

Revisiting Countermeasures Against NDN Interest Flooding

Samir Al-Sheikh
Freie Universität Berlin
samir.al-sheikh@fu-berlin.de

Matthias Wählisch
Freie Universität Berlin
m.waehlich@fu-berlin.de

Thomas C. Schmidt
HAW Hamburg
t.schmidt@haw-hamburg.de

ABSTRACT

Interest flooding has been identified as a major threat for the NDN infrastructure. Since then several approaches have been proposed to identify and to mitigate this attack. In this paper, we (a) classify nine existing countermeasures and (b) compare them in a consistent evaluation setup. We discuss the application of pure prefix-based as well as pure interface-based mitigation strategies in different network scenarios.

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]: General—Security and protection; C.2.2 [Computer-Communication Networks]: Network Protocols—Routing Protocols

Keywords

ICN; NDN; Interest flooding; attack; mitigation

1. COUNTERMEASURES IN A NUTSHELL

Interest flooding describes a denial-of-service attack in which a malicious node attempts to overload the distribution infrastructure by sending Interest packets [5]. The easiest implementation is the request of non-existing content as entries need to expire until they are removed. However, even the request of existing content may harm the infrastructure when the entries in the Pending Interest Table (PIT) exceed the content delivery rate (e.g., due to large network delays) [5].

Current countermeasures try to limit the number of incoming Interests, either per prefix, per interface, or per router. The main challenge is to distinguish valid from malicious Interests. As there is no clear notion of malicious Interests, several heuristics are proposed.

Considering *all* PIT entries of the (local) router, *Token Bucket* [1] and *Resource Allocation* [2] proactively try to balance the resources. The *Interest Traceback* [3] approach does not only include the size of the PIT but also the increase of entries over time. Furthermore, if a predefined threshold is exceeded, dummy data packets are sent towards all such

consumers that are responsible for stale Interests—to release states within the network and finally limit Interests at the upstreams of the supposed attackers. Note that those approaches do not explicitly try to locate the attacker at the local router.

In contrast to this, interface-based countermeasures apply thresholds and limits per interface to narrow the attack down to an interface. The *Satisfaction* [1] approaches only consider the ratio of Interest packets and data packets to identify requests for non-existing content, whereas *Poseidon* [2] additionally correlates the number of current PIT entries. Those approaches lack the option to isolate more specifically because all nodes behind the throttled interface will be affected by the limitation.

The last class of approaches that we discuss in this paper are countermeasures that analyze PIT consumption per name prefix. *Threshold-based Detecting and Mitigating (TDM)* [6] classifies valid and malicious Interests based on the number of expired Interests with respect to a specific prefix. *Prefix Pushback* [4] focuses on the overall number of Interests per prefix and alarms downstream peers.

Table 1 summarizes the countermeasures which we analyze. We argue that the concrete heuristic to identify an attack is less important. Instead, the applicability of the approaches depends significantly on the topology and the principal detection point (i.e., prefix, interface, or router).

2. PRELIMINARY RESULTS & OUTLOOK

Setup The objective of this paper is the consistent comparison of different countermeasures. Thus, we decided to extend ndnSIM, the common NDN support in NS-3. Only three [1] out of the nine approaches were supported by default. To verify the existing and our new implementations, we reproduced the measurements described in the original publications and compared the results. This extensive testing helped us to improve the quality of both code bases. Our simulation code is publicly available via <http://interest-flooding.realmv6.org>.

We analyzed the PIT load, caching capabilities, and the ratio of Interest request and data delivery for different topologies. In the following, we will concentrate on the Interest request ratio as this measures fairness of the mitigation strategy with respect to legitimate consumers. We deploy a one hop star topology and a scorpion topology where the producer represents the sting and (legitimate and malicious) consumers represent the feet. Attackers request non-existing content, which exhibits distinct prefix (*/evil/**) compared

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
ICN'15, September 30–October 2, 2015, San Francisco, CA, USA.
© 2015 ACM. ISBN 978-1-4503-3855-4/15/09 ...\$15.00.
DOI: <http://dx.doi.org/10.1145/2810156.2812604>.

Approach / Acronym		Detection	Trigger	Mitigation
Token Bucket	TB [1]	per router	–	Round Robin over all interfaces
Resource Allocation	RA [2]	per router	Adaptive PIT size threshold	Drop subsequent Interests
Interest Traceback	IT [3]	per router	PIT size threshold & gain	Downstream traceback
Poseidon Local	PL [2]	per interface	Interest-data ratio	Limit PIT size per interface
Poseidon Distributed	PD [2]	per interface	Interest-data ratio	PL + alarm downstream peers
Satisfaction-based Accept	SA [1]	per interface	Interest-data ratio	Decrease probability of forwarding Interests
Satisfaction-based Pushback	SP [1]	per interface	Interest-data ratio	SA + distributed Pushback
Prefix Pushback	PP [4]	per prefix	Absolute # Interests	Drop specific ratio of Interests + alarm downstream peer
TDM	TDM [6]	per prefix	# expired PIT entries	capacity threshold

Table 1: Proposals to detect and mitigate Interest flooding attacks.

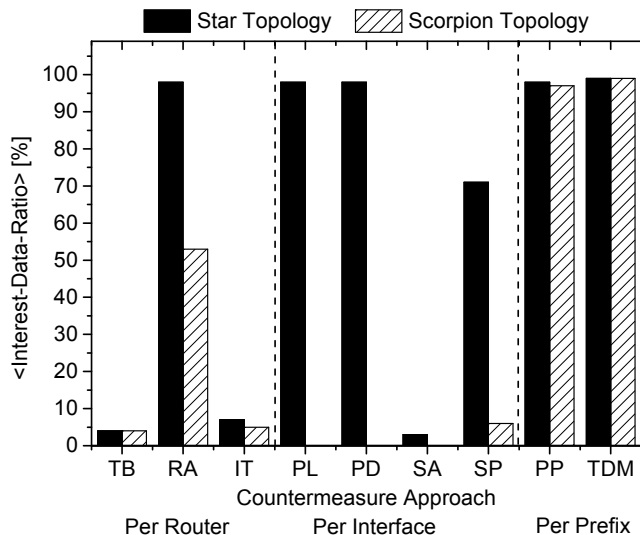


Figure 1: Fairness of different countermeasures.

to valid content (/good/*). Each simulation is sampled with the same parameter settings until it is converged.

Results Figure 1 shows the average Interest-data ratio over all runs and consumers, per countermeasure and topology. Higher y-values thus indicate better data delivery. We observe two key effects: (a) the results significantly depend on the underlying topology and (b) the detection point (per router, etc.). This is somewhat surprising as globally deployable countermeasures need to cope with heterogeneous networks, such as the Internet.

Prefix-based countermeasures outperform all other approaches in our scenario since these detection mechanisms can clearly identify malicious interests when the valid and the malicious data are prefix-free. Note that this changes as soon as malicious and valid Interest share a longest common prefix. This observation can be easily misused by an attacker when launching real-world attacks.

Interface-based approaches tend to exhibit similar behavior when deployed in star topologies, i.e., consumer, attacker, and producer are connected to the same router. In those scenarios, the interface complies with the maliciously requested prefix. However, in a scorpion graph the picture changes completely as the tail router cannot distinguish between

attacker and valid consumer—the complete downstream interface will be limited. Note that this single link property of the scorpion graph is very common in Internet backbone topology, in particular towards the edge networks.

Finally, we observe that most of the current per router mitigations are invariant of the specific topology but depend more on the drop Interests strategy. Strictly discarding Interests leads to worst case performance in all setups (see TB, IT). However, a more adaptive approach results in increased performance fairness. Analyzing this performance gain in more detail, will be part of our future work.

Outlook In this paper, we argue for both a comparable analysis of countermeasures against Interest flooding as well as the need for future work on this topic. In future work, we will investigate hybrid approaches, which combine different detection points (per interface and per prefix) to increase accuracy and performance while limiting PIT entries. Furthermore, we will extend our analysis, not only to cover more complex application scenarios but also to complement simulations by long-range real-world experiments.

Acknowledgments. This work was partially supported by the German BMBF within the projects SAFEST and Peeroskop.

3. REFERENCES

- [1] AFANASYEV, A., MAHADEVAN, P., MOISEENKO, I., UZUN, E., AND ZHANG, L. Interest Flooding Attack and Countermeasures in Named Data Networking. In *Proc. of IFIP Networking* (Piscataway, NJ, USA, 2013), IEEE Press.
- [2] COMPAGNO, A., CONTI, M., GASTI, P., AND TSUDIK, G. Poseidon: Mitigating Interest Flooding DDoS Attacks in Named Data Networking. arXiv 1303.4823, 2013.
- [3] DAI, H., WANG, Y., FAN, J., AND LIU, B. Mitigate DDoS Attacks in NDN by Interest Traceback. In *Proc. of IEEE INFOCOM NOMEN Workshop*, 2013, IEEE Press.
- [4] GASTI, P., TSUDIK, G., UZUN, E., AND ZHANG, L. DoS and DDoS in Named Data Networking. In *Proc. of ICCCN*, 2013.
- [5] WÄHLISCH, M., SCHMIDT, T. C., AND VAHLENKAMP, M. Backscatter from the Data Plane – Threats to Stability and Security in Information-Centric Network Infrastructure. *Computer Networks* 57, 16 (Nov. 2013), 3192–3206, (original version arXiv:1205.4778, May 2012).
- [6] WANG, K., ZHOU, H., LUO, H., GUAN, J., QIN, Y., AND ZHANG, H. Detecting and mitigating interest flooding attacks in content-centric network. *Security and Communication Networks* 7, 4 (April 2013), 685–699.