2016

Forwarding loop attacks and counter measures in content centric networks

S Sarat Chandra Velijala
Iowa State University

Forwarding loop attacks and counter measures in content centric networks

by

S Sarat Chandra Velijala

A thesis submitted to the graduate faculty

in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

Major: Computer Engineering

Program of Study Committee:
Yong Guan, Major Professor
Manimaran Govindarasu
Akhilesh Tyagi

Iowa State University

Ames, Iowa

2016

DEDICATION

I would like dedicate this thesis to my parents and my sister. I am immensely thankful for their support and encouragement throughout my thesis work. I would also like to thank my friends for their loving guidance during the writing of this work.

TABLE OF CONTENTS

LIST OF FIGURES

Page

NOMENCLATURE

ICN	Information Centric Networking
CCN	Content Centric Networking
NDN	Named Data Networking
CS	Content Store
FIB	Forwarding Information Base
PIT	Pending Interest Table
CR	Compromised Router
CC	Compromised Consumer
CP	Compromised Producer
IFA	Interest Flooding Attacks

ACKNOWLEDGMENTS

I would like to take this opportunity to express my thanks to those who helped me with various aspects of conducting research and the writing of this thesis. Firstly, I would like to thank Dr. Yong Guan for his guidance, patience and support throughout this research. His insights and words of encouragement have often inspired me and renewed my hopes for completing my graduate education.

I would also like to my committee members for their efforts and contributions to this work: Dr. Manimaran Govindarasu and Dr. Akhilesh Tyagi. Finally, I would like to thank all my family members and friends for their love and support.

ABSTRACT

Content Centric Networking (CCN) is a novel networking approach that aims at overcoming some of the limitations of the current Internet. In particular, CCN aims at providing better security and privacy by focusing on the data rather than on the location of data. However, this new networking concept opens up avenues for launching several new types of attacks including the “Forwarding Loop attacks”.

This paper describes how malicious customers can attack the availability of Content Centric Networks (CCNs) by creating forwarding loops. These loops cause one request to be processed repeatedly or even indefinitely, resulting in unwanted resource consumption and potential Denial-of-Service attacks. Next, we propose detection and mitigation techniques that will allow routers to identify and prevent the formation of such loops. To evaluate the practicality of such forwarding-loop attacks, we use the popular CCN simulation software, ndnSIM to simulate the occurrences of the loops and show how they can affect the overall service of the network.

CHAPTER I

INTRODUCTION

Today the Internet is mainly used for data dissemination to interested users, rather than for connecting hosts. The user is interested in data itself, while the location of data is of minor importance. However, the Internet was formerly designed and has evolved according to the host-centric communication paradigm. Recent studies have shown that the poor performance of the traditional Internet, in the areas of security, efficient content dissemination, content delivery etc., lies in its host-centric nature.

Information Centric Networking (ICN) is a new effort that aims to eliminate the traditional Internet's limitations. Content Centric Networking (CCN) is one of the proposed ICN approaches. In CCN the content or data is requested by the Client in the form of an “Interest” packet and the host, which can either be the primary source or the intermediate nodes who have cache capability, provide the data in the form of a “Content Object”. Thus in this new paradigm focus is shifted from location specifications of the content to the content delivery and encryption. As any new network protocol CCN is not free from malicious attacks. Interest Flooding attacks and Cache Poisoning attacks are more common and well explored.

In this work we present “forwarding-loop” attacks, which allow malicious CCN customers to attack CCN availability by creating looping requests within a single CCN domain(Intra-cluster) or across multiple CCN domains (Inter-cluster). Forwarding-loop attacks allow attackers to immensely consume CCN resources by building up a large number of requests (or responses) circling between CCN nodes. Although many CCN nodes have internal mechanisms to drop repeated content requests when they circle back, the attacker can ask for unique content requests

(with different segment names) thus preventing the request to be satisfied from the previous cached content on the intermediate nodes and continually avert the request to the primary content provider. In this way all the intermediate nodes are engaged in satisfying the requests from the attacker, indeterminably sustaining the loop.

1.1 Thesis Contributions

In this thesis we firstly define Content Centric Networking and its core principles and strategies We also discuss its implementation in Mobile Networks. Previous related work on the security issues and attacks in CCN are surveyed. We then present the forwarding loop attacks in CCN and discuss the three stages of its implementation in static, mobile-ad hoc and mobile-infrastructure networks. Later we describe a few strategies to mitigate these attacks. Finally, we use ndnSIM software tool to simulate and examine the impact of these looping attacks in single and multiple compromised node networks.

1.2 Thesis Organization

The rest of this paper as follows. Chapter 2 describes CCN operation, especially forwarding and flooding techniques. In Chapter 3, the security issues and related attacks in CCN are studied. Chapter 4 presents various forwarding loop attacks and discuss possible defenses to prevent or mitigate them. In Chapter 5 the aforementioned loop attacks are simulated using the ndnSIM software tool. We conclude in Chapter 6.

CHAPTER 2

CONTENT CENTRIC NETWORKING

2.1 Introduction

In the CCN architecture, a CCN node model is comparable to an IP node model. CCN nodes receive and send packets over multiple faces. A face in CCN is a connection to an application, or another CCN node, or some other kind of interface. A face may have attributes that indicate broadcast or multicast capability, expected latency and bandwidth, or other useful features.

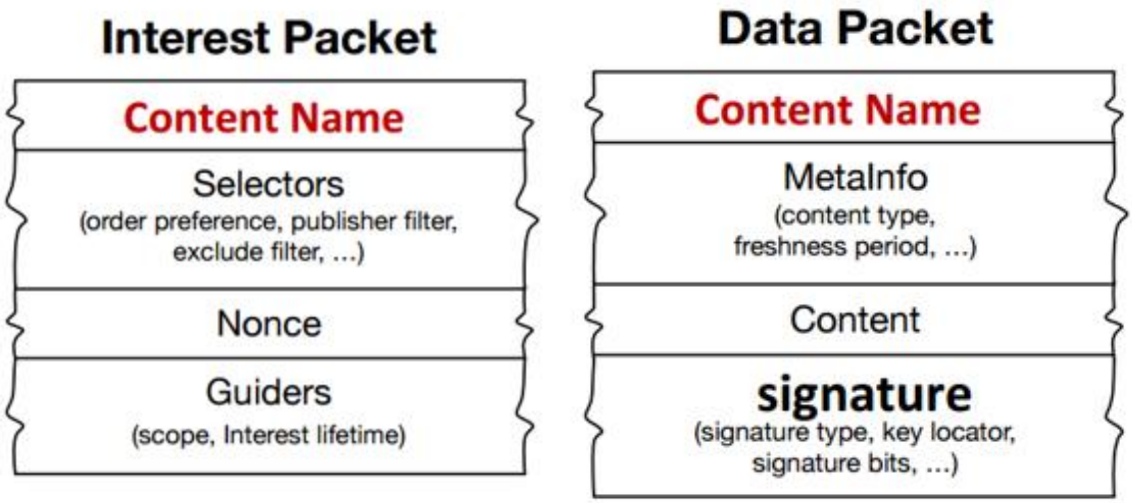

Figure 1. *Format of CCN Interest and Content Object*

A CCN node accepts Interest packets and either sends them out on an outgoing face, or directly replies with a matching Content object. If the node has forwarded an Interest packet, then it should normally receive a corresponding Content object which will be send to the original requesting face. A Content Consumer issues an Interest request for data over the network. The

Interest is transmitted through a set of intermediate Forwarders until the Content Object is found or the Interest's Life time expires.

A CCN node has three main data structures:

1) Content Store (CS)

2) Forwarding Information Base (FIB)

3) Pending Interest Table (PIT)

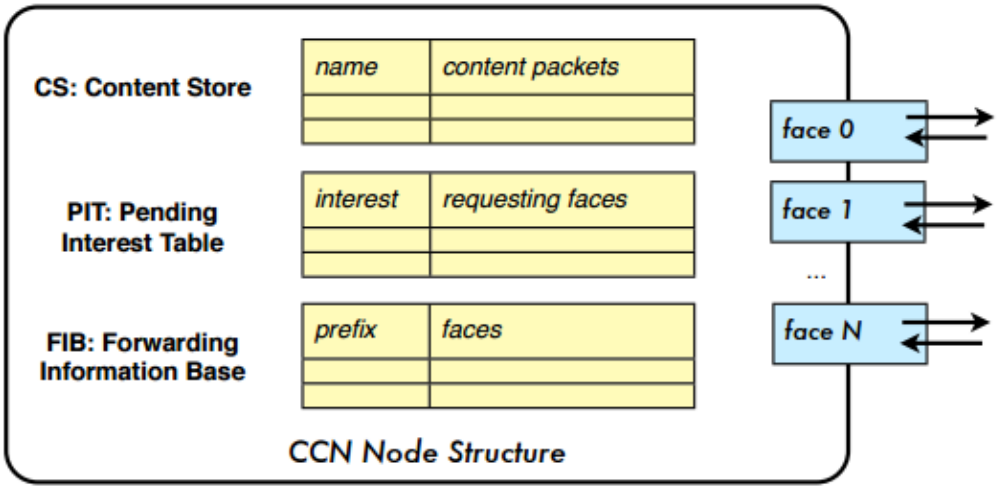

Figure 2.*CCN Node Structure*

The *Content Store (CS)* holds a table of previously seen (and optionally cached) Content objects indexed by the Name field of the packet. Besides providing communication buffering, the Content store also serves as a content cache. The CS is similar to the buffer memory of an IP node but has a different replacement policy. Since each IP packet corresponds to a single point-to-point conversation, it has no further value after it has been forwarded. Unlike IP, CCN packets are self-authenticating and self-identifying and can potentially be useful to many users. Thus, in order to

optimize the use of network bandwidth, and reduce user-perceived latency, CCN nodes store the Content objects in their CS. In CCN, only Interests are routed. Matching Content objects follow the reverse-path of the corresponding Interest packet.

The *Pending Interest Table (PIT)* is used to keep track of Interests forwarded upstream by that CCN node to- ward the content source so that Content objects later received can be sent back to their requestor(s). The PIT holds a set of entries, each entry containing a previously seen Interest packet and a list of the faces that received the interests. There may be additional information, such as timeout values, that affect the entries.

The *Forwarding Information Base (FIB*) forwards Interest packets towards potential data sources and is analogous to the FIB in IP routers. The FIB holds a set of entries, each entry containing a name prefix and a list of the faces that might provide content for that prefix. The FIB may be populated by an overlay routing protocol or with static routes There may be additional control information that affects forwarding.

2.2 Forwarding Strategy in CCN

We next describe the sequence of actions that are taken in order by CCN nodes when they receive an Interest or Content object.

CASE 1: *When an Interest Packet P with content name N is received on face F*

• Check for duplicate Interest packets:

----If there are recently seen Interests with the same name or P has reached its hop-limit, then P is discarded

• Check for existing data:

----If there is a Content Object in the CS whose name exactly matches N then send that Content object as a reply, and P is discarded

• Check for duplicate Interests waiting for replies:

----Else if N exactly matches the name in an entry of the PIT then F is added to the face list for that entry and P is discarded

• Forward the Interest towards a potential provider:

----Else if there are any entries in the FIB with names that are prefixes of N, then the face list L is taken from the entry with the longest matching prefix, and if L is not empty, then a new entry is made in the PIT for P and F, and is forwarded along one or more of the faces in L (possibly in parallel)

CASE 2: *When a Content Object D with content name N is received on face F,*

• Discard duplicate replies:

----If there is an exact match for N in the CS then D is discarded (as D is a duplicate)

• Discard unsolicited replies:

----Else if there is no match in the PIT then D is discarded (D is unsolicited)

• Forward replies to requestor:

----Else, the matching entry in the PIT has a face list L, and D is forwarded to every face in L, and the PIT entry is discarded

• Store in the CS:

-----Optionally choose to cache D in the CS

A Forwarder also implements **Interest aggregation**, so that multiple similar Interests do not get forwarded upstream. Aggregation is based on Interests that have the same Name, Key-Id-Restriction (if present) and Content-Object-Hash-Restriction (if present) and that fit within the same Interest Lifetime, where Interest-Lifetime is defined as the amount of time the consumer is willing to wait for the Content Object response.

Satisfying an Interest

A Content Object satisfies an Interest if:

(a) the Content Object name exactly matches the Interest name, and

(b) the Content-Object-Hash equals the Interest Content Object hash (if the restriction is given)

(c) the Content-Object-Key-Id exactly equals the Interest Key-Id (if the restriction is given). Only end systems (or Content Stores) need to verify cryptographic signatures, which decreases the computational load on each node and increases throughput. However, each hop may need to compute the Content-Object-Hash, if the pending Interest includes a Content-Object-Hash restriction.

2.3 Flooding Strategy and other concepts

(a) Arrival of New Interest

At the initial stage, the FIB on each CCN node is empty. An incoming Interest packet at Face 0 is propagated to all Faces (1, 2, and 3) except the incoming Face 0. This is initial flooding.

(b) Arrival of the First Data Packet

When the first Data packet arrives, the corresponding FIB entry is created. The prefix of the Data name is stored in the Prefix field and the arrival face of the Data is recorded in the Face(s) field, which can maintain N maximum elements (faces).

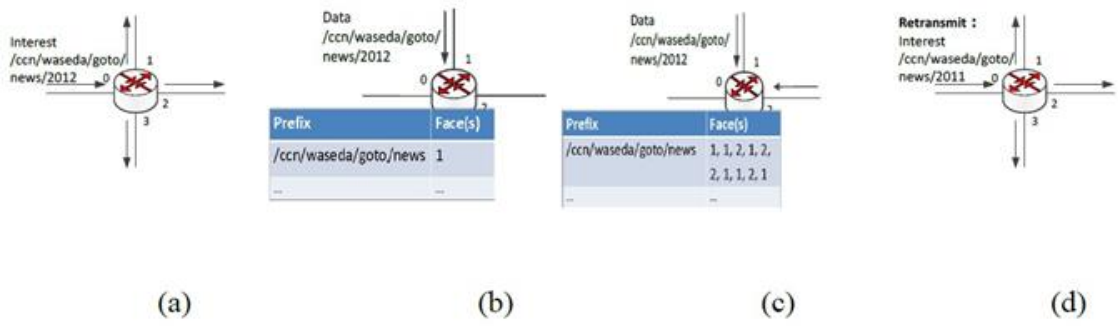

Figure 3. *Flooding Strategy in CCN*

(c) Arrival of Data Packet with the Same Prefix

When there is the second Data packet with the same prefix as an existing packet in the FIB, the arrival face will be added to the corresponding Face(s). If N faces are already stored in the set of Face(s), a FIFO operation is performed. That means that the first oldest arrival face will be deleted.

(d) Interest Forwarding Principle

As described in the previous section, when a new Interest packet arrives that matches Prefix entries in the FIB, Interest forwarding is performed by selecting one face among the elements in Face(s) according to the occurrence ratio of the faces. That is, the most successful face is selected.

Since the FIB can maintain at most N faces, the learning mechanism adjusts the face selection according to the recent data retrievals. In the case of link failure or packet loss, data does not return in time; therefore, the data requester or consumer returns to step [a], flooding the Interest to all available connected faces to discover a working path quickly.

Lifetime Strategy

CASE 1: If a new, similar, Interest comes from the same previous hop then:

• If the new Lifetime would extend PIT entry lifetime, update the PIT entry to the max of (old Lifetime, new Lifetime) and forward.

• if the new Lifetime would not extend PIT entry, forward (no PIT update required).

CASE 2: If a new, similar, Interest comes from a different previous hop then:

• If the new Lifetime would not extend the PIT entry, update PIT with new previous hop and don't forward.

• If the new Lifetime would extend the PIT entry, update PIT entry to the max of (old Lifetime, new Lifetime), add the new previous hop, and forward.

Loop suppression in CCN

CCN inherently supports multipath routing. IP routing adopts a single best path to prevent loops. In CCN, Interests cannot loop persistently, since the name plus a random nonce can effectively identify duplicates to discard. Data do not loop since they take the reverse path of Interests. Thus an CCN router can send out an Interest using multiple interfaces without worrying

about loops. The first Data coming back will satisfy the Interest and be cached locally; later arriving copies will be discarded. This built-in multipath capability elegantly supports load balancing as well as service selection. For example, a router may forward the first few Interests out via all possible interfaces, measure the performance based on returning Data, and select the best performing interface(s) for subsequent Interests.

Strategies for the FIB tabulation

The other faces in a FIB prefix entry are learned in different ways. Sources of data, such as the repositories, arrange to receive Interests for the prefixes they service by doing a *Register* operation to the local CCN core. This creates local FIB entries for the registered prefixes that have the repository application's face in their FIB face lists. The registered prefixes have optional flags that indicate if they have to be advertised outside the local machine. Announcement agents then read the registered prefix table on the local node (via CCN Interest-Data to a namespace reserved for local node communication) and advertise the flagged prefixes that meet their policy constraints.

Hop Limit

CCN distinguishes between "local" and "remote" hops. A local next hop is a directly attached application running locally on the system. A remote next hop is not local to the current system. These conventions along with how we use Hop-Limit result in expected behavior. If a local application sends an Interest with a Hop-Limit 0, that Interest will only go to other applications on the system. If it sends an Interest with a Hop-Limit 1, it will go to local applications

plus the 1-hop neighbors of the system. If a system receives a remote Interest with a Hop-Limit 1, it will be decremented to 0 and then only forwarded to local applications.

2.4 Mobility in CCN

Machines today typically have multiple network interfaces and are increasingly mobile. Since IP is restricted to forwarding on spanning trees, it is difficult for IP to take advantage of more than one interface or adapt to the changes produced by rapid mobility. CCN packets cannot loop so CCN can take full advantage of multiple interfaces. CCN talks *about* data, not *to* nodes, so it does not need to obtain or bind a layer 3 identity (IP address) to a layer 2 identity such as a MAC address. Even when connectivity is rapidly changing, CCN can always exchange data as soon as it is physically possible to do so. They can be broadly classified as mobile ad hoc and infrastructure networks. In the mobile ad hoc networks there is no central entity to oversee the transactions between the mobile nodes. Nodes join and leave the network in a random fashion. On the other hand, in the mobile-ad hoc networks, the central entity (generally the Access point) manages the interaction and movements of the mobile nodes. Mobile CCN networks are further discussed in the Chapter 5.

CHAPTER 3

RELATED WORK

3.1 Interest Flooding Attacks (IFA)

Malicious/compromised users may exploit the PIT-based forwarding mechanism of NDN to launch the IFA, which is considered as one of the most serious types of DDoS attacks on NDN [7]. According to IFA, the malicious user (or a group of users) will issue a large number of bogus Interest packets. Each router, upon receiving each of these packets, will create an entry in its PIT and will forward the packet to the next-hop node (router or content source). According to the NDN rules, an entry is removed from the PIT when the entry has expired or the router received the corresponding Data packet before the entry expiration.

According to the above, the best attacking strategy is to issue Interest packets for non-existent content. In this case, the bogus entries will stay in the PIT as much as possible. The goal of the attacker is to quickly fill in the PIT and to keep it full, so that the Interest packets originated from legitimate users will eventually be dropped.

For example, assume that the PIT capacity in each router is 3 entries. The attacker's strategy is to send 3 bogus Interest packets for (different) non-existent content. These packets will fill in the PITs of both routers. The source will drop these packets, since they request non-existent content. However, the corresponding entries will stay in the PITs until they expire. After the expiration, the attacker will issue 3 new Interest packets, aiming at keeping the PITs always full. This way, some, or even all, Interest packets of legitimate users will be dropped.

Interest flooding attack is the most common source of DoS attacks in CCN. Therefore, many countermeasures proposed against CCN DoS attacks primarily focus on mitigating the flooding attack. In CCN, flooding attacks are triggered to degrade a router's performance, network's responsiveness, or the performance of a content source. This CCN-DoS classification is different from traditional DoS classifications because of some of CCN's distinguished characteristics (flow balance, content integrity with publisher's signature, etc.). We classify CCN-DoS attacks considering three main attack methods, that is, (i) *flooding*, (ii) *forced computation,* and (iii) *cache/content manipulation* under which different forms of attacks are possible.

3.2 Undetected Interest Loops in CCN

The author in [2] defines an Interest loop in CCN and has postulated two theorems based on the loops.

Interest Loop: An Interest loop of h hops for NDO with name $n(j)$ occurs when one or more Interests asking for $n(j)$ are forwarded and aggregated by routers along a cycle $L = \{v1, v2, ..., vh, v1\}$ such that router vk receives an Interest

for NDO $n(j)$ from $vk-1$ while waiting for a response to the Interest it has forwarded to $vk+1$ for the same NDO, with $1 \leq k \leq h$, $vh+1 = v1$, and $v0 = vh$.

Two kinds of Interests loops are depicted below in Figure 4.

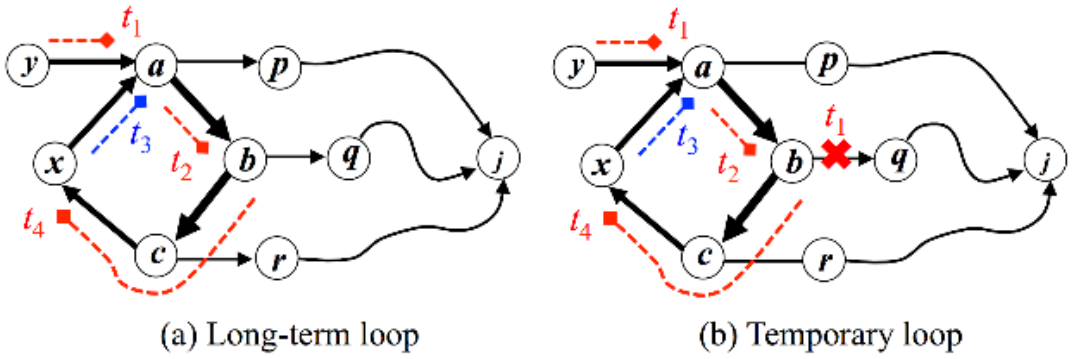

Figure 4. *Long-term and Temporary loops in CCN*

Theorem 1: CCN forwarding strategies specified in cannot ensure that Interest loops are detected when Interests are aggregated, even if the nonces were to denote Interests uniquely. The theorem assumes that all messages are sent correctly and that no routing-table changes occur to show that the CCN forwarding strategy can fail to return any content or NACK in response to Interests independently of network dynamics.

Theorem 2: No forwarding strategy can be correct if it allows Interest aggregation and attempts Interest-loop detection by the matching of Interest identification data.

3.3 Attacks in Mobile CCN

Mobile node(s), controlled by an attacker, can exploit the shared link in CCN while creating Denial of Service effect. Two possible ways to do this are as follows:

(i) A mobile node can issue a large number of interests to ask for bulk data while not maintaining interests itself. Subsequently, data (content objects) will arrive at the local

link without a receiver. A huge impact may be observed in such a case when there exists limited bandwidth in a mobile environment. A mobile attacker effectively jams a region using this technique;

(ii) A mobile node can issue a large number of interests to ask for bulk data while leaving the local link and traversing into neighboring regions using circular routes (the departure cannot easily be detected in a shared link environment). In this way, a mobile attacker continuously offloads its interest bundles and blocks the availability of network for legitimate users.

Jamming attack: A node on a shared link may issue a large number of content requests without maintaining the Interests at its own (losing interest). Content will then arrive at the local link without a receiver. This is particularly harmful in mobile environments of limited bandwidth. A mobile attacker can jam a region by traversing shared radio links while requesting bulk data.

Mobile blockade: A mobile node may issue a large number of invalid (or slow) Interests that block the state table of the access router for the period of state timeout. In a shared link-layer environment that cannot easily detect its departure, the mobile adversary can traverse neighboring networks on circular routes and continue to offload its interest bundle with the effect of a blockade of the regionally available networks. Initial countermeasures are difficult to apply, as the retransmission of Interests is part of the regular mobility pattern in ICN.

Overview of expected future work on countermeasures: Some existing work to address the mobility issues of CCN suggests special signaling that informs the upstream nodes before the mobile node leaves so that content distribution to the next access point can be established in advance. This is what fast mobile IP is doing. Nonetheless, the configurations of mobile nodes in CCN still demand special tuning to avoid CCN-DoS attacks in mobile environments. Currently, CCN routing and security in mobile environments are at an initial stage of research, and special schemes of defense will be required in the future to detect and mitigate CCN-DoS attacks as well as many other security and privacy threats of mobile CCN environments.

CHAPTER 4

FORWARDING LOOP ATTACKS IN CCN

Malicious customers of nodes (Consumers, Publishers or Routers) can deliberately manipulate the forwarding to create forwarding loops inside CCNs. Forwarding loops can cause CCNs to process one client request repetitively or even indefinitely. The consequent amplification effect allows malicious customers to launch, with little resources and cost, resource-consuming DoS attacks (Interest Flooding Attacks) against CCNs. In this thesis, we have identified approaches to create forwarding loops in static CCN network systems and extend it to Mobile CCN systems.

The process of Forwarding Loop is described in three main stages at each network system:

(i) Identification of the loop

(ii) Creation and Sustenance of the loop

(iii) Exploitation of the loop

4.1 Forwarding Loops in Static CCN Networks

We assume the physical topology of the underlying closed-static network is Mesh type and that we have a Compromised Router (CR) or a rogue-router (both terms used interchangeably) which is actively listening or transmitting on two or more faces. CR assisted by an independent set of consumer and producers adjacent to the router (one-hop away). This set of compromised producers (CP) and consumers (CC) aid in creating the required malicious Interest and Content

objects to be routed by the compromised router (CR). We have identified loops with one and two compromised routers, which can also be generalized to n-compromised nodes.

(A) Loop traversing one compromised router

Step 1: Identification of the Loop

The forwarding loop can either be identified by the CR independently or with the aid of another malicious consumer node, which is located in the same domain as that of the CR.

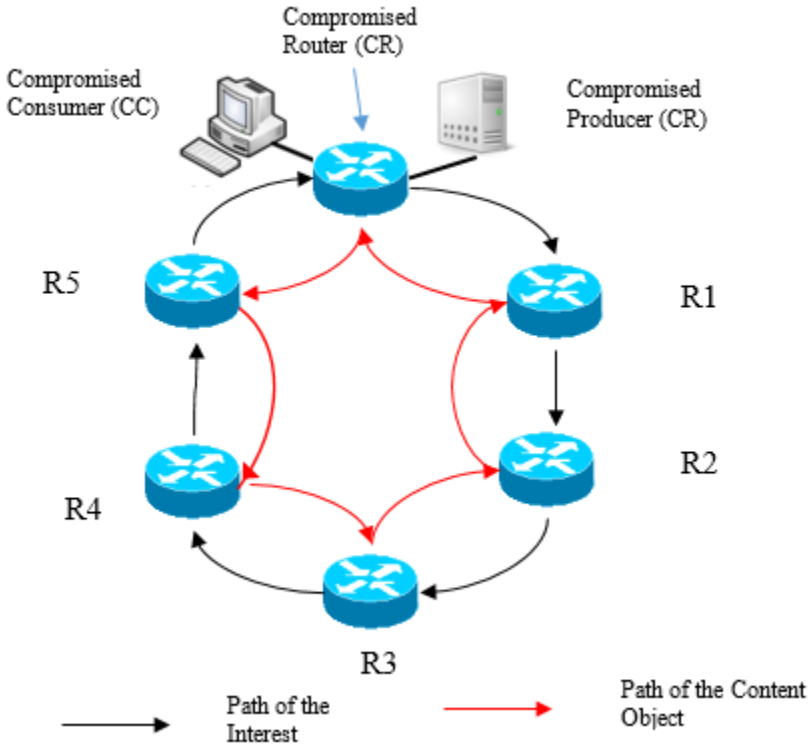

Figure 5. *Forwarding Loop with one CR in Static CCN*

In the first scenario, the CC sends out a fake Interest packet to the which in turn forwards it only on one of its outgoing faces. The specific face chosen can based on prior strategic analysis of the incoming interests on multiple interfaces of the routers, which is beyond the scope of this thesis. In the default case, as described in the previous sections, an Interest packet for unknown content is sent on all the broadcast capable faces (and later on the remaining faces). This fake interest packet is generated for a unique content name such that it can only be satisfied by the content object generated by the CP (or the inherent producer). The prefix of this unique content is purposefully not registered by the CR (can also be considered as the content being invoked/created for the Interest on the fly) in its CCN domain so as to trigger a broadcast flooding of the interest, in search of the satisfying content.

After a couple of iterations, the CR may receive the Interest back on one or more of its remaining faces, excluding the one it had previously sent the Interest out on. Figure 6 shows a simple scenario with a loop comprising of 6 nodes: CR → R1→R2→R3→R4→R5→CR. This proves the existence of a loop, traversing the CR, in the CCN domain. We can also assume that the loop created by the Interest packet received last on one of the faces, is the longest or has the highest cost, among all the loops detected and either of these properties can be further exploited by the attacker.

In the second scenario, another malicious consumer node, located in the same domain as that of the CR sends out the interest for unique content name which only be satisfied by the content object generated by the CP. This Interest packet is flooded in the domain and may be received on one or more faces of the CR. If the interest is indeed received on two or more faces of the CR, we

can safely assume the existence of the loop traversing the CR, which can be created and exploited as described below. The above scenario can also aid in the identification and creation of longer loops.

Step 2: Creation of the Loop

Now the attacker asks for an Interest with the same unique prefix as before, but with a different segment name, on the same router node. The router checks its FIB, and identifies the prefix name and the interface. It creates an entry in the PIT and forwards it along the interface on which the CO was received earlier. The Interest is forwarded along the singular path of nodes until it finally reaches the rogue router.

Consider the following scenario with a compromised router. The initial adjacent node (R1 in Figure 6) to which the attacker had asked for content, will be unable to find any cached content for this unique Interest in the CS. There wouldn't be any similar pending interests in the PIT nor any entries for the prefix in the FIB. According to the current methods, the initial edge node, floods the Interest on all its faces and adds an entry in its PIT for the same. The interest propagates in the network along the core nodes, following the above flooding strategy, until it reaches the rouge router node (CR → R1→R2→R3→R4→R5→CR). The rouge router is the only one which can satisfy the specific interest, such that it may receive the same interest on multiple interfaces from adjacent nodes in the same network. This strategy aids in interpreting the "loops" to the starting node on which the Attacker had asked for an interest.

But the Rogue node instead of sending back the Content Object on all the interfaces on which the interest was received (CR → R5→R4→R3→R2→R1→CR), it sends it only one particular face (of the so-formed loop). The CO is passed along all nodes on one singular path until it reaches the final node and forwarded to the Attacker. When the CO is received, the PIT entry for the unique interest is eliminated and an entry is added for that unique prefix in the FIBs (along with the ingress interface on which the CO was received) of all the nodes along that singular path. The content may also be cached in the CS of the nodes along the path. The PIT entry for the unique interest for the all other adjacent nodes, on which the CO was not received, will expire and no entry is added to the FIB or CS.

Instead of creating fake Interest packets, the compromised producer can also anticipate popular content and is ready with the content, before any other node can provide it.

Step 3: Exploitation of the Loop

Now without the help of the attacker, the Rogue router may send another Interest with the same prefix along the reverse loop. Next rogue router can send the same Interest with same name but a different nonce and with an Interest lifetime more than specified before.

So according to the strategy, the previous entry's Lifetime is increased. The routers keep increasing the Interest Lifetime of the Interest in the PIT entry, and this is propagated throughout the loop and the whole operation is stalled. In this scenario the genuine Interests packets get dropped and the attacker may employ Bots to increase the miss rate.

CC may also send out other many more Interests packets (either requesting unique fake content or a huge file which has multiple segments). These fake interest entries can then overflow the PIT table, which will subsequently drop the legitimate Interest packets from legitimate users, having reached the hardware memory capacity, thus leading to Interest flood attacks.

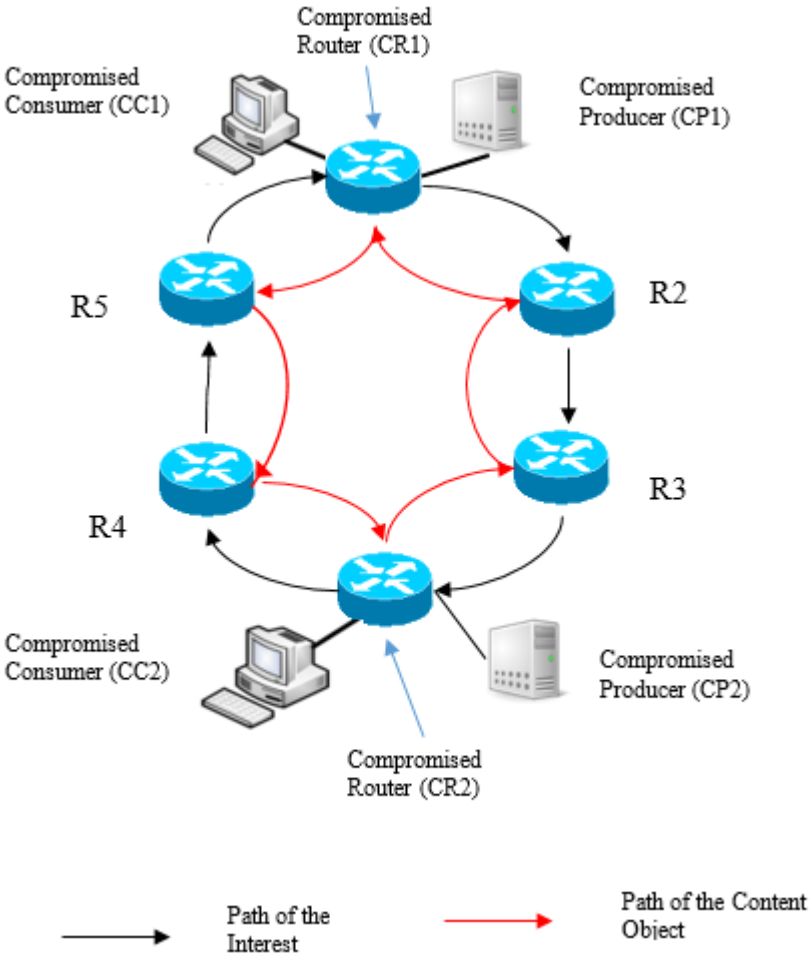

Figure 6. *Forwarding Loop with two CRs in Static CCN*

(B) **Loop traversing two compromised routers**

In this scenario, we assume we have two compromised routers (CR1 and CR2) that are located in the same domain preferably located multiple-hops apart from each other. First step would be to identify the presence of each other, by sending each other Interests messages for unique content which can only be satisfied by the other counter-part.

Both the CRs are listening on two or more faces. Initially one of the CRs sends an Interest on all its faces, for a prefix which has not been registered in the domain. By the default strategy for ad hoc networks the Interest packet is broadcasted to all the nodes in the domain, including the other CR. If the other CR receives the Interest on one or more of its faces, we can safely assume that they are multiple routes to the other CR.

Now the Rogue router does not send the CO for the new interest along the face on which it was received. Instead it floods an interest, with the same Interest name but with a different Nonce, along the other interface(s) to the adjacent nodes of the loop previously detected. Since the adjacent nodes still perceive the Interest to be unique, flood the Interest on all their respective faces and add an entry in their respective PITs for the same. This Interest is probated along the loop until it reaches the Initial edge node on which the attacker is asking for content, so completing the loop.

Now when the initial edge node receives the Interest with the same name but different nonce, it adds the ingress face to the PIT entry already waiting for the CO on the same name and deletes the interest. The rogue router now sends a CO along the initial singular path and this CO is satisfied along the loop, with the creation of the FIB entry and/or cache entry. This confirms the loop.

4.2 Forwarding Loops in Mobile CCN Networks

The forwarding loop attacks can be extended to the mobile CCN systems as discussed below. We assume the compromised node (CN) in the Ad Hoc networks has the feature of inherently creating both Interest and Content Objects, thus acting as the consumer, the producer (or publisher) and router independently.

In the below scenario (Fig.7a) the forwarding loop in a simple ad-hoc network with a single compromised node is depicted. We observe that the CN acts as both the CC and the CP. The loop follows similar stages as previously discussed in the Static CCNs in the above section. The only difference being the inherent capability of the Ad Hoc nodes acting as the CR, CP and CC.

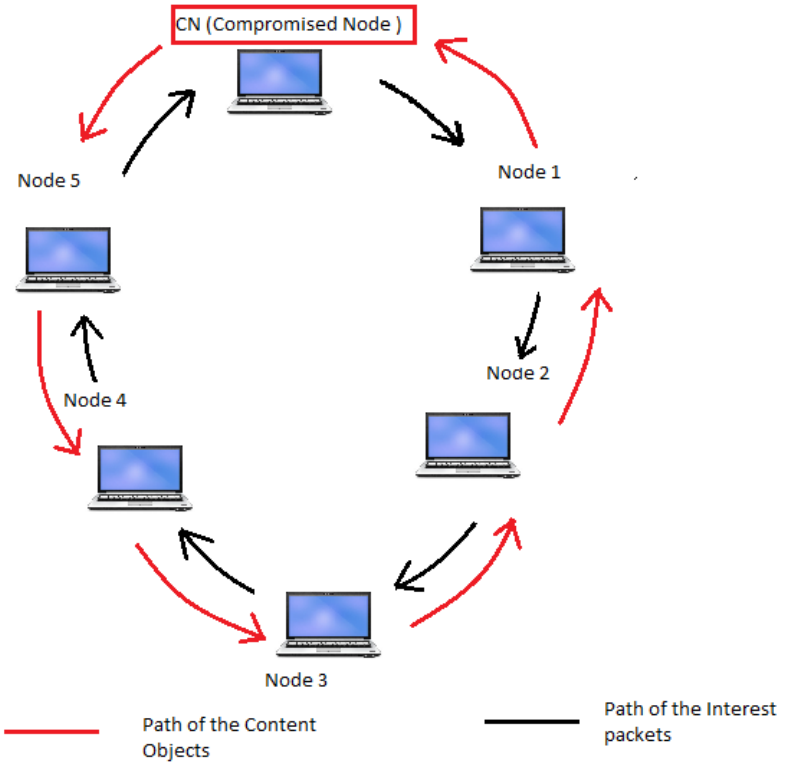

Figure 7a. *Forwarding Loop with one CR in Ad hoc CCN*

The implementation of two Compromised nodes CN1 and CN2 is depicted in Fig 7b. Both the nodes establish a connection and exploit the resources of the intermediate nodes. The loop can be initiated by either of the CRs and traverses all the intermediate nodes. We can further extend this scenario to loops created by the multiple compromised routers, which can identify larger loops extended between domains thus exploiting larger set of resources.

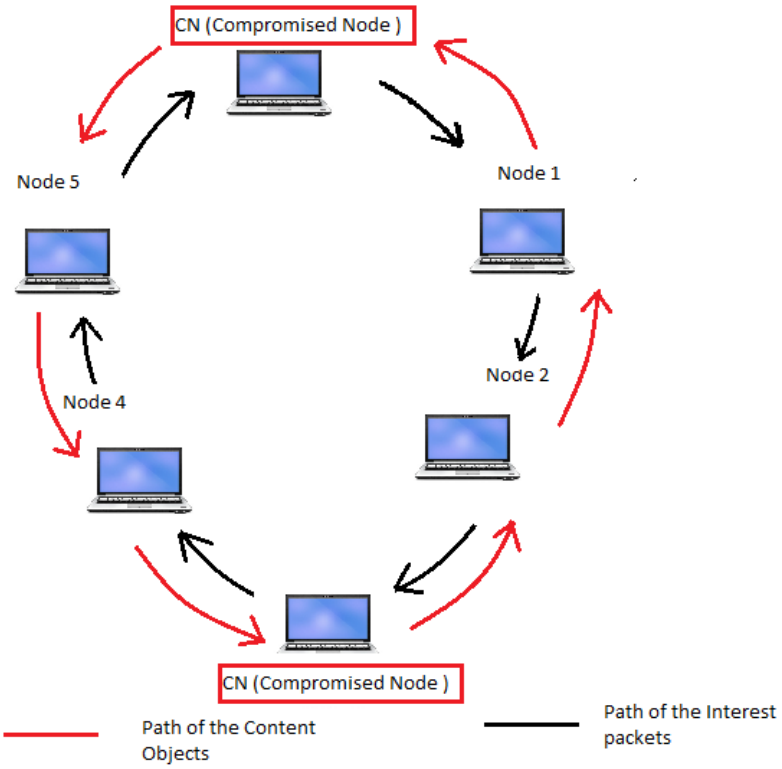

Figure 7b. *Forwarding Loop with two CRs in Ad hoc CCN*

We can also extend it to the mobile CCN infrastructure (wireless or cellular) networks with the underlying hierarchical topology. We present two instances in this scenario as depicted in Fig.8a and 8b. In the first instance the mobile node is travelling through a single domain (Intra-domain CCN system), where it could register as a Producer with one edge node (AP3) and

strategically choose another nearby edge node (AP4) where it can act as a Consumer. In the identification phase of the loop, the mobile node could register its fake unique name space acting as a mobile Producer on edge node and then send an Interest packet for that unique content on another nearby edge node now acting as a mobile Consumer node. If the interest is received on the producer side, we could safely assume the existence of a loop. This could then be preceded by the establishment and exploitation phases of the loop as discussed earlier.

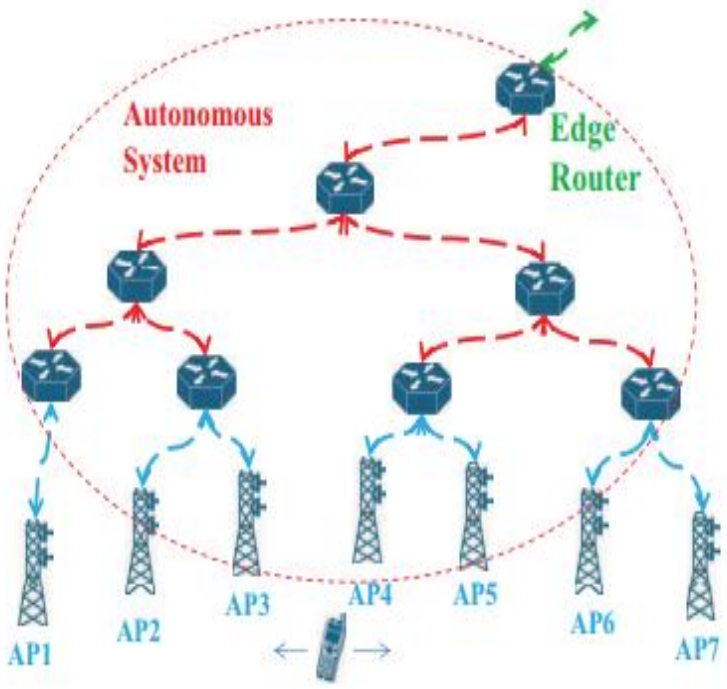

Figure 8a. *Forwarding Loop in Intra-Domain CCN*

In the second instance, the mobile node is travelling through two adjacent domains (Inter-Domain CCN system) where it can register as a Producer in one domain (AP4) and a Consumer in another domain (AP5). This could lead to larger loops traversing two domains engaging extensive resources of the intermediate nodes.

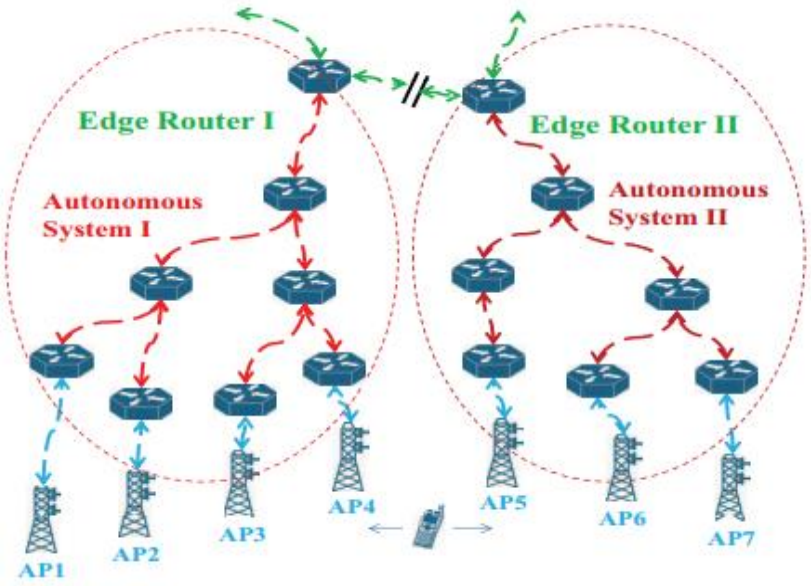

Figure 8b. *Forwarding Loop in Inter-Domain CCN*

4.3 Detection and Mitigation Techniques

In this section we discuss the possible detection and mitigation strategies of Forwarding Loops. The implementation and impact of these loops is studied in the Chapter 6.

Node ID/ Router ID fields

One of the main reasons for the creation of the forwarding loop is the retransmission of the Interest back to the to the compromised node. This can be avoided by the implementation of a Router-ID field. Using a Router-ID. along with the Interest packet. R-IDs are assigned when a node joins the Ad-Hoc network and identifies itself to the nearby node by sending a PROBE request and every node is aware of the R-IDs of its one-hop neighbors. When forwarding the

Interest for the unknown content, a particular node checks the R-ID of the Interest packet. If the R-ID matches to its one-hop neighbors it drops the packet, thus avoiding the formation of the loop.

Maximum Interest Life-time (MIL)

The Maximum Interest Life-time (MIL) assumed by a router before it deletes an Interest from its PIT should be large enough to preclude an excessive number of retransmissions. On the other hand, MIL should not be too large to cause the PITs to store too many Interests for which no NDO messages or NACKs will be sent due to failures or transmission errors. A few seconds would be a viable value for MIL. In practice, however, the consumer submitting an Interest to its local router could provide an initial value for the Interest lifetime estimated over a number of Interests submitted for NDOs in the same NDO group corresponding to a large piece of content (e.g., a movie). This is specially the case given our assumption that Interest retransmissions are carried out by content consumers, rather than by routers. Furthermore, because the CCN forwarding strategy does not detect loops when Interests are aggregated, many Interest entries in PITs may have to be stored until their lifetimes expire.

Suppression of Malicious Nodes

The methods of detecting and mitigating Interest Flood Attacks elaborated in [7] which are mainly focused on edge routers could be extended to the core routers in Static CCN systems to detect the compromised routers which generate exceedingly high traffic. The steps involved in the detection of the malicious node and its suppression is briefly stated in the following steps.

1) Attack detection:

During the detection phase, the edge router keeps statistics about the expired PIT entries per each user. Two thresholds are used to classify users into: legitimate, suspicious (possible attackers), and malicious (attackers). If the number of expired PIT entries per time unit, N of a user u is below the low threshold, T-low, user u is considered legitimate. If N is above T-low but below the high threshold, T-high, user u is considered suspicious. Finally, if N> T-high, user u is considered malicious.

2) Rate reduction and blocking phase:

During this phase, any user that has been classified as malicious, will be blocked, whereas the suspicious users will receive reduced data rate.

3) Attack notification phase:

If an edge router detects an ongoing attack, after blocking this user, it will notify other routers about the identity of the malicious user, by sending the attack notification packet. This is done to prevent the Mobile Interest Flooding Attack, where a mobile user periodically visits different routers and floods them with Interest packets. In this context, the notion of router is extended and refers to any data-forwarding network element, such as a Wi-Fi Access Point (AP) or a Base Station (BS) in a cellular network.

CHAPTER 5

PERFORMANCE EVALUATION

The fundamental departure of the CCN communication paradigm from the Internet Protocol principles requires extensive evaluation through experimentation, and simulation is an essential tool to enable the experimentation at scale. For this purpose, we rely on the open source simulation software ndnSIM.

We consider a fixed set of nodes in the network. In the first scenario we evaluate the packet dropping probability of the entire network, with respect to the length of the forwarding loop involving a single compromised router (Self-loop). We consider a mesh network in matrix representation of 5X5, with 24 ordinary routers and 1 compromised router. The CR, located at the center of one edge of the matrix, has both the CP and CC assisting it in administering the forwarding loop. We consider a simple scenario where we have 9 consumer nodes (C1-C9) and 9 producer nodes (P1=P9). Each of these consumer nodes sends Interests requesting content that can be satisfied by the corresponding Producer node (C1 for P1, C2 for P2 and so on). The producers listen on specific prefixes (we chose a random prefix: /producer*n*/ for each of the n- producers) previously announced in the network.

To show the negative impact of the forwarding loops, we evaluate the packet dropping probability of actual users (producers and consumers) due to PIT overflow. We consider the PIT capacity equal to 250 MB and PIT entry expiration time of 500ms. The rest of the simulation parameters remain the same as described in the previous paragraph. In the first scenario depicted in Fig 9, we have considered loops of six different lengths. Here the length of the loop is

determined by the number of intermediate nodes engaged in the loop including the compromised node. We consider loops of length 5,7,9,11,13 and 15. The length of the loop can be varied by manipulating the Hop-Limit attribute in the Interest packet sent by the CC. For example if the hop-limit field is set to 4 in the Interest packet the loop would engage 5 nodes including the CR. The hop-limit limits the number of nodes the Interest packet can be routed to before being discarded. For each individual loop we calculate the packet dropping probability.

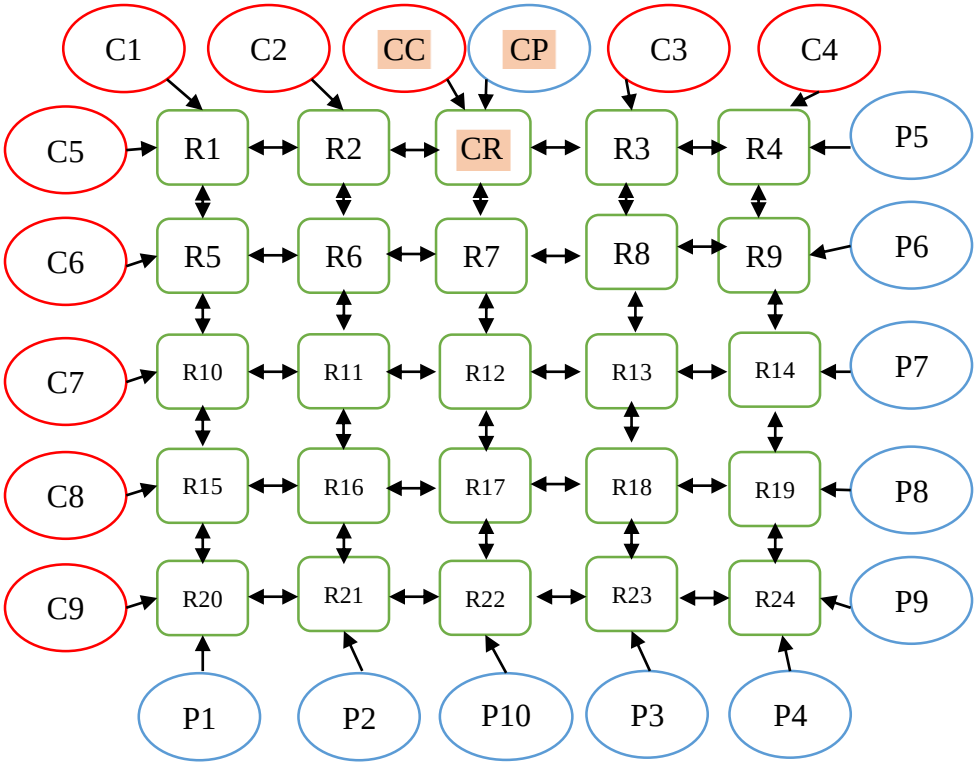

Figure 9. S*imulation Scenario 1: One Compromised Node*

The results are outlined in the graph (Fig 10) were we plot the rate of the packet dropping probability over time for the six individual loops. We observe that a large number of legitimate requests will not be satisfied, since the corresponding PIT entry of each request will expire before the content arrives. We can also observe that dropping probability increases with increase in the length of the loop. These results show that it is relatively easy even for attacking nodes of limited capabilities to quickly occupy large amounts of router's storage and processing resources.

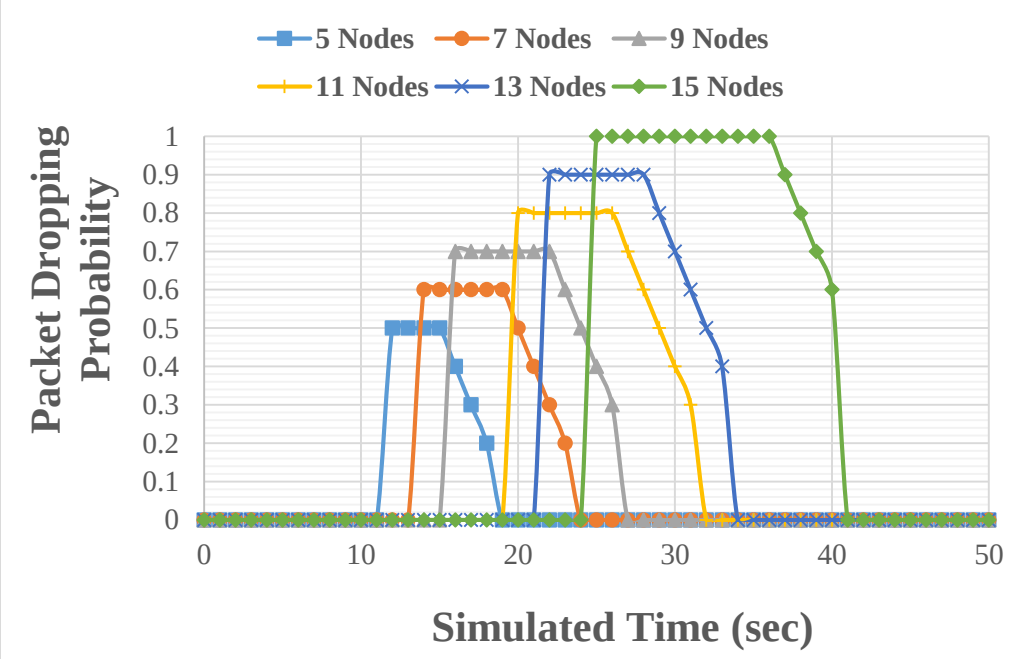

Figure 10. *Graph depicting the variation of Packet Dropping Probability with the Loop Lengths*

In the second scenario we extend the above evaluation to assess and compare the situation with two and four compromised routers with the same underlying 5X5 matrix. In this scenario a fixed length loop of 15 (Hop-limit of 14) is considered. The rest of the simulation parameters remain the same as described in the previous scenario. In the first case we deploy two compromised

router nodes, CR1 and CR3 (with their corresponding CCs and CPs), which are located at center of the opposite edges of the matrix. We then calculate the change the packet dropping probability of the entire network over the period of 50 seconds. In the second case we deploy four compromised router nodes, CR1, CR2, CR3 and CR4 (with their corresponding CCs and CPs) located at the center of all the edges of the matrix. As above, we calculate the change the packet dropping probability of the entire network over the period of 50 seconds. The results of the evaluation are depicted in the graph of Fig. 12.

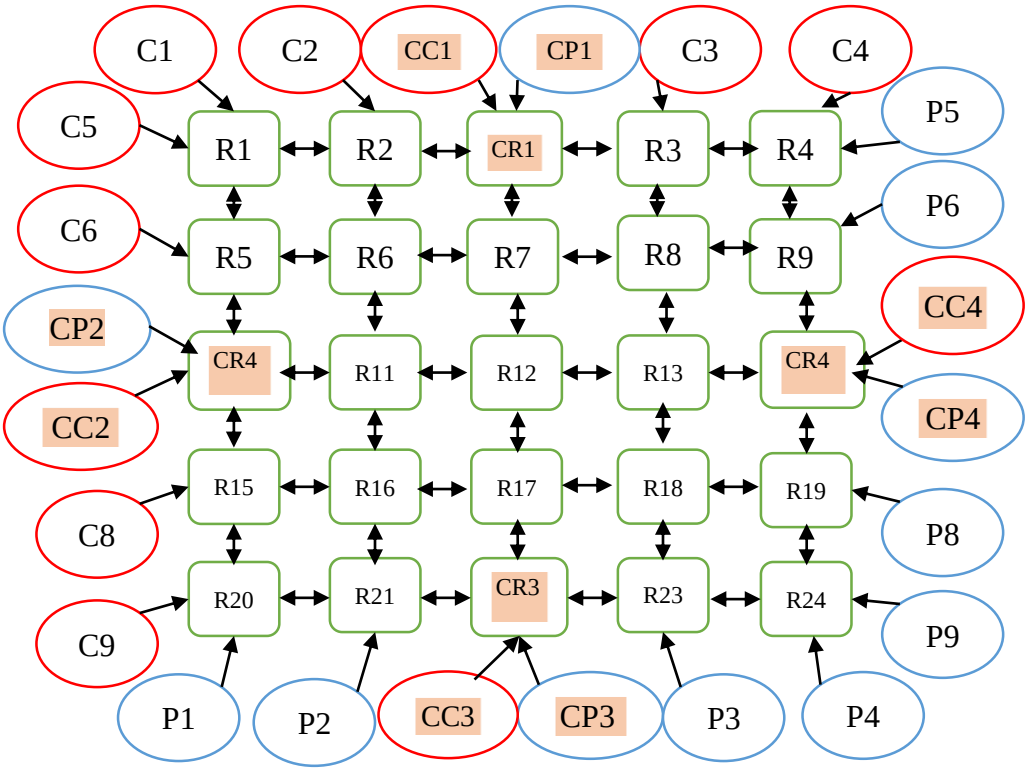

Figure 11. *Simulation Scenario 2: Two and Four Compromised Routers*

We then try to implement the Blocking mechanism as a counter measure to eliminate or prevent the occurrence of the loop. We propose a blocking technique where we send an exclusive broadcast packet called ‘Block’ which is an Interest packet listing the namespaces that are to be blocked i.e. all the interest packets with those restricted namespaces are deemed to be malicious and should be dropped by the router. All the future Interest packets querying for content with the restricted namespace are dropped and the loop is eliminated. This technique is implemented in the second case with four compromised routers and the results are populated in Fig 12. We observe that the packet dropping probability is decreased with the blocking technique in the second case.

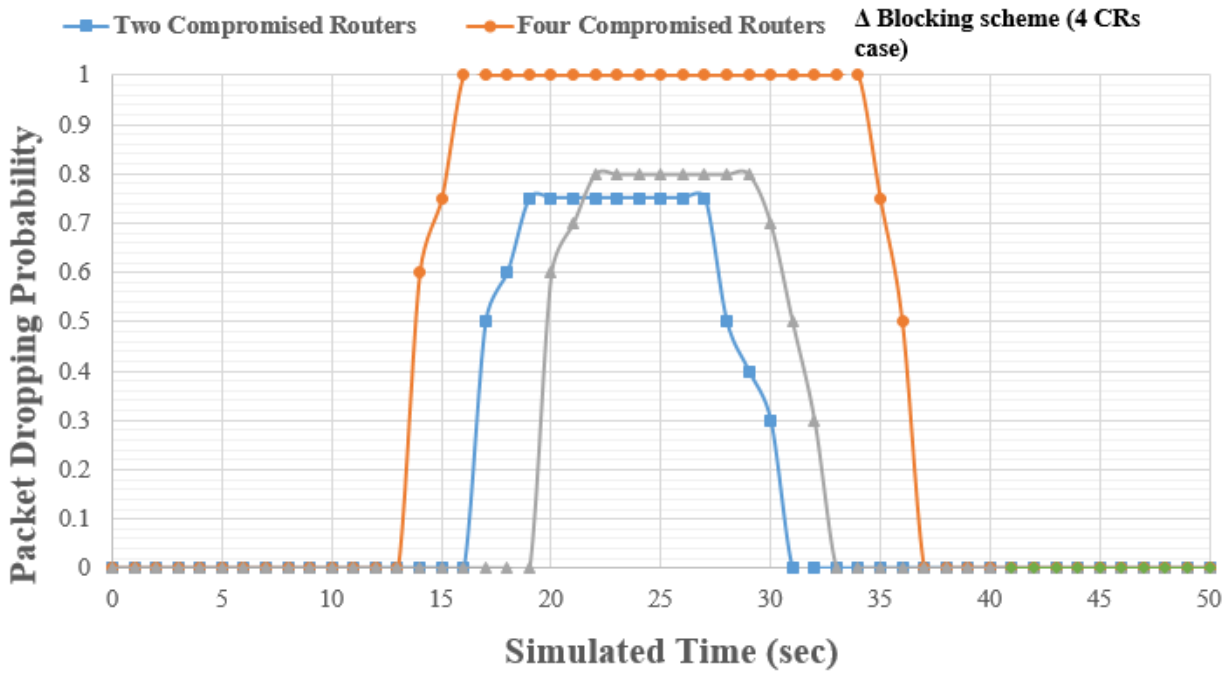

Figure 12. *Graph depicting the variation of Packet Dropping Probability with two pairs of CRs.*

The simulation can also be described with the bipartite graph G = (C ∪ P, E) where the vertex set is composed of two parts Consumer set C and Producer set P and every edge has one end point in C and the other in P. That is there is no edge with both its endpoints in C or P. The set R depicts the various combinations of the intermediate paths (of nodes) that can be traversed to reach from the consumer set to the producer set. A simple case of this implementation is depicted in Fig 13.

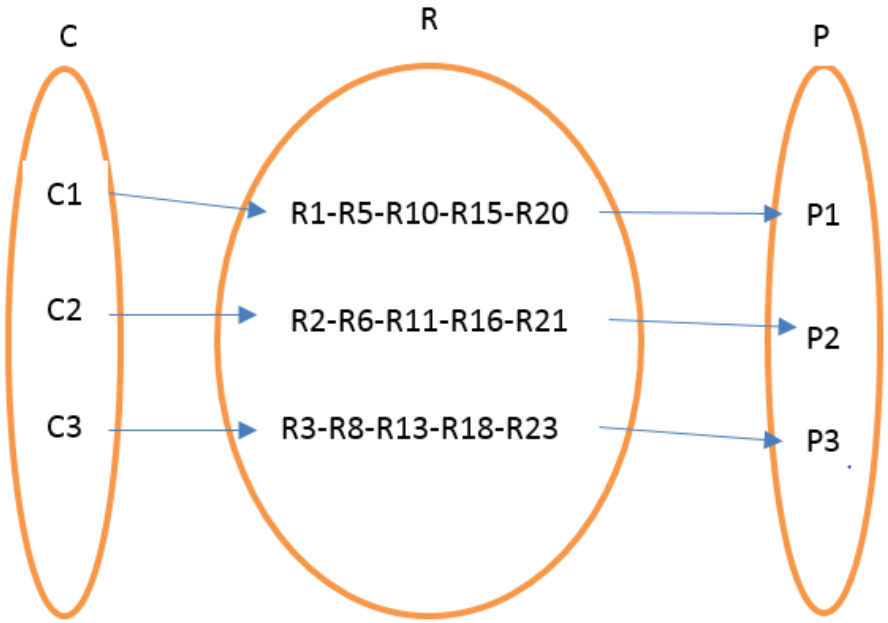

Figure 13. *Bipartite Graph connecting Consumer and Producer*

We can generalize this scenario to multiple consumers which can retrieve content from any producer listening on the corresponding name space as depicted in Fig 14.

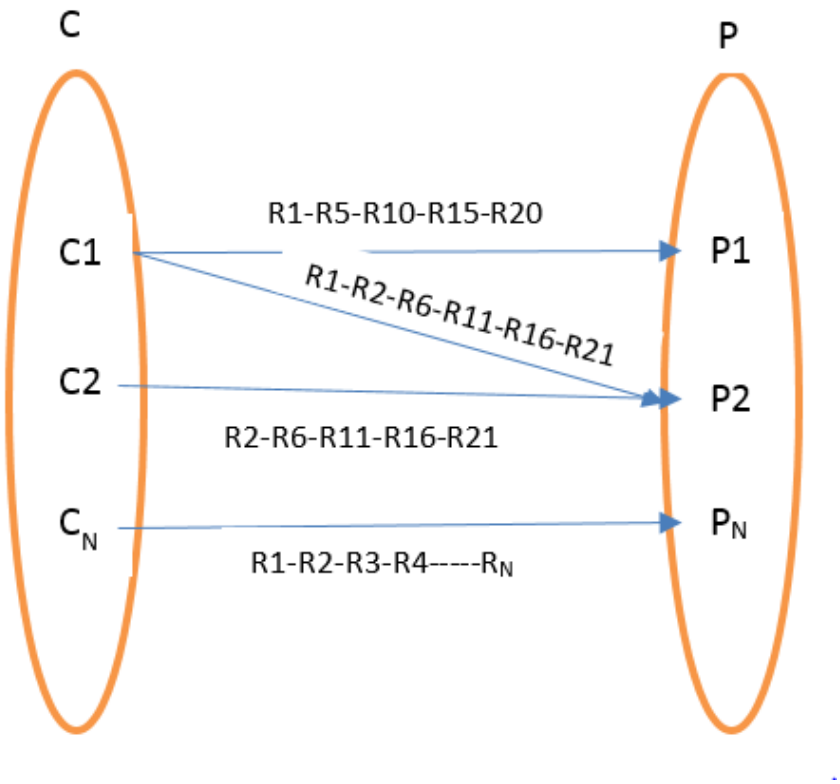

Figure 14. *Bipartite Graph extended to n-nodes*

The above results clearly depict the extent of the impact the forwarding loops have on the Content Centric Networks. Various techniques discussed in Chapter 5, aid in mitigating these attacks. Occurrence of these attacks in other novel networking protocols can be further explored in the future.

CHAPTER 6

CONCLUSION AND FUTURE WORK

In this paper we have described how malicious customers can attack the availability of Content Centric Networks (CCNs) by creating forwarding loops inside one CCN cluster or across multiple Clusters. We have discussed how the forwarding loops cause one request to be processed repeatedly or even indefinitely, resulting in undesired resource consumption and potential Denial-of-Service attacks. We have discussed the phases of these loop attacks in Static CCN systems and extended it to the mobile CCN systems. Next, we proposed a few detection and mitigation techniques that will allow routers to identify and prevent the formation of such loops. Finally, to evaluate the practicality of such forwarding-loop attacks, we used ndnSIM to simulate these loops in various scenarios discussed earlier.

We have observed that the Forwarding loops have a considerable impact on the normal services provided by the CCN. Future work could focus on more efficient proposals for the detection and mitigation strategies of these loops. Forwarding loop attacks can also be evaluated further in other scenarios of ICN in the future.

REFERENCES

[1] A. Feldmann, “Internet clean-slate design: What and why?” ACM SIGCOMM Computer Commun. Review, vol. 37, 2007, pp. 59-64.

[2] J. J. Garcia-Luna-Aceves, "Eliminating undetected interest looping in content-centric networks," Network of the Future (NOF), 2015 6th International Conference on the, Montreal, QC, 2015, pp. 1-6.

[3] V. Jacobson, D. Smetters, J. Thornton, M. Plass, N. Briggs, and R.Braynard, “Networking named content,” Proc. CoNEXT, Rome, Italy, Dec. 2009, pp. 11-18.

[4] H. Yuan and P. Crowley, “Scalable pending interest table design: From principles to practice,” Proc. IEEE INFOCOM, 2014, pp. 2049-2057.

[5] S. Arianfar, T. Koponen, B. Raghavan, and S. Shenker, “On preserving privacy in content-oriented networks,” Proc. ACM SIGCOMM Work- shop on Information-Centric Networking, 2011, pp. 19-24.

[6] C. Ghali, G. Tsudik, and E. Uzun, “Network-layer trust in named-data networking,” ACM SIGCOMM Computer Communication Review, vol.44, no. 5, 2014, pp. 12-19.

[7] Vassilios G. Vassilakis, Ioannis D. Moscholios, Bashar A. Alohali, Michael D. Logothetis “Mitigating Distributed Denial-of-Service Attacks in Named Data Networking,” The Eleventh Advanced International Conference on Telecommunications, AICT, 2015

[8] S. Choi, K. Kim, S. Kim, and B. Roh, “Threat of DoS by interest flooding attack in content-centric networking,” Proc. IEEE International Conference on Information Networking (ICOIN), 2013, pp. 315-319.

[9] F. Li, F. Chen, J. Wu, and H. Xie, “Longest prefix lookup in named data networking: How fast can it be?,” Proc. 9th IEEE NAS, 2014, pp.186-190.

[10] Amadeo, M., Campolo, C., & Molinaro, A. (2014). Forwarding strategies in named data wireless ad hoc networks: Design and evaluation. Journal of Network and Computer Applications, doi:10.1016/j.jnca.2014.06.007