



HAL
open science

Advanced Interest Flooding Attacks in Named-Data Networking

Salvatore Signorello, Samuel Marchal, Jerome Francois, Olivier Festor, Radu State

► **To cite this version:**

Salvatore Signorello, Samuel Marchal, Jerome Francois, Olivier Festor, Radu State. Advanced Interest Flooding Attacks in Named-Data Networking. NCA 2017 - IEEE International Symposium on Network Computing and Applications, Oct 2017, Cambridge, United States. hal-01636494

HAL Id: hal-01636494

<https://inria.hal.science/hal-01636494>

Submitted on 16 Nov 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Advanced Interest Flooding Attacks in Named-Data Networking

Salvatore Signorello*, Samuel Marchal[†], Jérôme François[‡], Olivier Festor^{‡§} and Radu State*

*SnT, University of Luxembourg, email: firstname.lastname@uni.lu

[†]Aalto University, Finland, email: samuel.marchal@aalto.fi

[‡]Inria Nancy Grand Est, France, email: firstname.lastname@inria.fr

[§]Telecom Nancy, Université de Lorraine, France

Abstract—The Named-Data Networking (NDN) has emerged as a clean-slate Internet proposal on the wave of Information-Centric Networking. Although the NDN’s data-plane seems to offer many advantages, e.g., native support for multicast communications and flow balance, it also makes the network infrastructure vulnerable to a specific DDoS attack, the Interest Flooding Attack (IFA). In IFAs, a botnet issuing unsatisfiable content requests can be set up effortlessly to exhaust routers’ resources and cause a severe performance drop to legitimate users. So far several countermeasures have addressed this security threat, however, their efficacy was proved by means of simplistic assumptions on the attack model. Therefore, we propose a more complete attack model and design an advanced IFA. We show the efficiency of our novel attack scheme by extensively assessing some of the state-of-the-art countermeasures. Further, we release the software to perform this attack as open source tool to help design future more robust defense mechanisms.

I. INTRODUCTION

With regard to the design of future Internet architectures, the Information-Centric Networking (ICN) [1] is a clean-slate approach which strives for leveraging content-awareness at the layer-3. To that purpose, ICN architectures usually feature named self-secured content objects, in-network storage and anycast name-based routing. So far, several ICN architectures have emerged. Among those, the Named-Data Networking (NDN) [2] has been widely adopted by researchers worldwide as a platform for experimentation on ICN concepts.

The NDN protocol is based on the asynchronous exchange of two different network packets for content requests and content themselves, namely, Interests and Data. Both NDN packets carry human-readable names, made of globally-routable prefixes and provider-specific content identifiers, which are used by a stateful forwarding plane in routers. In fact, processed Interests are stored by their name in routers’ tables and kept pending for homonymous Data packets to be routed on backward paths. On one hand, keeping such status in the network offers advantages like loop-free multipath content retrieval, flow balance, prompt recovery from sudden network problems, etc. [3]. On the other hand, it exposes the network to Interest Flooding Attacks (IFAs), that is, attacks issuing a huge quantity of unsatisfiable requests with the aim to exhaust network’s and content providers’ resources [4], [5], [6]. The severity of IFAs for NDN-like networks and the questionable efficacy of the proposed solutions threaten the stateful forwarding plane existence in ICNs [7].

While countermeasures have been proposed to mitigate IFAs, simplistic assumptions in attack scenarios undermine their potential. First, all of them are *reactive*, i.e., the mitigation of the attack inevitably starts after a time interval in which routers collect statistics about the processed traffic. This makes these countermeasures vulnerable to attackers which may adjust their Interest generation to influence collected statistics or exploit routers’ monitoring time windows. Second, the evaluation scenarios for those countermeasures often assume malicious and legitimate Interests belong to *trivial disjoint prefix sets* and/or use *randomly-generated content identifiers* to generate requests for non-existent contents. These assumptions make it easy to drop malicious traffic once a certain prefix has been detected as infected.

In contrast, we believe that stealthy attack methods would mimic real content names and vary the attack characteristics over time in order to remain undetected by routers. Therefore, we have designed an advanced IFA and reassessed the most effective state-of-the-art countermeasures. Our evaluation shows that those countermeasures fail to cope with the novel attack model and that different countermeasures need to be designed as future works.

To summarize, the paper’s contributions are the following:

- an advanced attack model for Interest Flooding Attacks in NDN-like networks (Sec. III-B).
- an analysis of the pitfalls of the state-of-the-art countermeasures against IFAs and a list of good practices for the design of future ones (Sec. IV).
- a novel IFA which leverages the attack model presented in III-B and exploits the pitfalls of existing countermeasures outlined in IV (Sec. V). The source code and the dataset to generate this IFA have been made public for researchers to reproduce the results reported hereby and assess newly-designed countermeasures.
- a re-evaluation of the three *collaborative* countermeasures proposed in [4], [8], [9] which highlights the efficiency of our novel attack model (Sec. VI).

II. NAMED-DATA NETWORKING

Named-Data Networking [2] is an open source ICN instance. The NDN layer-3 protocol processes two different packets: Interest and Data. An Interest packet is meant to carry a request for a specific content; while, a Data packet is

meant to carry the content itself and a digital signature. Both packets have URI-like variable length names, so an Interest and a Data for a copy of this article could have a name like: *ndn:/en.wikipedia.org/wiki/technology/media/IFAs_in_NDN*.

An NDN content name is usually made of two main components. The first part, e.g., "en.wikipedia.org", is usually referred to as prefix name. A prefix name constitutes a globally routable identifier which is looked up in router's tables to move the packet across the network. The second part of the name, e.g., "wiki/technology/media/IFAs_in_NDN", is usually referred to as content identifier and it refers to a specific resource provided by a certain provider. A part of the name which is included between two delimiters, e.g., "wiki" and "media", is called a name component or a name segment. Both prefix names and content identifiers may include one or more name segments.

When an Interest is received at an NDN-enabled router, the three following look-up operations are performed at worst. First, the router checks a local storage, called Content Store (CS), for a copy of the requested content. Next, if there are no local copies, a second look-up is performed to a table called Pending Interest Table (PIT). The PIT contains records of requests processed, yet not expired, lingering in the router. If the PIT holds a record with the same name, then that record is updated when the Interest comes from a previously-unseen incoming network interface, otherwise the Interest is dropped. When no PIT records exist for that Interest name, a further look-up is done on a Forwarding Information Base (FIB) table to forward the Interest to a next hop. The FIB is a name-based forwarding table providing output interfaces per prefix name. The entry that has more name components in common with the Interest name is selected over all the existing FIB entries. The router forwards the Interest out of the interface provided by the looked-up FIB entry; lastly, it creates a new PIT entry to record the forwarded Interest. FIB entries are managed by routing protocols in the control plane, while PIT entries last until related Data packets consume them or expire after a certain time.

III. INTEREST FLOODING ATTACK IN NDN

The NDN data-plane makes the architecture resilient to many of the DoS attacks affecting the current Internet [10], [3]. Nevertheless, the presence of per-Interest state in PIT introduces some new critical vulnerabilities. In fact, since PIT modifications are packet driven, tailor-made packets and packet generation frequencies may be misused to overload routers with unnecessary processing or exhaust their resources, e.g., tables memory or CPU cycles. As consequence, regular traffic will suffer from a degradation of performance at best or from a denial of service at worst. Preliminary evidence of the vulnerability, as well as of the possibility to easily exploit it, was shown in [5]. Among several vulnerabilities identified therein, the authors showed that resource exhaustion on routers can be achieved by issuing a large number of Interests for non-existent content. However, the term *Interest flooding* was coined for the first time in [11] to identify an attack mirroring




Fig. 1: Illustration of an IFA targeting a Wikipedia content provider on an NDN network. Malicious users produce closely-spaced Interests for non-existent contents (alias Fake Interests) which persist on routers' PITs until expiring. Affected routers are across the path to the content provider. Thus, as a result, Legitimate Interests (LIs) may be frequently dropped by some router whose PIT is full (e.g., routers R4 and R5).

traditional DoS attacks by sending large numbers of Interests hard to aggregate and be served by caches. A more precise definition was later given in [6], where the attack goal to overflow routers' PIT is also clearly stated. Finally, the first seminal proof that IFAs can severely disrupt network services was proven empirically in [12] where a modest number of attacking nodes sharply decreased the throughput delivered to legitimate consumers.

A. General characteristics and IFA types

It is important to distinguish the Interest purpose to characterize IFAs. Hence, we resume the preliminary tentative taxonomy presented in [10] to classify unambiguously Interest types and respective IFAs.

Interests whose aim is to fetch Data for a legitimate purpose are commonly referred to Legitimate Interests (LIs). While, Interests aiming to achieve network or producers service degradation are commonly referred to Malicious Interests (MIs). A malicious Interest can be satisfiable, when it refers to an existent content or to a dynamically generated content, or fake, when it refers to a non-existent content (FI). FIs can be easily created by trailing randomly generated strings to valid routable prefix names, e.g., *ndn:/en.wikipedia.org/wiki/technology/media/512h3jh10u*. A valid routable prefix name, e.g. "en.wikipedia.org", allows FIs to create PIT entries in routers across the paths to content sources in the network. Further, assuming there is no specific eviction policy in the PIT, FIs maximize related PIT entries lifetime, since those cannot be consumed by any Data packet. An IFA consists of the generation of a large number of closely time-spaced MIs targeting one or few prefix names. The attack rate (i.e., the number of MIs per second) is important since MIs saturate router's PIT if their frequency is greater

TABLE I: Summary of the IFA impact on different targets and of the Interest types used for each purpose.

Attack target	Attack effects	Interest Type
Content providers	overwhelmed with requests demanding computation. Some Interests may require more computation from the provider, e.g., for the signature generation [14], while FIs can be filtered out by lightweight mechanisms.	MI
Routers	suffering from resources exhaustion, especially of PIT space, which causes further Interests to be dropped	MI
Network infrastructure	links get saturated due to a certain asymmetry in size between Interests and Data packets.	$MI - FI = \{x \in MI : x \notin FI\}$

than the one at which the Interests are erased from the PIT, either consumed by Data or expired. So far, two main attack variants have been reported in the related scientific literature. First, an IFA can be launched by a weakly-coordinated botnet of consumers issuing FIs. This attack, which is illustrated in Fig. 1, has so far received more attention by other works in literature. Second, an IFA can be performed by a botnet of collusive consumers and producers. The consumers issue malicious yet satisfiable Interests, while collusive producers delay at maximum the delivery of the corresponding Data packets [13].

B. Attacker Model

This work introduces a robust attack model whose attacker capabilities and goals are described in the following. Furthermore, Table I summarizes the attack impact on different targets and the Interest type that can be used to achieve it.

Attacker capabilities Interests are regular requests expressed by users without any specific privilege. Hence, controlled hosts are sufficient to launch an IFA, since no control over the network infrastructure is required. However, according to the previous literature [9], a botnet of infected end-devices emitting FIs can maximize the attack efficacy.

In general, attackers have the following capabilities:

- ability to produce Interests which both do not correspond to any existent content and correspond to existent content,
- some knowledge about countermeasures put in place by routers,
- ability to affect multiple targets at the same time by means of certain Interest types and different prefix names,
- ability to influence monitoring statistics collected by routers.

Attacker goals IFAs may be targeting either content providers or the network infrastructure or both. So far it has appeared that they can achieve the following effects:

- the exhaustion of routers' resources. Especially, they tend to fully occupy the PIT space causing further Interests to be dropped.
- the overload of requests on target providers. In order to be more efficient, such requests should ask for existent contents since those require more computation from the provider, e.g., signatures generation. Instead, FIs can be more easily filtered out by content producers.
- the saturation of network links due to a certain asymmetry in size between Interests and Data packets. Although this requires colluding parties performing the attack or asking for many different existent contents, flooding with Interests upstream links may cause Data packets received downstream to saturate the bandwidth.

IV. IFA COUNTERMEASURES & GOOD PRACTICES

Since the identification of the IFA, many researchers have investigated its effectiveness and proposed some countermeasures. Usually, countermeasures against IFAs feature a two-phase, detection and mitigation, mechanism. The detection phase aims to identify the attack source (generally, a specific interface) and/or target prefix names. While the reaction, which is triggered by a successful identification, tries to either stop the attack or reduce the attack's impact.

This section presents the state-of-the-art IFA countermeasures with a focus respectively on *detection techniques* in IV-A, *mitigation mechanisms* in IV-B, *evaluation settings* in IV-C. Finally, Sec IV-D outlines pitfalls of the presented techniques with regard to those *three aspects*.

A. IFA detection

Both percentage of expired Interests and PIT usage may be periodically observed to detect IFAs. On one hand, routers know whether forwarded Interests are satisfied by Data or expired after a timeout because of the strict flow balance between Interest and Data packets in NDN (One Interest is consumed by only one Data). Therefore, many expired Interests along a certain time window may reveal the presence of anomalies. For example, the Interest Satisfaction Ratio (ISR), which is defined as the ratio between expired Interests over satisfied ones, is used in [4]. Similarly, the ratio of incoming Interests over outgoing Data is monitored to detect an IFA in [15]. More simplistically, counters on expired Interests can be set to obtain similar indicators [16], [17]. On the other hand, the PIT usage can always provide an indication of the load managed by a router at a certain time. Thus, for example, PIT size in bytes is monitored in [18], [9], while, PIT utilization rate over the entire PIT or per name-prefix are monitored respectively in [19] and in [8], [13].

Overall, both ISR-oriented and PIT-oriented are good indicators of anomalies in the network, yet assuming the related ratios are computed over meaningful time windows. Nevertheless, both detection metrics require thresholds to trigger alarms when exceeded. Ideally, thresholds are dynamically adjusted to reflect the network load like in [9]. In reality, most of the times threshold values are the outcome of empirical

TABLE II: Summary of the Collaborative IFA Countermeasures re-evaluated in this work

Technique	Detection	Reaction	Prefix name	Content Identifier	Protocol changes
SBP [4]	none	Probabilistic Forwarding based on ISR values advertised by neighbor routers for their interfaces	trivial disjoint sets	random integers	appending information to the name of an Interest with a local scope; queuing and scheduling mechanisms at output interfaces
DP [9]	local ISR & PIT usage values & neighbors' alerts reception	rate limit on the infected interface & dissemination of alerts to neighbors	unreported	unreported	reserved name-space for the communication between neighbor routers; processing of unsolicited Data packets
CNMR [8]	locally at router: per-prefix ISR & PIT usage values. Globally at the controller: PIT usage per-prefix name	Probabilistic Forwarding based on ISR values locally at each router	non-disjoint sets	random integers	i) placement algorithm for the selection of monitoring nodes; ii) ad-hoc routing and forwarding algorithms to maximize traffic coverage by the monitoring nodes; iii) ad-hoc messages and namespaces between the controller and the monitoring nodes; iv) custom changes to the Interest packet format

observations biased by topologies and traffic distributions used in the evaluation settings [8].

B. IFA reaction

With regard to the reaction mechanisms, different strategies can be applied locally at each router:

- rate limits of accepted Interests can be enforced on interfaces based on different parameters, e.g. expired Interests for FIB-records [16], at access routers [18], alert messages from neighbor routers [9], [4], tokenized link capacity [4].
- detection and reaction can be consolidated in a single phase. For example, the satisfaction-based mechanisms of [4] use the ISR for a prefix name or an interface as probability to forward or drop Interests.
- prefix names, detected as target of an attack, can be processed differently to prevent them from filling the PIT, as done, for example, in [17].

However, according to literature, decisions taken independently at every router suffer from two issues. First, they cause *overreaction*, which means Interests are unnecessarily processed multiple times by the same defense mechanism on all the routers along the path from a consumer to the producer (for example, in Fig. 1, the routers R3-R4-R5 along the path from the user to the wikipedia content provider may independently apply their own defense mechanism). Overreaction may also lower the probability of forwarding legitimate Interests from one source to a destination [4]. Second, a decision local to a router does not necessarily *detect or mitigate the attack globally*. In fact, the effect of the attack is expected to be stronger in proximity of producers of a targeted prefix name (see the diameter of the circles around routers R4 and R5 in Fig. 1). While, routers which are more distant from those target producers (e.g., router R1 in Fig. 1) may carry malicious traffic as well but not be able to detect it properly.

For the above reasons, collaborative defense mechanisms [18], [9], [19], [8], [4] result to be more efficient. In collaborative techniques, the dissemination of routers' local status allows defense mechanisms to be triggered even where the attack effect is not sufficiently strong to be detected (e.g., in Fig. 3

alert messages produced by router R4 may trigger mitigation at R1). Three different *collaborative* techniques are hereby briefly described, while a summary of their characteristics and differences is provided in Table II. These countermeasures are later tested (see Sec. VI) against the IFAs proposed in this paper.

Satisfaction-Based Pushback (SBP) [4] uses the ISR of an interface as probability to either accept or drop the traffic coming from that interface. The SBP also adjusts the forwarding rate on an interface according to the bandwidth limit announced by the neighbor routers. Moreover, the SBP keeps separate queues per each output interface to establish fairness among different input interfaces.

In *Distributed Poseidon* (DP) [9], routers whose PIT usage and ISR exceed their respective thresholds send alarms to their downstream neighbors. In this way, the information about an attack is propagated back as close as possible to its source. The alert propagation lowers detection thresholds in traversed routers, making routers less-affected able to detect the attack and react too.

Coordination Monitoring (CNMR) [8] features a small set of routers which monitor the forwarding of Interest packets within the same Autonomous System. In CNMR, the monitoring routers may either detect an attack locally or rely on the alarms raised by a centralized controller, which aggregates and analyzes information from all the monitoring routers. At each monitoring router, detection is triggered by PIT usage and ISR-based metrics, while reaction is based on a probabilistic forwarding driven by local ISR values. The monitoring routers avoid *overreaction* by flagging already-monitored Interests. Globally, a ratio of expired Interests is computed over all the monitoring nodes by the controller per every reported prefix name. If the overall ratio exceeds a certain threshold, then the controller alerts the monitoring routers that certain prefix names could be infected.

C. Evaluation

The name sets used for the evaluation of countermeasures in previous works are reported in Table III. Those works are classified in three different categories. Unfortunately, a large portion of previous works, which are listed in the row

”unreported”, provide no information about the content names. A further division in two subcategories is applied to the works reporting indication about content names, either in their publications or in the publicly available implementations referenced therein. On one hand, the works using separate prefix names for attackers and legitimate consumers have been listed in the subcategory ”disjoint sets”. On the other hand, the works in the subcategory ”non-disjoint” assume malicious Interests and legitimate ones may be expressed with the same prefix name.

As for the content name composition, in the vast majority of works, where reported, content names consist of two name segments, one segment for the prefix name and one for the content identifier. The prefix name is usually a short string, whether referring to some original domain name, like ”/google.com/”, or aiming to be self-explanatory, like ”/good/” to indicate that Interests with that prefix are legitimate. As regards to the content identifier, the consumer applications provided by the widely used ndnSim simulator [20] use integers.

D. Pitfalls

Some of the flaws outlined in this section inspired the design of more effective IFAs presented in Sec. V-C. Further, findings reported hereby can be considered as good practices for the design and evaluation of future countermeasures.

Detection As explained in Sec. IV-A, FIs have so far been identified by looking at the ISR and/or at PIT usage in routers, yet both detection metrics have some downsides:

- ISR-like metrics are influenced by all types of Interest (sometimes at the name-prefix granularity): legitimate Interests satisfied, legitimate Interests expired, fake Interests expired, legitimate, yet issued by the attacker, Interests satisfied. Therefore, the metric can be intentionally polluted to weaken the detection as exploited by the first IFA variant presented in Sec. V-C.
- abnormal PIT usage values may be also caused by network conditions where the load grows more easily, e.g., bursts of Interests or congestion; therefore, a large PIT usage itself does not give any indication about the kind of Interests which populate the PIT.

The above issues might be overcome by using both metrics, like done in [9], wherein the authors assume low ISR values may prevent the generation of false alarms when there is a high PIT usage due to a huge load of legitimate traffic. However, attackers can influence the ISR values and remain undetected even under heavy PIT load, as achieved, for example, by the first IFA variant of Sec. V-C.

Finally, detection metrics are observed over a certain time window and, beyond that, need to be applied to a name-prefix granularity to penalize less legitimate traffic. This implies respectively that i) reaction is performed with a certain delay, ii) specific prefix names have to be identified as infected. Attackers can exploit both conditions to stay undetected by changing target prefix names over short time intervals as done by the second attack variant presented in Sec. V-C.

TABLE III: Summary of prefix names used in literature.

Unreported	[21], [22], [9], [12], [18]	
Reported	disjoint	[4], [16], [17]
	non-disjoint	[8]

Reaction As outlined in Sec. IV-B, distributed defense mechanisms are more efficient in the detection and mitigation, yet they open up other issues:

- they often require reserved prefix names and ad-hoc data-plane modifications (see the column ”Protocol changes” of Tab. II). For example, routers in [9] exchange unsolicited Data packets or routers in [18] generate spoofed Data to trace an attack back to its closest source. Those practical requirements limit any straightforward deployment of the related solutions and impede their interoperability.
- they introduce overhead which is not always considered and reported in the evaluation of the countermeasures.
- they may counter-productively disseminate wrong information when detection metrics are purposely polluted by the attackers.

Evaluation The works in the subcategory ”disjoint sets” of Tab. III use separate prefix names for attackers and legitimate consumers. This assumption allows mitigation techniques applied at prefix-level to be extremely effective since they do not penalize any legitimate Interest. Indeed, according to the attacker model described in Sec. III-B, attackers have the ability to carefully hide their Interests in-between the legitimate ones to remain undetected by routers. Therefore, mitigation techniques dropping Interests with an infected prefix name inevitably affect some legitimate ones too.

With regard to the content identifiers, having simplistic names in a simulation environment represents a risk of underestimation of several operations done for packet processing [23]. In fact, although the size of the PIT entries, which will be eventually determined by the PIT implementation, could be independent of the Interest name size, longer content names are definitely going to influence many other processing tasks performed in routers and by content providers. Overall, at the time being the precise implications of the content name composition on the evaluation of countermeasures against IFAs stay unexplored; however, the IFA using Wikipedia page titles, which is presented in this paper, has showed to be more efficient compared to what is reported in literature.

V. FLOODING AN IMAGINARY NDN-BASED WIKIPEDIA

This section describes a novel IFA which targets large websites whose lists of contents are publicly available and are only subject to incremental changes. The attack i) is built upon the attacker model presented in Sec. III-B, and ii) exploits drawbacks of existing defense mechanisms illustrated in Sec. IV-D. The choice of a specific target website used for the evaluation of the attack impact is motivated in Sec. V-A. The main attack is presented in Sec. V-B, while two attack variants are illustrated in Sec. V-C.

A. NDN content names for Wikipedia pages

As outlined in Sec. IV-C state-of-the-art countermeasures have so far neglected the importance of the namespace used for the evaluation. In contrast, more realistic sets of content names are important for two reasons. First, those allow research to quantify better the resources needed by routers and the impact of an attack whose main purpose is to abuse them. Second, they help simulate better attack models where attackers aim to produce difficult-to-detect fake Interests (FIs).

In the attack model presented in Sec. III-B, prospective attackers can refer to publicly available content lists to create spoofed names looking very similar to valid ones. Further, attackers can precisely emit Interests for existent and non-existent contents. This latter knowledge makes the attackers able to dose the emitted Interests for specific purposes, like in the attack variants presented in Sec. V-C.

For the purpose of the evaluation of IFAs presented in this work, we have selected a specific website, i.e., the free online encyclopedia Wikipedia (however, the attack is widely applicable to any large website whose list of contents is available). A daily updated list of page names is provided by the Wikimedia Foundation [24]. All the content names used in our evaluation share the same valid routable 1-component-long prefix name, which is `"/wikipedia.org/"`. Then, legitimate Interests include one of the page titles as 1-component-long content identifier, which is a variable-length alphanumeric string; while, malicious Interests have a 1-component-long fake content identifier generated as described in Sec. V-B.

B. IFA on an NDN-based Wikipedia server

We introduce an Interest Flooding Attack targeting a specific NDN content provider and we name it *pure IFA* (pIFA). We assume NDN contents to have the same naming property of DNS names, namely, content names from a same provider expose a common semantically-related naming scheme, as shown in [25]. Our pIFA consists in modeling the composition of content names from the targeted provider using a set of known existent content names. The learned model is later used to generate new and non-existent content names following the provider naming scheme. Considering Wikipedia as target content provider, the full list of existent contents is publicly available. Wikipedia page titles are composed in the following manner: `"word1_word2_..._wordn"` with $n \in [1, x]$ and `wordi` is meaningful. We generate new content names for Wikipedia using a generator of semantically related words: DISCO [26]. Given an input word w and a number m , DISCO returns a maximum of m most related words to w . Using an existing seed content name w we generate the list l_1 of 200 most related words to w , e.g., *computing, hardware, desktop, etc.* are generated from the existing Wikipedia content name *"computer"*. We launch again the related word generation on each obtained word in l_1 with parameter $m = 200$ to produce the word list l_2 . Each obtained word in l_1 and l_2 is used as content name for a malicious Interest in our pIFA. It is worth noting that this generation process produces less than $200 + 200 \times 200 = 40,200$ words since seed names

and duplicates are removed from l_1 and l_2 and only unique words are considered. Our pIFA uses a fixed number of attacking nodes holding each a list of 5,000 seed Wikipedia content names. Each node generates off-line content names for malicious interests using the introduced generation technique. Generated existent content names can be discarded if the full list of existent contents is publicly available, as for Wikipedia. The off-line generation process produced 180,800 new content names per node on average, from which 50.3% were non-existent.

C. Attack variants


Two simple variants of the pIFA presented in Sec. V-B are proposed in this work. The *blended IFA* (bIFA) includes attackers generating Interest for both existent contents and non-existent ones. The main aim of a bIFA is to influence the detection metrics observed by routers, i.e., the ISR-like ones and the PIT usage, in order to both stay undetected and lower the probability for malicious Interests to be dropped. The *chameleonic IFA* (cIFA) includes attackers which change target prefix name after a certain time window. The goal of a cIFA is to avoid prefix name based countermeasures which are applied by routers after an observation window where certain name prefixes are marked as infected.

VI. EVALUATION


In this section we report results obtained from the evaluation of three state-of-the-art countermeasures against our novel IFA. We prove that existing IFA countermeasures are less effective than claimed in literature against our steadier attack model. We use the same simulation settings reported in Sec. 6B of [8] which are briefly summarized in Sec. VI-A. The countermeasures have been evaluated against the pIFA, the bIFA and the cIFA of Sec. V; the related results are reported respectively in Sec. VI-B, Sec. VI-C and Sec. VI-D. All the experiments have been conducted on the version 1.0 of the open-source ndnSIM [20] module to fairly compare to the state-of-the-art countermeasures which were implemented and evaluated on that version of the simulator. We have made public the ndnSIM extensions implementing the novel attacks and all the tools needed to reproduce the reported experiments at [27].

A. Set-up and Metrics

Set-up: the network topology simulated is the same as in [8], that is, the Exodus ISP (AS 3967) inferred by the Rocket-fuel project [28]. In every scenario, 25% of the client nodes are selected as attackers. Normal clients emit 100 Interest packets per second (ipps), while attackers Interest generation varies from 500 ipps to 1000 ipps. We report hereby only results for the 1000 ipps attack frequency since this frequency is representative enough for the conducted experiments (further, it allows the reader to visualize more clearly the results in each plot). Each simulation lasts nine minutes with the attack starting at 60 seconds and ending at 360 seconds. All the scenarios have been simulated 10 times and average values are



(a) Average ISR of legitimate clients



(b) Average PUR of all the routers

Fig. 2: Interest Satisfaction Ratio (ISR) in 2a and PIT Utilization Ratio (PUR) in 2b for the experiments presented in [8] and reproduced by this work.

reported each time. The curves labeled with 'ND' correspond to a baseline scenario with no defense mechanism. The curves labeled SBP, CNMR, DP correspond to scenarios with selected countermeasures, [4], [8] and [9] respectively on.

Metrics: We evaluate the impact of our IFAs on two metrics which have been widely used in the related works. First, the Interest Satisfaction Ratio (ISR) at clients, computed as ratio between satisfied Interests and expired ones over a time window, gives an indication of the quality of service perceived by end-users while the network is under attack. Second, the PIT utilization ratio (PUR) reports about the available capacity on routers to process legitimate traffic during an attack. The PUR is computed as the ratio between the number of PIT entries and the maximum PIT size in entries over a time window.

State-of-the-art (SoA) results: we first replicated the experiments presented in [8], where only CNMR and SBP were evaluated¹. Surprisingly, although the trend for both the evaluation metrics is almost consistent with what reported therein, we have got slightly different results. The ISR values for the ND and the CNMR scenarios, plotted in Fig. 2a, are 10 to 20% higher than the ones in [8]. Fig. 2a also shows that SBP and DP worsen the ISR compared to the baseline. This last behavior was only reported in [8] for higher attack frequencies. Indeed, this result was expected since SBP and DP apply rate limits of accepted interests on detected interfaces independently from the prefix names. In fact, if multiple prefix names are used in the evaluation, as done for the experiments in [8], then SBP and DP may accidentally drop many Legitimate Interests. CNMR, differently, mitigates only specific prefix names, which are detected as infected. Therefore, CNMR overperforms SBP in this evaluation scenario because the former never drops legitimate Interests with other prefix names except the one under attack.

With regard to the PUR plotted in Fig. 2b, none of the countermeasures alleviates the average PIT consumption in routers contrary to what reported in [8]. This is expected for

SBP and DP, which do not mitigate the attack efficiently as already shown in Fig. 2a (low efficiency of SBP in large network topologies with many attackers was already reported in [4]). As regards CNMR, the average global PIT usage is higher than the one measured in the ND scenario. This is due to a notable increment in the path stretch introduced by CNMR to steer every Interest towards a monitoring router.

B. Pure IFA - pIFA

In pIFAs, attackers emit Interests only for non-existent contents, while legitimate clients only ask for existent ones. Moreover, all the nodes use the same single prefix name. Both above conditions did not hold in the SoA experiments of [8]. The measured ISR values are reported in Fig. 3a (for sake of clarity, from now on the DP curves are omitted since those are very similar to the SBP ones). As for the SoA results, SBP poorly mitigates the attack, while CNMR improves the metric compared to the baseline. However, when comparing to the SoA results, Fig. 3g shows a 17% ISR decrease for both the baseline scenario and the CNMR one. This result confirms that our attack model generates more effective attacks compared to the ones used in the state-of-the-art related works. Fig. 3d also shows a 10% increment of the PUR in CNMR's monitoring routers compared to the SoA results, consistently with the ISR decrease.

C. Blended IFA - bIFA

The ISR values of the bIFA scenario are reported in Fig. 3b. In bIFAs, attackers also emit Interests for existing contents (a fixed percentage of the attack frequency per second, which we call *purity level*). Fig. 3b shows a 22% ISR decrease of the SoA results and a 5% decrease of the pIFA one for CNMR. Moreover, pIFA attackers' behavior totally neutralizes this countermeasure. In fact, the CNMR curve always lays above the baseline in Fig. 3b. The efficiency of the SBP technique is lowered too in case of bIFA, since the detection metric monitored by the routers is polluted and allows more fake interests to create entries in the PITs. Thus, it is also important to see how the efficacy of a bIFA increases with the increment of the number of legitimate Interests used by the attackers. Fig. 3i shows a decrease of the ISR for CNMR when the bIFA's

¹DP is not considered in [8], most likely because the implementation has never been made available since that work was published. We have implemented DP according to the description in [9] and made the code available at [27].

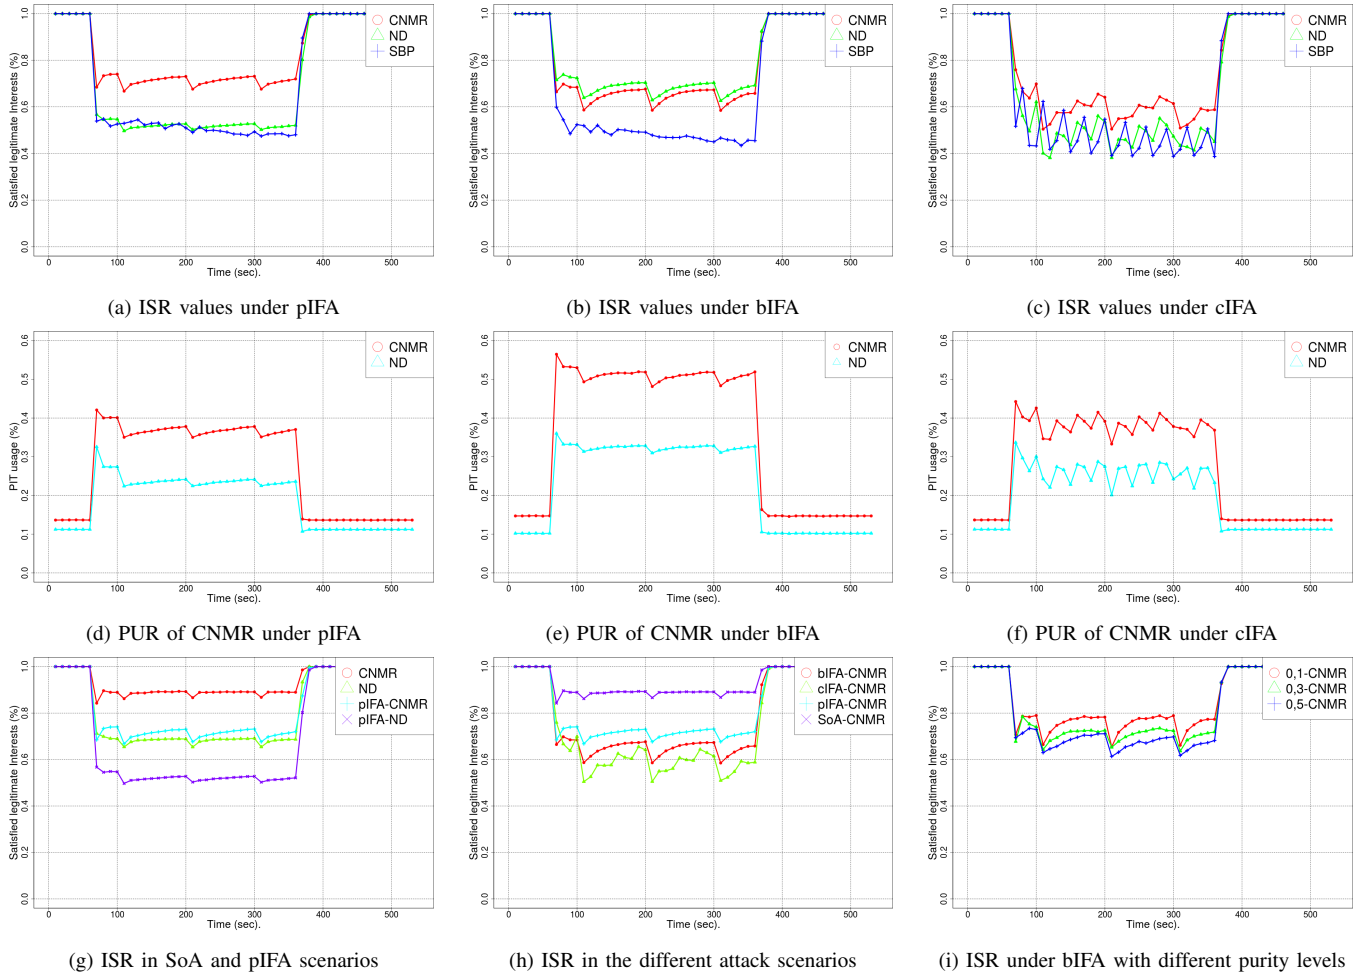


Fig. 3: On the first row from left to right, Interest Satisfaction Ratios (ISR) under pIFA, bIFA and cIFA; on the second row from left to right, PIT Utilization Ratio (PUR) in monitoring routers under pIFA, bIFA and cIFA; on the third row from left to right, ISR of CNMR in pIFA and SoA scenarios in Fig. 3g, ISR of CNMR under the three different proposed attacks in Fig. 3h, ISR of CNMR under bIFA with different purity levels in Fig. 3i.

purity level increases. The latter result is a consequence of forwarding decisions taken independently at each monitoring router. In fact, although the successful detection made by the CNMR controller node, then the monitoring routers' local statistics (see definition in Sec. IV-A) are used as probability to either forward or drop every Interest with the detected prefix name. Therefore, legitimate Interests emitted by pIFA attackers increase routers' local statistics and, by consequence, the probability for Fake Interests to occupy PIT space and be forwarded to the next hops.

In Fig. 3e for CNMR, the attack succeeds at occupying more space across monitoring routers' PITs, as a consequence of the poor detection achieved by the applied countermeasure.

D. Chameleonic IFA - cIFA

In cIFAs, the attackers either know or guess (e.g., by measuring the RTTs of issued legitimate Interests) the observation time window used by detection techniques in routers. The attackers' goal is to avoid prefix name-based mitigation techniques by switching target prefix name at every observation window (in evaluation scenarios where content

producers serve several prefix names). CNMR monitoring routers used a fixed observation time window of 10 seconds in [8]. Hence, cIFA attackers in our scenario switch target prefix name at every such time window (this justifies the periodic oscillations seen in Fig. 3c and 3f during the attack period). As expected, this attack reduces the ISR of CNMR compared to the SoA results as shown in Fig. 3c. As for the previous scenarios, the SBP still poorly detects the attack. As regard to the PUR of CNMR's monitoring routers, shown in Fig. 3f, the values are consistent with the ISR ones, showing an increase of the former metric when the latter one is reduced.

Summary of our findings: the IFAs, built upon our novel attack model, degrade the performance of the best state-of-the-art countermeasures. In particular, the attack model reduces the quality of service perceived by clients (lower ISR values mean less client requests satisfied) and increases the load on network's nodes (higher PUR values correspond to more

computational and memory resources used in routers). Fig. 3h shows how the model can be leveraged to mount increasingly strong attacks which affect the most effective state-of-the-art countermeasure. This trend holds for all the tested attack frequencies and the proposed attack variants. We recapitulate our findings in the following points:

- by using non-disjoint prefix name sets all the tested defense mechanisms drop some legitimate traffic too. In fact, none of these countermeasures distinguishes between Fake and Legitimate Interests, rather, when either a prefix name or an interface is detected as infected, all the related traffic is mitigated.
- It is possible for attackers to pollute metrics collected by defense mechanisms in routers and increase the efficacy of their attacks. This was successfully achieved by both bIFA and cIFA variants in our experiments.
- Under all the tested attacks, CNMR generates a higher load on routers compared to SBP. This means that SBP protects better the network infrastructure while CNMR better the clients (this result was not reported by the previous work in [8]).
- SBP performs aggressive mitigation on infected interfaces which inevitably affects many legitimate Interests; thus, overall, clients experience an ISR degradation. However, SBP better shields the routers, whose PUR values do not increase as sharply as for CNMR (see the last two columns of Tab. IV). This would be beneficial to legitimate traffic coming from non-infected interfaces, which would find the necessary PIT's space along the path to contents.

Table IV recapitulates the improvements achieved by the different attacks on both the evaluation metrics considered. Overall, the cIFA causes the worst ISR at clients as reported in the first two columns of Tab. IV. However, the cIFA is not the attack variant which generates the highest load on routers for all the countermeasures. In fact, the bIFA variant causes higher PUR when CNMR is on, as reported in the third column of Tab IV. Therefore, the ultimate choice between bIFA and cIFA could be driven by knowledge available to the attackers (e.g., ability to guess observation windows used by routers' defense mechanisms) and desired targets (either clients or network infrastructure or both).

TABLE IV: Summary of improvements achieved by our attacks compared to the values in the SoA results.

	CNMR-ISR	SBP/DP-ISR	CNMR-PUR	SBP/DP-PUR
pIFA	-17%	-36%	+10%	+2%
bIFA	-22%	-38%	+18%	+1%
cIFA	-28%	-40%	+11%	+3%

VII. CONCLUSION & FUTURE WORK

We have proposed a steadier attacker model for IFAs in NDN and leveraged it to design a novel attack by exploiting pitfalls of existing defense mechanisms. We have tested our attack against state-of-the-art countermeasures which fail to

detect and mitigate that properly. The source code to mount the attack has been made available to the community with the aim to have a common tool to design future more robust countermeasures. Finally, we believe that proactive countermeasures should be designed instead to be less vulnerable to the IFAs proposed in this work. Therefore, as future work, we plan to design countermeasures based on semantic analysis of the Interest names.

ACKNOWLEDGMENT

This work was undertaken under the Pollux II IDSECOM project supported by the National Research Fund Luxembourg (FNR) and the National Centre for Research and Development (NCBiR) in Poland. This work was also supported in part by the SELIoT project by the Academy of Finland under the WiFiUS program (grant 309994). Moreover, we thank the authors of [8] for sharing their implementation with us.

REFERENCES

- [1] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman, "A survey of information-centric networking," *IEEE Communications Magazine*, vol. 50, no. 7, 2012.
- [2] L. Zhang, A. Afanasyev, J. Burke, V. Jacobson, k. claffy, P. Crowley, C. Papadopoulos, L. Wang, and B. Zhang, "Named data networking," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 3, pp. 66–73, Jul. 2014.
- [3] C. Yi, A. Afanasyev, I. Moiseenko, L. Wang, B. Zhang, and L. Zhang, "A case for stateful forwarding plane," *Comput. Commun.*, vol. 36, no. 7, pp. 779–791, Apr. 2013.
- [4] A. Afanasyev, P. Mahadevan, I. Moiseenko, E. Uzun, and L. Zhang, "Interest flooding attack and countermeasures in named data networking," in *IFIP Networking Conference, 2013, Brooklyn, New York, USA, 22-24 May, 2013*, 2013, pp. 1–9.
- [5] M. Wählisch, T. C. Schmidt, and M. Vahlenkamp, "Backscatter from the data plane — threats to stability and security in Information-Centric Networking," *CoRR*, vol. abs/1205.4778, 2012.
- [6] P. Gasti, G. Tsudik, E. Uzun, and L. Zhang, "DoS and DDoS in named-data networking," *CoRR*, vol. abs/1208.0952, 2012.
- [7] C. Ghali, G. Tsudik, E. Uzun, and C. A. Wood, "Living in a pit-less world: A case against stateful forwarding in content-centric networking," *arXiv preprint arXiv:1512.07755*, 2015.
- [8] H. Salah, J. Wulfheide, and T. Strufe, "Coordination supports security: A new defence mechanism against interest flooding in NDN," in *40th IEEE Conference on Local Computer Networks*, 2015, pp. 73–81.
- [9] A. Compagno, M. Conti, P. Gasti, and G. Tsudik, "Poseidon: Mitigating interest flooding DDoS attacks in named data networking," in *Local Computer Networks (LCN), 2013 IEEE 38th Conference on*. IEEE, 2013, pp. 630–638.
- [10] P. Gasti, G. Tsudik, E. Uzun, and L. Zhang, "DoS and DDoS in named data networking," in *Proceeding of the 22nd IEEE International Conference on Computer Communications and Networks, ICCCN, 2013*, pp. 1–7.
- [11] L. Zhang, D. Estrin, J. Burke, V. Jacobson, J. Thorton, D. K. Smetters, B. Zhang, G. Tsudik, K. Claffy, D. Krioukov, D. Massey, C. Papadopoulos, T. Abdelzaher, L. Wang, P. Crowley, and E. Yeh, "Named Data Networking Tech Report 001," the NDN project team, Technical Report NDN-0001, October 2010.
- [12] A. Compagno, M. Conti, P. Gasti, and G. Tsudik, "NDN interest flooding attacks and countermeasures," in *Annual Computer Security Applications Conference*, 2012.
- [13] H. Salah and T. Strufe, "Evaluating and mitigating a collusive version of the interest flooding attack in NDN," in *Proceedings of the IEEE Symposium on Computers and Communication, ISCC-16*. Messina, Italy: IEEE Computer Society, 2016, pp. 938–945.
- [14] X. Marchal, T. Cholez, and O. Festor, "Server-side performance evaluation of NDN," in *Proceedings of the 3rd ACM Conference on Information-Centric Networking*, ser. ACM-ICN '16. New York, NY, USA: ACM, 2016, pp. 148–153.

- [15] J. Tang, Z. Zhang, Y. Liu, and H. Zhang, "Identifying interest flooding in named data networking," in *Green Computing and Communications (GreenCom), 2013 IEEE and Internet of Things (iThings/CPSCOM), IEEE International Conference on and IEEE Cyber, Physical and Social Computing*. IEEE, 2013, pp. 306–310.
- [16] K. Wang, H. Zhou, H. Luo, J. Guan, Y. Qin, and H. Zhang, "Detecting and mitigating interest flooding attacks in content-centric network," *Security and Communication Networks*, vol. 7, no. 4, pp. 685–699, 2014.
- [17] K. Wang, H. Zhou, Y. Qin, J. Chen, and H. Zhang, "Decoupling malicious interests from pending interest table to mitigate interest flooding attacks," in *Globecom Workshops (GC Wkshps), 2013 IEEE*. IEEE, 2013, pp. 963–968.
- [18] H. Dai, Y. Wang, J. Fan, and B. Liu, "Mitigate DDoS attacks in NDN by interest traceback," in *IEEE INFOCOM NOMEN Workshop, 2013*, 2013.
- [19] K. Wang, H. Zhou, Y. Qin, and H. Zhang, "Cooperative-filter: countering interest flooding attacks in named data networking," *Soft Computing*, vol. 18, no. 9, pp. 1803–1813, 2014.
- [20] A. Afanasyev, I. Moiseenko, and L. Zhang, "ndnSIM: NDN simulator for NS-3," the NDN project team, Technical Report NDN-0005, October 2012.
- [21] T. N. Nguyen, R. Cogranne, and G. Doyen, "An optimal statistical test for robust detection against interest flooding attacks in ccn," in *2015 IFIP/IEEE International Symposium on Integrated Network Management (IM)*. IEEE, 2015, pp. 252–260.
- [22] T. N. Nguyen, R. Cogranne, G. Doyen, and F. Retraint, "Detection of interest flooding attacks in named data networking using hypothesis testing," in *Information Forensics and Security (WIFS), 2015 IEEE International Workshop on*. IEEE, 2015, pp. 1–6.
- [23] H. Yuan, T. Song, and P. Crowley, "Scalable NDN forwarding: Concepts, issues and principles," in *Computer Communications and Networks (ICCCN), 2012 21st International Conference on*. IEEE, 2012, pp. 1–9.
- [24] Wikimedia downloads. <https://dumps.wikimedia.org/>. Accessed: 2017-03-04.
- [25] S. Marchal, J. François, C. Wagner, and T. Engel, "Semantic exploration of dns," in *Proceeding of the 11th International IFIP TC 6 Networking Conference*. Springer Berlin Heidelberg, 2012, pp. 370–384.
- [26] P. Kolb, "DISCO: A Multilingual Database of Distributionally Similar Words," in *KONVENS 2008 – Ergänzungsband: Textressourcen und lexikalisches Wissen*, A. Storrer, A. Geyken, A. Siebert, and K.-M. Würzner, Eds., 2008, pp. 37–44.
- [27] "Ndn-ifa signorello's github repository," <https://github.com/signorello/NDN-IFA>, accessed: 2017-08-31.
- [28] N. Spring, R. Mahajan, and D. Wetherall, "Measuring isp topologies with rocketfuel," *ACM SIGCOMM Computer Communication Review*, vol. 32, no. 4, pp. 133–145, 2002.

APPENDIX SUMMARY OF ACRONYMS

<i>Architecture</i>	
NDN	Named-Data Networking
CS	NDN's Content Store
PIT	NDN's Pending Interest Table
FIB	NDN's Forwarding Information Base
<i>Attacks</i>	
IFA	Interest Flooding Attack
pIFA	pure Interest Flooding Attack of VI-B
bIFA	blended Interest Flooding Attack of V-C
cIFA	chamaleonic Interest Flooding Attack of V-C
<i>Countermeasures</i>	
CNMR	CoordiNation MonitoRing [8]
DP	Distributed Poseidon [9]
SBP	Satisfaction-Based Pushback [4]
<i>Interest Types</i>	
LI	Legitimate Interest
MI	Malicious Interest
FI	Fake Interest
<i>Metrics</i>	
ISR	Interest Satisfaction Ratio
PUR	PIT Utilization Ratio