L. Zhang, A. Afanasyev, J. Burke, V. Jacobson, P. Crowley, C. Papadopoulos, L. Wang, B. Zhang *et al.* “Named data networking,” *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 3, pp. 66–73, 2014.
- [5] T. Lauinger, “Security & scalability of content-centric networking,” M.S. Thesis Dissertation, TU Darmstadt, 2010.
- [6] N. T. Nguyen, R. Cogranne, and G. Doyen, “An optimal statistical test for robust detection against interest flooding attacks in ccn,” in *IFIP/IEEE Intl’ Symposium on Integrated Network Management (IM)*, 2015.
- [7] T. N. Nguyen, R. Cogranne, G. Doyen, and F. Retraint, “Detection of interest flooding attacks in named data networking using hypothesis testing,” in *Intl’ Workshop on Information Forensics and Security (WIFS)*, IEEE, Nov 2015, pp. 1–6.
- [8] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica, “A data-oriented (and beyond) network architecture,” in *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 4. ACM, 2007, pp. 181–192.
- [9] N. Fotiou, P. Nikander, D. Trossen and G.C. Polyzos “Developing Information Networking Further: From PSIRP to PURSUIT,” in *Broadnets*, pages 1–13. Springer, 2010.
- [10] C. Dannewitz, D. Kutscher, B. Ohlman, S. Farrell, B. Ahlgren, and H. Karl, “Network of information (netinf)—an information-centric network architecture,” *Computer Communications*, vol. 36, no. 7, pp. 721–735, 2013.
- [11] A. Afanasyev, I. Moiseenko, L. Zhang *et al.*, “ndnsim: Ndn simulator for NS-3,” *University of California, Los Angeles, Tech. Rep*, 2012.
- [12] X. Marchal, T. Cholez, and O. Festor, “Pit matching from unregistered remote faces: A critical ndn vulnerability,” in *Proc. of the 3rd ACM Conference on Information-Centric Networking*, ACM-ICN’16, pp. 211–212, 2016.
- [13] D. Kim, J. Bi, A. V. Vasilakos, and I. Yeom, “Security of cached content in ndn,” *IEEE Transactions on Information Forensics and Security*, vol. 12, no. 12, pp. 2933–2944, Dec 2017.
- [14] Z. Zhang, Y. Yu, H. Zhang, E. Newberry, S. Mastorakis, Y. Li, A. Afanasyev, and L. Zhang, “An overview of security support in named data networking,” Technical Report NDN-0057, NDN, Tech. Rep., 2018.
- [15] Z. Zhang, H. Zhang, E. Newberry, S. Mastorakis, Y. Li, A. Afanasyev, and L. Zhang, “Security support in named data networking,” Technical Report. Available online: <https://named-data.net/wp-content/uploads/2018/03/ndn-0057-1-ndn-security.pdf> (accessed on 18 March 2018), Tech. Rep., 2018.
- [16] G. Acs, M. Conti, P. Gasti, C. Ghali, G. Tsudik, and C. Wood, “Privacy-Aware Caching in Information-Centric Networking,” *IEEE Transactions on Dependable and Secure Computing*, (IEEE early access, to be published, doi:10.1109/TDSC.2017.2679711), 2018.
- [17] S. Rai and D. Dhakal, “A survey on detection and mitigation of interest flooding attack in named data networking,” in *Advanced Computational and Communication Paradigms*. Singapore: Springer Singapore, 2018, pp. 523–531.
- [18] R. Tourani, S. Misra, T. Mick, and G. Panwar, “Security, privacy, and access control in information-centric networking: A survey,” *IEEE Communications Surveys Tutorials*, vol. 20, no. 1, pp. 566–600, Firstquarter 2018.
- [19] T. Chatterjee, S. Ruj, and S. D. Bit, “Security issues in named data networks,” *Computer, IEEE*, vol. 51, no. 1, pp. 66–75, January 2018.
- [20] P. Gasti, G. Tsudik, E. Uzun, and L. Zhang, “Dos and ddos in named data networking,” in *Computer Communications and Networks (ICCCN), 22nd Intl’ Conf. on. IEEE*, 2013, pp. 1–7.
- [21] H. Dai, Y. Wang, J. Fan, and B. Liu, “Mitigate ddos attacks in ndn by interest traceback,” in *Conf. on Computer Communications Workshops (INFOCOM WKSHPS)*, IEEE, 2013, pp. 381–386.
- [22] J. Tang, Z. Zhang, Y. Liu, and H. Zhang, “Identifying interest flooding in named data networking,” in *Proc. Joint conf. on Green Computing and*

- Communications (GreenCom), Internet of Things (iThings) and Cyber, Physical and Social Computing (CPSCom)*. IEEE, 2013, pp. 306–310.
- [23] A. Compagno, M. Conti, P. Gasti, and G. Tsudik, “Poseidon: Mitigating interest flooding ddos attacks in named data networking,” in *Local Computer Networks (LCN), Intl’ Conf. on*. IEEE, 2013, pp. 630–638.
- [24] A. Afanasyev, P. Mahadevan, I. Moiseenko, E. Uzun, and L. Zhang, “Interest flooding attack and countermeasures in named data networking,” in *IFIP Networking Conference*. IEEE, 2013, pp. 1–9.
- [25] C. Ghali, G. Tsudik, E. Uzun, and C. A. Wood, “Closing the floodgate with stateless content-centric networking,” in *2017 26th International Conference on Computer Communication and Networks (ICCCN)*, July 2017, pp. 1–10.
- [26] T. Zhi, H. Luo, and Y. Liu, “A gini impurity-based interest flooding attack defence mechanism in ndn,” *IEEE Communications Letters*, vol. 22, no. 3, pp. 538–541, March 2018.
- [27] Y. Xin, Y. Li, W. Wang, W. Li, and X. Chen, “Detection of collusive interest flooding attacks in named data networking using wavelet analysis,” in *MILCOM 2017 - 2017 IEEE Military Communications Conference (MILCOM)*, Oct 2017, pp. 557–562.
- [28] C. Yi, A. Afanasyev, I. Moiseenko, L. Wang, B. Zhang, and L. Zhang, “A case for stateful forwarding plane,” *Computer Communications*, vol. 36, no. 7, pp. 779–791, 2013.
- [29] K. Wang, Y. Zhaon, S. Liu and X. Tong, “On the urgency of implementing Interest NACK into CCN: from the perspective of countering advanced interest flooding attacks” in *IET Networks*, vol. 7, no. 3, pp. 136–140, 2018.
- [30] S. DiBenedetto and C. Papadopoulos, “Mitigating poisoned content with forwarding strategy,” in *2016 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, April 2016, pp. 164–169.
- [31] H. Guo, X. Wang, K. Chang, and Y. Tian, “Exploiting path diversity for thwarting pollution attacks in named data networking,” *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 9, pp. 2077–2090, Sept 2016.
- [32] H. L. Mai, N. T. Nguyen, G. Doyen, A. Ploix, and R. Cograanne, “On the readiness of ndn for a secure deployment: The case of pending interest table,” in *IFIP Intl’ Conf. on Autonomous Infrastructure, Management and Security*. Springer, 2016, pp. 98–110.
- [33] M. Virgilio, G. Marchetto and R. Sisto, “PIT overload analysis in content centric networks,” in *Proceedings of the 3rd ACM SIGCOMM workshop on Information-centric networking*, pages 67–72. ACM, 2013.
- [34] X. Marchal, M. El Aoun, B. Mathieu, T. Cholez, G. Doyen, W. Mallouli and O. Festor, “Leveraging NFV for the deployment of NDN: Application to HTTP traffic transport” in *Proc. IEEE/IFIP Network Operations and Management Symposium (NOMS)*, IEEE, 2018, pp. 1–5.
- [35] E. Lehmann and J. Romano, *Testing Statistical Hypotheses, Second Edition*, 3rd ed. Springer, 2005.
- [36] J.-C. Bolot, “Characterizing end-to-end packet delay and loss in the internet,” *J. High Speed Networks*, vol. 2, no. 3, pp. 305–323, 1993.
- [37] E. Altman, K. Avrachenkov, and C. Barakat, “A stochastic model of tcp/ip with stationary random losses,” *Networking, IEEE/ACM Transactions on*, vol. 13, no. 2, pp. 356–369, 2005.
- [38] H. Yin, C. Lin, B. Sebastien, B. Li, and G. Min, “Network traffic prediction based on a new time series model,” *Intl’ Journal of Communication Systems*, vol. 18, no. 8, pp. 711–729, 2005.
- [39] R. Cograanne, G. Doyen, N. Ghabban and B. Hamni, “Detecting Bot-clouds at Large Scale: a Decentralized and Robust Detection Method for Multi-Tenant Virtualized Environments,” *Network Service Management, IEEE Transactions on*, Special Issue on Advances in Big Data Analytics for Management, vol. 15, no. 1, pp. 68–82, January 2018.
- [40] R. Cograanne and F. Retraint, “An asymptotically uniformly most powerful test for LSB matching detection,” *Information Forensics and Security, IEEE Transactions on*, vol. 8, no. 3, pp. 464–476, March 2013.
- [41] R. Cograanne and F. Retraint, “Detection of defects in radiographic images using an adaptive parametric model,” *Signal Processing*, vol. 96, Part B, pp. 173–189, March 2014.
- [42] A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential Analysis: Hypothesis Testing and Changepoint Detection*. Taylor & Francis, 2014.
- [43] H. Poor and O. Hadjiladis, *Quickest Detection*. Cambridge University Press, 2008.
- [44] B. K. Gupié, L. Fillatre, and I. Nikiforov, “Sequential detection of transient changes,” *Sequential Analysis*, vol. 31, no. 4, pp. 528–547, 2012.
- [45] R. Cograanne, “A Sequential Method for Online Steganalysis,” in *Proc. IEEE Workshop on Information Forensics and Security (WIFS)*, 2015.
- [46] E. Page, “Continuous inspection schemes,” *Biometrika*, pp. 100–115, 1954.
- [47] G. Lorden, “Procedures for reacting to a change in distribution,” *The Annals of Mathematical Statistics*, pp. 1897–1908, 1971.
- [48] S. Basu, A. Mukherjee, and S. Klivansky, “Time series models for internet traffic,” in *15th Intl’ Conf. on Computer Communications (INFOCOM)*, vol. 2. IEEE, 1996, pp. 611–620.
- [49] L. Breslau & al., “Web caching and zipf-like distributions: Evidence and implications,” in *18th Intl’ Conf. on Computer Communications (INFOCOM)*, vol. 1. IEEE, 1999, pp. 126–134.
- [50] V. S. Frost and B. Melamed, “Traffic modeling for telecommunications networks,” *Communications Magazine, IEEE*, vol. 32, no. 3, pp. 70–81, 1994.
- [51] E. Parzen, “On estimation of a probability density function and mode,” *The annals of mathematical statistics*, vol. 33, no. 3, pp. 1065–1076, 1962.
- [52] T. Nguyen, X. Marchal, G. Doyen, T. Cholez and R. Cograanne, “Content poisoning in named data networking: Comprehensive characterization of real deployment,” in *IFIP/IEEE Intl’ Symposium on Integrated Network Management (IM)*, pp. 72–80, 2017.
- [53] H.L. Mai, N.T. Nguyen, G. Doyen, R. Cograanne, W. Mallouli, E. Montes de Oca and O. Festor, “Towards a Security Monitoring Plane for Named Data Networking: Application to Content Poisoning Attack,” *Proc. IEEE/IFIP Network Operations and Management Symposium (NOMS)*, IEEE, 2018, pp. 1–9.
- [54] T. Nguyen, H.L. Mai, G. Doyen, R. Cograanne, W. Mallouli, E. Montes-de-oca and O. Festor, “A Security Monitoring Plane for Named Data Networking Deployment,” *IEEE Communications Magazine (ComMag), Feature topic on Information-Centric Networking Security*, vol. 56, no. 11, pp. 88–94, November 2018.

Tan Nguyen received the PhD from Troyes University of Technology (UTT), France, in July 2018. His PhD was co-supervised by Dr. Rmi COGRANNE and Dr. Guillaume DOYEN. His research area focuses on security issues in Information Centric Networks and especially the NDN proposal. His PhD takes part of the DOCTOR project, started in December 2014 and funded by the French National Agency of Research (ANR).

Hoang-Long Mai received his master in Information Systems Security from Troyes University of Technology in 2016. He is currently a Ph.D. student in a CIFRE (Industrial Convention of Formation by Research) contract between Montimage, Troyes University of Technology and INRIA Lorraine. His Ph.D. topic focuses on the Autonomous Monitoring and Control of Virtualized Network Functions with an application to Named Data Networking.

Rémi Cograanne is an Associate Professor at Troyes University of Technology (UTT), France, since 2013. He has regularly been a visiting scholar at Binghamton University between 2014 and 2017. He received his PhD in Systems Safety and Optimization from UTT in 2011, since on, his research focus on hypothesis testing applied to image forensics, steganalysis, steganography and computer network anomaly detection which lead to more than 55 papers and 3 International patents.

Guillaume Doyen is an Associate Professor at Troyes University of Technology, France, since 2006. His current research focuses on the design of autonomous management solutions for the performance and security of content distribution and virtualized infrastructures. He published more than 50 papers in the network and service management community. He is a TPC member of high-venue conferences (IFIP/IEEE CNSM, NOMS, IM) and a co-chair of several events (AIMS, ManSDN/NFV).

Olivier Festor is Professor in Computer Science at the University of Lorraine and Director of the TELECOM Nancy, the Graduate Engineering School in Computer Science. Chair of IFIP TC6 WG 6.6 and IEEE COMSOC member, he is active for more than 25 years in the Network and Service Management scientific community. His research interest are in Network Security Monitoring and Configuration.

Wissam Mallouli is currently a research and development project manager at Montimage, France. He received his PhD in computer science from Telecom and Management SudParis (France) in 2008. His topics of interest cover formal methods for monitoring of functional, performance and security aspects of networks and applications. He is working in several European and French research projects. He also participates to the program/organizing committees of numerous national and international conferences.

Luong Nguyen received his master in Software Engineering for Ambient Intelligence from Telecom SudParis in 2014. He is currently a research and development engineer at Montimage, France. His research area focuses on network monitoring and especially deep packet inspection.

Moustapha El Aoun TBD.

Edgardo Montes de Oca graduated as a Computer and Electronics engineer 1985 from Paris XI, Orsay and DEA in Computers from Paris VI, Jussieu 1986. He was research engineer and leader in Euriware, and Alcatel's and Ericsson's Research centres. In 2004 he founded Montimage, a research oriented SME. His main interests include monitoring the security and performance of 4G/5G networks. He has published more than 30 papers, book chapters and patents.