

GATE and FENCE: Geo-Blocking Protocols for Named Data Networking

Alberto Compagno
Cisco Systems, Paris, France
alberto.compagno@cisco.com

Mauro Conti, Stefano Munari
University of Padua, Padua, Italy
conti@math.unipd.it,
stefanomunari.sm@gmail.com

Riccardo Lazzeretti
Sapienza University of Rome, Rome, Italy
lazzeretti@diag.uniroma1.it

Abstract—Named Data Networking (NDN) is a novel Internet architecture which focuses on content distribution by exploiting in-network caching and name-based forwarding. Contrary to today's Internet, NDN has been designed from the ground up to be secure. From a content provider perspective (e.g., YouTube, Netflix), NDN offers appealing advantages in terms of network load and traffic reduction at producer side through in-network requests aggregation and content caching. As a side effect, content providers lose control on content dissemination when consumers' requests are aggregated or satisfied by the network. This hinders the correct application of copyright and licensing agreements: only specific regions are allowed to consume a subset of the distributed contents. In attempt to address this problem, the existing TCP/IP approaches exploit requests' source addresses (at server side) to identify the geographic origin of each request. In NDN these solutions are unfeasible for two reasons: consumers' requests do not carry any source address, and a request will never reach content providers when aggregated or satisfied in the network. We solve this problem by proposing two lightweight and distributed geo-blocking protocols (GATE and FENCE) which use packet marking to identify and validate network regions at network edges. We perform experiments both on a network simulator and by extending the NDN implementation. Through our results we prove the proposed protocols are feasible, i.e., all the regions blacklisted by content providers are blocked and their network costs, in terms of space and router processing overhead, are negligible.

Index Terms—Named Data Networking, Content-Centric Networking, Geo-blocking, Content distribution, Content provider, Packet marking

I. INTRODUCTION

In the last decade, content distribution has exponentially grown until covering more than half of the overall Internet traffic [1]: Facebook, Netflix and Youtube are just few examples of such content-based services. Nevertheless, the Internet architecture has not been able to evolve quickly enough to embrace this paradigm shift. As a consequence, two main approaches have been proposed to address its shortcomings [2]: (1) building a content-centric communication above the transport layer, e.g., Content Delivery Network [3], and (2) replacing TCP/IP with a different paradigm, i.e., Information Centric Networking [4].

Information Centric Networking (ICN) is a networking paradigm which has been proposed as an alternative to the current IP-based Internet architecture. Unlike IP, which explicitly refers hosts in datagram packets, ICN addresses named-contents by promoting data (rather than hosts) as first-class network-layer entities. Moreover, ICN focuses on mobility by decou-

pling publishers from their contents, privacy by decoupling senders from receivers, and security by enabling the sender to embed its trust directly in the content and fully supporting in-network filtering mechanisms. Named Data Networking (NDN) is one among several projects which implement the ICN principles (e.g., CICN [5], MobilityFirst [6], NetInf [7]). In NDN, communications are named-based and reflect the pull-model: a consumer can retrieve a content by issuing a request, namely an interest, which contains the content name; the interest is then forwarded up to a producer able to satisfy it. To facilitate efficient content distribution, NDN introduces in-network caching. It enables routers to directly satisfy incoming interests if a copy of the requested content exists in their cache, thus reducing communication latency and network traffic.

Content providers (CPs) of video on demand and streaming services are prominent beneficiaries of this new architecture. Nevertheless, they have to abide by copyright and licensing agreements established with the content owners. Often, such agreements restrict a CP from serving a content (e.g., movie, web resource) in specific areas [8]. An example is given by Netflix which maintains different catalogs based on the region where the user resides¹. Current mechanisms restrict the distribution of geo-blocked contents when the geographical position of the end user is detected. A common solution consists in filtering unauthorized regions by dropping their requests. In IP-based networks, it is possible to identify two main geo-location approaches: (1) using global positioning system (GPS) and (2) exploiting the consumer IP address as a geographic locator. While the former is barely used due to the GPS hardware dependency, the latter provides a convenient way, even though not infallible [9, 10], to apply restrictions. Unfortunately, the lack of consumers' addresses (or identifiers) in ICN packets prevents the adoption of IP-like solutions [11].

In this paper we address the problem of geo-blocking in NDN, i.e., restricting content distribution only to consumers connected to specific network segments/regions. To the best of our knowledge, this is the first work which *directly* address this problem by providing a lightweight solution. Indeed, state-of-the-art works focus on access control [12, 13, 14] while our proposal specifically targets geo-blocking: former solutions cover a different problem (i.e., generic access control) implying

¹Netflix Terms of Use: <https://help.netflix.com/legal/termsofuse>

significant deployment, storage and performances overhead respect to our solution which is orthogonal to access control. Our contribution consists in a distributed, lightweight and optional policy which can be leveraged by CPs to control geographical content distribution in the NDN network. *With our solution, interest and/or data packets can be dropped according to geographical properties, resulting in simple filtering mechanisms that can be placed over existing access control mechanisms.* To this aim we design, implement and evaluate two protocols: Geo-blocking At The Edge (GATE) and Flexible Name Code (FENCE). GATE is a lightweight and fully-decentralized protocol based on a list of region-id markers released by producers and checked by ISP edge routers. Instead, FENCE takes advantage of the existing packet name to distribute the markers in the network. Both solutions introduce low space overhead in the packets guaranteeing 100% accuracy, i.e., all blacklisted regions are effectively blocked and outperforming existing TCP/IP filtering-based counterparts. In terms of average overhead in upstream packets processing we measure $26.95 \mu s$ for GATE and $99.78 \mu s$ for FENCE, whilst for downstream processing GATE requires only $0.18 \mu s$ and FENCE does not introduce any overhead.

II. NDN ARCHITECTURE

NDN [15] supports two kinds of packets: *interest*, which contain a request for a named content, and *data*, the requested named content. These packets are exchanged between three types of NDN entities: *consumers* which issue interests for contents, *producers* which create, digitally sign and publish contents in the network, and *routers* which forward interests and data packets, and cache both of them. Each entity maintains three data structures: a *Content Store (CS)* used to cache forwarded data packets, a *Forwarding Information Base (FIB)* which is a forwarding table populated by a name-prefix based routing protocol, a *Pending Interest Table (PIT)* indexed on pending interests' names. Each PIT entry has a set of incoming and outgoing interfaces where multiple interests which reference the same content are collapsed under the same entry.

In the upstream flow of the NDN forwarding process (Fig. 1a) (1) A consumer issues an interest to a router which checks its CS for a corresponding content: in case of cache hit it replies with the content, otherwise it performs a lookup in its PIT; (2) If a PIT hit occurs, it aggregates the interest with the matching entry, otherwise it adds a new entry to the table. In case of a PIT miss, the router forwards the interest to the next hop entity according to its FIB; (3) On selecting the next hop, the router applies a forwarding strategy to the result of the FIB lookup: it may decide to drop an interest in certain situations, e.g., all upstream links are congested or the interest is suspected to be part of a DoS attack. Finally, the packet reaches the producer which replies with the corresponding digitally signed content or with an application-level NACK.

Data packets follow the reverse path of the interest and their forwarding process can be split into the following steps (Fig. 1b). (4) A router checks the data packet name against its PIT: if a match occurs the corresponding entry is removed from

the PIT and the content is stored in the CS (5), otherwise the content is dropped; (6) Finally the content is forwarded on all the down-stream faces stored in the matched PIT entry.

Instead of referring an host explicitly, an NDN consumer refers to a named content. In NDN, names are hierarchically structured, similarly to URLs (e.g., `/ndn/games/chess`), and opaque to the network [16]. Hence, the naming scheme can evolve independently from the network allowing each application to choose the one that best fits its needs.




Fig. 1: Data forwarding in NDN

III. RELATED WORK

To the best of our knowledge, this is the first work in ICN on geo-blocking. Some path identification mechanisms have been recently proposed [17, 18], although all of them focus on traceroute and ping utilities and are ineffective for geo-blocking. As for path identification, even traceback mechanisms are available in the ICN literature but useless in the context of geo-blocking. Dai et al. [11] provide an *interest traceback* mechanism to mitigate DDoS attacks in NDN. When the network detects any spoofed name, it notifies the related edge router through the interest's reverse path so the latter subsequently limits the attacker's interest rate. Despite this solution limits DDoS, it cannot be applied to the geo-blocking problem because it does not enable the CP to identify malicious users, who ask for contents unauthorized in their area.

There is plenty of NDN literature on access control [19]. These works can be seen as complementary to ours: each of them focuses in denying the consumer from *accessing* the data (e.g., through encryption) while we prevent the consumer to even *receive* the data. For example, Li et al. [13] propose LIVE, a signature-based and centralized access control scheme for ICN. The CP generates per-content public and private tokens and assigns them to different network entities according

to their access level. The private token enables legitimate entities to access the published content while content integrity verification is done through the public token. This scheme reveals scalability issues due to its tokens-storage and token-synchronization overhead (in order to receive fresh tokens NDN entities interact frequently with CPs). LIVE shows an average communication overhead between 15% and 50% while the overhead range of our solution is between 0% and 4.5%.

IV. SYSTEM AND THREAT MODEL

In this section we introduce our system (Section IV-A) and the threat model (Section IV-B) which respectively define the core characteristics of the environment on which our protocols operate and the possible types of attackers.

A. System Model

In our model we consider a set of CPs which own several NDN producers. From the CPs' perspective, each producer is a data supplier which publishes popular contents, such as a live video streaming of a Major League Baseball match. For each CP, there exists a set of consumers which request the same sequence of content objects, e.g., the same popular video.

GATE and FENCE are applied only at the edges of the network by Internet Service Provider (ISP) edge routers [20], producers and consumers. Thus, creating an overlay network with respect to all the other nodes which act transparently to it (Fig. 2). In our model we assume that the home router of a subscriber belongs to the same network region of the ISP edge router. For example, the network traffic of an Italian home router directed to a French recipient will always pass through an Italian ISP edge router before reaching the French network.

We assume that ISPs are trusted and apply our protocols. We believe that by agreeing to implement GATE and FENCE on their edge routers, ISPs will create a favorable situation for them in which consumers' requests can be aggregated and satisfied by the network, thus reducing the overall traffic load on their infrastructures. If ISPs choose to not implement GATE or FENCE, CPs will likely adopt solutions that prevent content caching and request aggregation (e.g., by using TLS-like connections), resulting in a greater overhead for ISPs. In Section VI, we show that our solution has almost zero impact on the processing performances of a router which is a further incentive to persuade ISPs to adopt it.

In our model NDN entities need to know the available region identifiers, so each entity involved in the protocol adheres to a common *region-id schema* (not necessarily statically defined). This can be easily achieved by using an existing standard schema. Unless otherwise stated, each router caches content packets and applies the best route forwarding strategy, as in the default NDN routers settings [21].

B. Threat Model

In our analysis we assume the following:

1) *Routers are not widely compromised*: we consider unfeasible for a malicious adversary to compromise part of the network routers. In particular, we assume core, border and ISP edge

routers [20] as trusted entities, while subscriber home routers can be compromised. Even if a very powerful adversary (e.g., NSA) could compromise the former set of routers, it would not fit our model because it would have no reason to put effort in breaking GATE or FENCE. Indeed, geo-blocking protocols mainly concern CPs and their end users.

2) *Producers are trusted*: an NDN producer which redistributes a content in a blocked network region becomes effectively a producer of a new content. This is due to the NDN protocol, which mandates to digitally sign each content. Also, we do not consider legitimate producers owned by a CP and compromised by an external adversary as part of our model.

3) *Consumers are not trusted*: *E* is an evil user who wants to circumvent the geo-blocking restrictions defined by the CP on the distributed contents.

- *E can fake her location* before sending the interest in the network. For example, if *E* is in UK and wants to access a content distributed only in IT, she can claim to be in Italy by crafting a fake interest. In Section V we discuss how our protocols proactively prevent this kind of attack.
- *E can use a Virtual Private Network (VPN) or a proxy server* to act as a consumer who is connecting from an authorized region. In this scenario the CP is responsible of monitoring the VPN providers and blocking them. This attack exists also in the current Internet architecture [22] where the aforementioned approach is applied.




Fig. 2: Logical view of GATE overlay network

V. GEO-BLOCKING PROTOCOLS

In Section V-A we introduce GATE while in Section V-B we describe FENCE. Both of them, with the help of the ISP edge routers, enable an NDN CP to restrict the content access only to consumers of selected network regions. The proposals differ in the required processing and space overhead, but they

can coexist simultaneously in the same network so that a CP can choose which one to apply for each of its contents.

To identify the region in which a consumer resides, we define *region-id* markers, which are inserted in NDN packets at the edge of the network. Thus, we strive to preserve the end-to-end principle [23] expressed both by NDN and TCP/IP. In particular, we choose the *ISP edge router* to perform the marker verification instead of core, border, subscriber home router for the following reasons: (1) First of all, *any consumer's data pass through a ISP edge router before reaching a producer*. Indeed, any consumer has a contract with an ISP whose edge router acts as checkpoint for the licensed contents. On the contrary, this is not true for core or border routers. (2) The ISP edge router is *the first network entity* which can easily verify the *region-id* of the consumer and consequently apply geo-blocking: indeed, both the consumer and the router belong to the same network region. This approach limits the number of packets exchanged between the NDN entities to locate the consumer's region and, similarly to [11], the ISP edge router can immediately block unauthorized traffic generated by a cheating consumer.

As mentioned in Section IV, each entity who participates in the protocol shares a schema of region identifiers. For example, by adopting the ISO 3166-1 numeric standard for country codes [24] the geo-blocking protocols can be applied on a global scale with each *region-id* marker identifying a different country.

A. GATE

Fig. 3 describes the GATE protocol. Interest forwarding is performed in 6 steps. (1) First of all, the CP's application executing on the consumer's node inserts its region-id marker in the corresponding field of the interest packet. (2) The CP's application issues an interest. (3) If the interest has a valid name, i.e., there exists an entry for its prefix in the router's FIB, the first ISP edge router applies a forwarding strategy which comprises GATE as internal step (Fig. 4a). The router checks if the interest belongs to an upstream face in use by the consumer and if the *Interest.RegionId* field is set. If not, the policy is not applied. Otherwise, the router also verifies if the interest contains a correct region-id by comparing the latter with its own. If not, the packet is dropped. (4) The interest is forwarded according to the chosen NDN forwarding strategy and possibly reaches another region, different from the source region. (5) Finally, the interest reaches a producer owned by a CP or a router which has the requested content in its CS.




Fig. 3: GATE protocol




Fig. 4: GATE - Packet processing. Optional step of a router's forwarding strategy

The downstream flow of a geo-blocked data packet is described in the following (dashed arrows in Fig. 3). (6) The producer performs a lookup of *Interest.RegionId* in her list of unauthorized regions. A cheating consumer who does not set the *Interest.RegionId* field is directly blocked by the producer. (7) According to the comparison result: if the *Interest.RegionId* is in an *authorized region* the producer replies with a data packet and a region list², which is embedded as a field in the data packet, otherwise it replies with an application level NACK [25]. (8) The data packet is forwarded following the reverse path of the interest until reaching an ISP edge router. (9) The ISP edge router checks if the content is subject to geo-blocking by observing if it is intended to a consumer face and if the *Data.RegionList* field is set (Fig. 4b). If at least one condition is not satisfied, the content can be propagated to the next entity (another router or the consumer), otherwise the edge router checks if the content can be delivered. It performs a lookup on the region list embedded in the data packet using its region-id as list index. Depending on the comparison outcome, the router applies the next forwarding step if it is in an *authorized region*, otherwise it drops the data packet. (10) If not dropped, the data packet is eventually delivered to the consumer. Each region list is specified by the CP and associated with a set of one or more contents. So, *a region list is actually a blacklist of regions* which enables the edge routers to prevent the malicious behavior discussed in Section IV-B.

If the content is cached by some router along the forwarding path, then a consumer can retrieve it directly from one of them. In this case, the region list enables the ISP edge router

²We refer the term "list" in its generic acceptance, it does not reflect the data structure used for the region list implementation.

to enforce the geo-blocking policy (Fig. 4b and step 9 in downstream flow of Fig. 3). Indeed, the region list has been previously inserted in the data packet by the producer, before issuing the content in the network, so that each geo-blocked data packet contains the correct region list. In its processing steps, the ISP edge router (Fig. 4) has to check if the packets are related to consumer's faces to distinguish between standard NDN packets and geo-blocked packets, which are processed before being possibly delivered.

B. FENCE

Flexible Name Code (FENCE) is an alternative lightweight solution to GATE. It leverages the packet's name to store the region-id marker instead of extending the NDN packet specification like GATE does.

Basically, the CP's application running on the consumer node appends its region-id to the `Interest.Name` field (step 1 in Fig. 3) while the ISP edge router drops all the interests with a wrong region-id³ (step 3 in Fig. 3). Hence, the producer verifies only if the consumer belongs to an authorized region by matching the region-id contained in `Interest.Name` with its local region list (step 6 in Fig. 3). Other protocol steps are equivalent to the one already discussed for GATE (Section V-A) except for the following, which concerns the downstream flow: in step 7 (Fig. 3) the `Data.RegionList` field does not exist and step 9 (Fig. 3) is not applied by FENCE because all the data packets have been distributed under a valid name which contains only allowed region-ids. Indeed, unauthorized consumers are proactively geo-blocked by the producers which discard all of their interests.

VI. EVALUATION

The simulation environments along with the experiments are described in Section VI-A. The related results are discussed in Section VI-B. We single out the following metrics to evaluate our implementation:

- *Router processing overhead* (μ s): the per-router processing time overhead introduced by our protocols.
- *Per-packet space overhead* (B): the space overhead introduced for each packet.
- *Goodput* (kbps): the application-level throughput from the consumer's perspective. We purposely exclude the header overhead because the packet specification of GATE differs from the NDN standard.
- *RTT* (ms): the time between the issuing of a consumer's request and the receiving of the corresponding content object (data packet).

A. Simulations

We simulate network scenarios through tools widely adopted by the NDN and the research communities [26, 25, 27]: we use *ns-3*⁴, an open-source discrete-event network simulator, as base framework and *ndnSIM* 2.3 [28], an NDN simulator (based

on *ns-3*) which provides all the NDN routing and forwarding strategies along with routers' building components. In addition to the simulators, we also configure the NDN stack extended with GATE and FENCE on 6 virtual machines. Each virtual machine has been configured with 1 virtual processor, up to 100% usage capacity and 1.1GB of RAM. To extract different metrics out of our implementation we consider 2 types of topologies. A set of simple topologies, which we deploy locally as an overlay network of virtual machines, and a large one, which we deploy on the *ndnSIM* network simulator.

a) *Simple Topologies*: To preliminary study GATE and FENCE, we test our work on the following topologies.

- *Line topology*: a topology of 6 nodes disposed linearly (Fig. 5).




Fig. 5: Line topology

- *Tree topology*: it represents a simple use case with the producers at the root of the tree and a consumer associated to each child node (Fig. 6). In this topology, the CP geo-blocks the C4 consumer and there is a time drift between two parts of the network.




Fig. 6: Tree topology

b) *AT&T Topology*: This topology and the network configuration parameters are taken from Rocketfuel ISP mapping engine [29] measurements. To better simulate a popular CP scenario, we extend the original topology with additional edge nodes and use them as consumers and producers (Fig. 7). Indeed, the original topology has only few edge nodes while a popular provider like Netflix has over 50 million subscribers in the US (we scale them to the proportions of our simulated

³Thwarting cheating consumers.

⁴<https://www.nsnam.org/>

network). Although the mapping subscriber to network node is not one-to-one, we can expect a lot of consumer nodes involved. Fig. 7 depicts a network of 864 nodes with the 202 red nodes split in 50 producers ($\sim 25\%$) and 152 consumers ($\sim 75\%$).




Fig. 7: AT&T topology

c) Scenarios: In *Scenario A*, the line topology (Fig. 5) is configured with all the routers' content store (CS) set to 0 and the consumer requesting 50000 content objects at a constant rate of 200 interests/s. Since the consumer is located in an authorized region, the producer at the other edge of the topology always replies with the requested content.

In *Scenario B1*, the tree topology (Fig. 6) has the consumers of the left branch (C0 and C2) which apply GATE and the consumers of the right branch (C1 and C3) which apply FENCE. In the example, C2 and C3 request the content 123 respectively with the names P0/movie/123/ and P1/movie/123/C. Therefore, to balance the load of requests, we configure two producers: one serving the left subtree (e.g., P0/movie/123), and the other the right one (e.g., P1/movie/123/C). Also, we set the CS size of routers belonging to Region C to 5000. This favors the consumers of Region D, which start executing with a drift⁵ of 10 s respect to all the other nodes (the timezone drift of Fig. 6 is scaled). The *Scenario B2* is the same as B1 except for the consumer in Region B, in B1 it applies GATE while in B2 it applies FENCE (Section V-B).

Since United States (US) has the largest deployment of most popular CPs, we consider its map along with the *AT&T topology* to test geo-blocking (*Scenario C*). We deal with two CPs, namely *N* and *H*, who control several producers located in different US regions. Each US consumer requests a content which corresponds to a sequence of video chunks (different content objects) distributed by *N* or *H*. The consumers represented in Fig. 7 belong to specific US regions which depends upon the ISP edge routers' positions in the topology; the latter were given in the original topology configuration. Throughout the experiments we consider the 18 given regions

to apply GATE, deciding to block different sets of them. The blocking percentages are computed considering the 18 regions as 100%, with this percentage no consumer will ever receive the requested content. To synthesize the general behavior of the protocols we define the following blocking percentages: 27.78% blocks 5 out of 18 regions, 50% blocks 9 out of 18 regions, 77.78% blocks 14 out of 18 regions and 94.44% blocks 17 out of 18 regions. The last percentage represents the case of a CP which restricts the content distribution to a single region.

B. Results

We use *Scenario A* to measure the *router processing overhead*: blocking or authorizing consumers does not affect our measurements because the router would anyway process geo-blocked named packets by applying a geo-blocking policy. The experiments are carried out by tracing the computational overhead (μs) required by each router to perform the specific operations defined by the geo-blocking policies. In particular, we measure two distinct overheads: one related to the processing of interest packets (upstream traffic) and one related to the processing of data packets (downstream traffic).

We execute the experiment 30 times computing the mean, standard deviation, minimum and maximum value (Table I). GATE outperforms FENCE on all the parameters: it is more than 3x faster than FENCE as well as more stable. It may be due to the fact that FENCE extracts and manipulates a substring while GATE directly accesses a packet field. The difference between the two protocols for the downstream traffic is negligible: GATE requires less than 1 μs .

We compute the *space overhead* for both interest and data packets considering the NDN packet implementation extended with GATE's fields. With FENCE the resulting overhead is 1 byte both for interest and data packet: the region-id is a component of the packet name encoded with Type-Length-Value format (a nonNegativeInteger type with length 1 and value of size 1). With GATE the resulting overhead consists in 1 byte required by the RegionId field encoded in the interest packet and 32 bytes to encode the RegionList in the data packet. The data packet overhead is due to the region list, which needs a flag (a bit) for each of the 249 regions encoded in the ISO standard [24].

To estimate the *average goodput at the consumer side*, we perform 15 experiments for each Scenario (B1 and B2), each of them lasting 20 minutes. Fig. 8a reports the results of Scenario B1: we compare the two subtrees, the one which applies GATE with the one which applies FENCE. Since the average goodput of C0 (red line) partially overlaps with the one of C2 (blue line) we gather that the cache of C2 is partially exploited by

TABLE I: Router processing overhead - Scenario A

Traffic	Protocol	Avg (μs)	SD (μs)	Min (μs)	Max (μs)
Upstream	GATE	26.95	28.46	12.72	41.48
	FENCE	99.78	72.22	63.66	135.89
Downstream	GATE	0.18	0.38	0.00	1.00
	FENCE	-	-	-	-

⁵The drift indicates the time elapsed since the beginning of the experiment.


C0. About that, the average goodput of C1 (green line) is less than the one measured on C3 (black line): this is explained by the highest number of hops which separates C1 and the producer P1. Indeed, C1 can not exploit the cache of C3 because it belongs to a different region and it uses FENCE as forwarding strategy. In this scenario GATE outperforms FENCE. C4 is not reported in Fig. 8a because it is a blocked consumer which only receives NACKs. In terms of standard deviation both C2 and C4 experience between 51 and 53 kbps while C1 and C3 measure respectively 36 and 49 kbps. The latter result shows

that FENCE is a little more stable than GATE.


The results of Scenario B2 (Fig. 8b) confirm the one obtained with Scenario B1. The gap between C1 and C3 is 127.85 kbps while the difference between the average goodput of C0 and C2 is negligible. In terms of standard deviation, the most stable protocol is again FENCE. By observing both the scenarios, we conclude that (1) GATE and FENCE do not affect the other network flows: the different geo-blocking protocol applied to consumer C4 in the forwarding path of the other consumers does not reveal significant differences in the goodput estimated for these consumers; (2) geo-blocking is feasible: our protocols work as expected because P0 and P1 safely deliver their contents to the regions C, D while blocking B.

With Scenario C we perform 30 experiments applying GATE and 30 experiments applying FENCE for each blocking percentage, with each simulation lasting 60 seconds. Fig. 9 describes the average goodput estimated for each consumer of the network. From Fig. 9 we can see that a lower blocking percentage corresponds to an higher average goodput. A comparison between the two protocols reveals that GATE exploits the in-network caching feature of NDN better than FENCE. Indeed, with FENCE different regions request different contents because the region-id is appended by the consumer to the interest's name before issuing it. Conversely, with GATE the named object is the same for all the regions because the region-id is encapsulated in a specific field of the packet. The behavior of the two protocols overlaps with 94.44% blocking percentage, in-network caching is leveraged exactly in the same way by GATE and FENCE.

Table II summarizes the results about the average goodput and the standard deviation computed for the consumers of the AT&T topology. For 27.78% and 50% blocking percentages



(a) Scenario B1



(b) Scenario B2

Fig. 8: Average goodput. In the legend (n)hop refers to the distance in number n of hops from the producer while (m)CS indicates the content store size.




Fig. 9: Average goodput - AT&T topology. Each color represents a different blocking percentage.

TABLE II: Average goodput - AT&T topology

Blocking (%)	Protocol	Average (kbps)	SD (kbps)
27.78	GATE	2914.51	15.72
	FENCE	2643.44	13.02
50.00	GATE	1833.47	11.61
	FENCE	1646.93	8.70
77.78	GATE	894.09	4.71
	FENCE	780.95	4.70
94.44	GATE	102.78	1.88
	FENCE	102.50	1.89

TABLE III: RTT - Scenario A

Protocol	Average (ms)	SD (ms)
GATE	123.77	47.77
FENCE	123.89	47.82
NDN-Std	123.75	47.75

FENCE is more stable than GATE while for greater percentages the standard deviation is the same for both protocols. Considering the average goodput in terms of absolute values, the most significant gap between the protocols is represented by 271.07 kbps observed with 27.78% blocking percentage.

Finally, we measure the *RTT* in Scenario A by performing 30 simulations (10 for each protocol) lasting 20 minutes each. Table III shows that FENCE and GATE require nearly the same ms of the NDN best route strategy without introducing significant overhead from the consumer's perspective.

VII. CONCLUSION AND FUTURE WORK

This work focuses on geo-blocking content distribution in NDN: more specifically, we design two protocols which enable CPs to selectively restrict content access to specific network regions. The first protocol, Geo-blocking At The Edge (GATE), is a lightweight and fully distributed solution which leverages region-id markers inserted in named packets to exchange information about consumer's location between NDN network entities. The second protocol, Flexible Name Code (FENCE), directly exploits the existing packet name by appending its region-id marker to it. To the best of our knowledge, this is the first work which specifically addresses geo-blocking in NDN.

We conduct several experiments, observing that GATE requires little overhead in terms of packet size and router processing time. FENCE has shown to require less space occupation than GATE but more overhead in terms of packet processing. The latter needs 32 bytes to encode the whole region-list in a packet while the former uses only 1 byte to encode 1 region-id. About packet processing, our implementation of GATE requires on average less than 27 μ s, outperforming FENCE which costs around 100 μ s. Moreover, we verify through the simulations that GATE exploits in-network caching while FENCE can do it only partially. Although GATE is a more efficient solution than FENCE, the latter provides greater flexibility by encoding a marker-id in a name component. FENCE can thus be employed when the number of regions is not known a priori or if the latter is expected to be extended in the future.

As future work, we plan to improve the performances and the expressiveness perspectives of our solutions. In particular, we

will perform our experiments on a real distributed testbed [28] and on a faster ICN forwarder [5].

ACKNOWLEDGMENT

This work is partially supported by the EU TagItSmart! Project (agreement H2020-ICT30-2015-688061), the EU-India REACH Project (agreement ICI+/2014/342-896), the grant n. 2017-166478 (3696) from Cisco University Research Program Fund and Silicon Valley Community Foundation, and the Security and Privacy of Biometrics for Mobile Authentication (SPoB-MA) project (grant n. RM11715C7878B045 of Sapienza University of Rome Bando Ricerca 2017).

REFERENCES

- [1] "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2016-2021," <https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/complete-white-paper-c11-481360.pdf>, accessed: 2017-01-03.
- [2] Q. Jia, R. Xie, T. Huang, J. Liu, and Y. Liu, "The collaboration for content delivery and network infrastructures: A survey," *IEEE Access*, vol. 5, pp. 18 088–18 106, 2017.
- [3] A.-M. K. Pathan and R. Buyya, "A taxonomy and survey of content delivery networks," *Grid Computing and Distributed Systems Laboratory, University of Melbourne, Technical Report*, vol. 4, 2007.
- [4] G. Xylomenos, C. N. Ververdis, V. A. Siris, N. Fotiou, C. Tsilopoulos, X. Vasilakos, K. V. Katsaros, and G. C. Polyzos, "A survey of information-centric networking research," vol. 16, no. 2, pp. 1024–1049, 2014.
- [5] "Fd.io community icn (cicn) project," <https://wiki.fd.io/view/Cicn>.
- [6] I. Seskar, K. Nagaraja, S. Nelson, and D. Raychaudhuri, "MobilityFirst Future Internet Architecture Project," in *Proceedings of the 7th Asian Internet Engineering Conference*, ser. AINTEC '11, 2011, pp. 1–3.
- [7] C. Dannewitz, D. Kutscher, B. Ohlman, S. Farrell, B. Ahlgren, and H. Karl, "Network of information (netinf)—an information-centric networking architecture," *Computer Communications*, vol. 36, no. 7, pp. 721–735, 2013.
- [8] Q. Liu, R. Safavi-Naini, and N. P. Sheppard, "Digital rights management for content distribution," in *Proc. ACSW Frontiers*. Australian Computer Society, Inc., 2003.
- [9] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," in *Proc. USENIX Sec. Symp.*. USENIX Association, 2004.
- [10] Z. Lu, Z. Li, J. Yang, T. Xu, E. Zhai, Y. Liu, and C. Wilson, "Accessing Google Scholar under extreme Internet censorship: A legal avenue," in *ACM/FIP/USENIX Middleware*. ACM, 2017.
- [11] H. Dai and Y. Wang and J. Fan and B. Liu, "Mitigate DDoS attacks in NDN by interest traceback," in *IEEE INFOCOM Workshops*, 2013.
- [12] Q. Li, R. Sandhu, X. Zhang, and M. Xu, "Mandatory content access control for privacy protection in information centric networks," *IEEE Transactions on Dependable and Secure Computing*, vol. 14, no. 5, pp. 494–506, Sept 2017.
- [13] Q. Li, X. Zhang, Q. Zheng, R. Sandhu, and X. Fu, "Live: Lightweight integrity verification and content access control for named data networking," *IEEE Transactions on Information Forensics and Security*, vol. 10, no. 2, pp. 308–320, Feb 2015.
- [14] S. Misra, R. Tourani, F. Natividad, T. Mick, N. E. Majd, and H. Huang, "AccConF: An Access Control Framework for Leveraging In-network Cached Data in the ICN-Enabled Wireless Edge," *IEEE Transactions on Dependable and Secure Computing*, pp. 1–1, 2017.
- [15] L. Zhang, A. Afanasyev, J. Burke, V. Jacobson, k. claffy, P. Crowley, C. Papadopoulos, L. Wang, and B. Zhang, "Named data networking," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 3, 2014.
- [16] A. Ghodsi, T. Koponen, J. Rajahalme, P. Sarolahti, and S. Shenker, "Naming in content-oriented architectures," in *Proc. ACM ICN*. ACM, 2011.
- [17] R. Droms and D. Oran, "ICNRG S. Mastorakis Internet-Draft UCLA Intended status: Experimental J. Gibson I. Moiseenko Cisco Systems," 2017.
- [18] S. Khoussi, D. Pesavento, L. Benmohamed, and A. Battou, "Ndn-trace: A path tracing utility for named data networking," in *Proc. ACM ICN*. ACM, 2017.

- [19] R. Tourani, S. Misra, T. Mick, and G. Panwar, "Security, privacy, and access control in information-centric networking: A survey," *IEEE Communications Surveys Tutorials*, vol. PP, no. 99, pp. 1–1, 2017.
- [20] H. Berkowitz, E. E. Davies, S. Hares, P. Krishnaswamy, and M. Lepp, "Terminology for Benchmarking BGP Device Convergence in the Control Plane," Internet Requests for Comments, RFC Editor, RFC 4098, June 2005. [Online]. Available: <https://www.rfc-editor.org/info/rfc4098>
- [21] A. Afanasyev, J. Shi, B. Zhang, L. Zhang, I. Moiseenko, Y. Yu, W. Shang, Y. Huang, J. P. Abraham *et al.*, "NFD developer's guide," *Dept. Comput. Sci., Univ. California, Los Angeles, CA, USA, Tech. Rep. NDN-0021*, 2014.
- [22] J. A. Muir and P. C. V. Oorschot, "internet geolocation: Evasion and counterevasion."
- [23] J. H. Saltzer, D. P. Reed, and D. D. Clark, "End-to-end Arguments in System Design," *ACM Trans. Comput. Syst.*, vol. 2, no. 4, pp. 277–288, 1984.
- [24] ISO Standard 3166-1:2013, "Codes for the representation of names of countries and their subdivisions – Part 1: Country codes," International Organization for Standardization, Geneva, CH, Standard, Nov. 2013.
- [25] A. Compagno, M. Conti, C. Ghali, and G. Tsudik, "To NACK or Not to NACK? Negative Acknowledgments in Information-Centric Networking," in *IEEE ICCCN*, 2015, pp. 1–10.
- [26] A. Compagno, M. Conti, P. Gasti, and G. Tsudik, "Poseidon: Mitigating interest flooding DDoS attacks in Named Data Networking," in *IEEE LCN*, 2013.
- [27] C. Ghali, G. Tsudik, and E. Uzun, "Needle in a haystack: Mitigating content poisoning in named-data networking," in *Proceedings of NDSS Workshop on Security of Emerging Networking Technologies (SENT)*, 2014.
- [28] S. Mastorakis, A. Afanasyev, and L. Zhang, "On the Evolution of ndnSIM: An Open-Source Simulator for NDN Experimentation," *SIGCOMM Comput. Commun. Rev.*, vol. 47, no. 3, pp. 19–33, 2017.
- [29] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson, "Measuring ISP topologies with Rocketfuel," vol. 12, no. 1, 2004.