Analyzing Gateways' Impact on Caching for Micro CDNs based on CCN

César Bernardini[1] and Bruno Crispo[2,3]
[1]*DPS, University of Innsbruck, A6020, Innsbruck, Austria*
[2]*DISI, University of Trento, 38123, Povo, Italy*
[3]*DistriNet, KU Leuven, 3001, Leuven, Belgium*

Keywords: ICN, CCN, Content, Centric, Information, Caching, Cache, Gateways, Single Gateways, Multiple Gateways, CDN, Catalog.

Abstract: Content Centric Networking (CCN) is a new architecture for a future Internet. CCN is a clean-state architecture that targets the distribution of content. As such, content is located at the heart of the architecture and CCN includes two main features: communication led by names and caches everywhere. Nevertheless, CCN has been criticized due to the economical cost of replacing every IP router with a CCN router. As such, we assume that CCN will be used for small content delivery networks –Micro CDNs– located in the ISP facilities: it has already been shown that with only 100MB of caches in the ISP facilities, the ISP traffic to the Internet can be reduced by 25%. As a matter of fact, if CCN is deployed as a Micro CDN, gateways must exist to interconnect the CDN network with the Internet. In this paper, we study the advantages of using multiple gateways against a single gateway and its impact on the caching features. Our results show that multiple gateways are beneficial not only because they improve the performance of caches but also because the load of the network get distributed across several nodes.

1 INTRODUCTION

Today, users largely consume videos over the Internet. Indeed, the overall consumption of video streaming media accounts for 66% of all the Internet traffic (Cisco, 2013). New devices are already pushing these numbers forward. Smartphones, tablets, game consoles and Smart TV are all used for streaming video and by 2017 the number of smart devices will double. In particular, web-enabled TV will have a fourfold increase. Mobile video will increase 14-fold between 2013 and 2018 (Cisco, 2013). Therefore, the consumption of video streaming is expected to keep growing.

The traffic is already affecting the ISP capabilities. In fact, it has a twofold impact on the ISP: economical and networking aspects. With regards to economical aspects, today, many content providers provide services and access to content that depends on someone else infrastructure. Clearly, video on demand services such as YouTube or Vimeo could not be used without ISP infrastructures. ISPs expect that somebody pays for the extra traffic generated by these services. In the one hand, these content providers are not prone to pay for the dedicated infrastructures that they require. A clear example is the dispute between Verizon & AT&T in the USA or the French Free ISP and Google (http://blog.netflix.com/2014/04/the-case-against-isp-tolls.html; http://www.cnet.com/news/france-orders-internet-provider-to-stop-blocking-google-ads/). In the other hand, end users are not likely to pay more to solve this dispute. With regards to networking technology, Content Delivery Networks (CDNs) are used today to cope with large bulks of data being transmitted. CDNs are large storage network distributed across multiple datacenters, some of them located at the ISP facilities. Their goal is to serve content, applications, live-streaming media from nearby locations. The video on demand services has skyrocketed and the use of CDN has become a requirement for correct functioning of the services. However, every CDN company has implemented their own architectures and protocols. As a consequence, content delivery services require to hire multiple CDNs services which are incompatible between them. Furthermore, ISPs can not link together two similar contents being distributed through different CDNs. Thus, the construction of a new architecture targeting the content delivery is required to alleviate the charge of similar content being requested to the ISP.

Bernardini, C. and Crispo, B.
Analyzing Gateways' Impact on Caching for Micro CDNs based on CCN.
DOI: 10.5220/0005946400190026
In *Proceedings of the 13th International Joint Conference on e-Business and Telecommunications (ICETE 2016) - Volume 1: DCNET*, pages 19-26
ISBN: 978-989-758-196-0

Content Centric Networking has appeared as an emerging architecture for content delivery. Its communication paradigm based on names seems to fit better the current Internet needs: users are interested in content rather than its location, as IP does. The main features of CCN are communication based on names and in-network caching features everywhere. Especially, the in-network caching features appear as the masterpiece to solve the current problem of ISPs. It has been proven that with only 100MB of memory allocated into the ISP facilities, the traffic from the ISP to the Internet can be reduced by 25%. Adding 100GB of memory can reduce the load further by 35% (Imbrenda et al.,). This extra memory can be organized in particular locations of the ISP, organized as Micro CDNs using the CCN technology. As such, these Micro CDNs can help to achieve three goals: remove redundant traffic, increase the performance for end users and reduce the traffic generated towards the Internet.

Micro CDNs are expected to bridge CCN networks to the Internet. As such, many entry and exit points must exist between these two networks. We refer to entry and exit points as Gateways. These Gateways are used to transform packets from one architecture (CCN) to another (TCP/IP). However, the transformation of packets from one network to another is not the only problem that the gateways cause. The design of Single or Multiple Gateways may have a major impact on the performance of the caches of the Micro CDN network. In this paper, we study the difference between implementing Single or Multiple Gateways and the impact on the caching features. With extensive simulations, we show that Multiple Gateways are a better design choice because they not only provide a better caching performance for the overall network, but also they distribute uniformly the load across the network.

The rest of the paper is organized as follows. In the Section 2, we describe the CCN architecture and present a use case for the architecture: Micro CDNs. The section 3 give an insight into Single and Multiple Gateways and details the hypothesis treated in this paper. Then, the simulation environment and the results are presented in Sections 4 and 5. Subsequently, the related work on the area is presented in Section 6. Section 7 summarizes the contributions of this work and defines future work.

2 CCN OVERVIEW

2.1 Details of the Architecture

CCN is a clean-slate approach for the future Internet. The communication paradigm is led by names, known as (*Content Names*), which are contained in the *Interest* and *Data* messages. Communication is based on these two primitives: *Interest* and *Data*. The *Interest* expresses the will of a user for a content, while a *Data* message contains the answer for that content.

To address *Content Names*, every CCN node holds three tables:

- *Content Store (CS)* is a caching structure that stores content temporarily.
- *Pending Interest Table (PIT)* keeps track of the currently non-satisfied interests. Basically, it serves as a trace for the reverse path once the requested content is found.
- *Forwarding Interest Base (FIB)* is a routing table used to determine the interface to transmit content.

When an *Interest* arrives, the node extracts the *Interest* name and looks up into its CS for content stored that matches the full name. If the content is found, it is automatically sent back through the interface over which the Interest arrived and the *Interest* is discarded. Otherwise, we lookup the PIT to decide whether the CCN node is already waiting for the requested content. If an entry is found, the *Interest* is discarded. Otherwise, the FIB table is checked to decide the interface to redirect the *Interest* to. In the FIB table, a longest prefix matching operation is performed and the *Interest* is forwarded to another CCN node where same procedure is repeated until the requested content is eventually found.

There exists a reluctancy from the major network equipment companies such as Cisco or Huawei to implement the CCN technology at large scale. Implementing CCN at Internet scale will mean replacing every IP router by a CCN router, which indeed will have a major economic cost. As such, we argue that CCN will be implemented in smaller networks where there exists a clear economical benefit. For instance, ISPs can profit of this technology with content delivery by answering traffic locally instead of redirecting it through the Internet. In the following section, we revisit a use case for CCN into the ISP infrastructures that target content delivery: Micro CDNs.

2.2 Micro CDN: A Use Case for CCN

Micro CDN is a content delivery system. It is a network of caches that serve content to the clients of the network. Micro CDN offers two main distinctive features. First, in-network high-speed caching features. Second, a content based routing protocol based on names. Its features are achieved by implementing a CCN network.

Using CCN as the underlying architecture for Micro CDNs presents the following advantages:

- *Caches for Every Protocol/Application.* Currently, different CDN architectures target different applications or protocols. For instance, a CDN may offer services for live streaming of content or static content delivery. With the use of CCN, different protocols and applications share the same information at the networking layer. At networking layer, a *Content Name* specifies the content being transmitted instead of the protocol or any IP address of the server destination. It means that every caching protocol and application can benefit of the in-network caching features of CCN.
- *Fine-Granularity of Cached Content. Content Names* are hierarchically structured and made up of components. Among the components, CCN stores segmentation and version of the packet. It means that CCN handles packets at a chunk level. This can be beneficial for particular cases of content delivery. Such is the case for video on demand services, where first chunks must be retrieved faster to start playback than the rest of the chunks that can be incrementally downloaded.
- *No More HTTP.* Nowadays, most of the content is transmitted through HTTP pages. Although the HTTP protocol was not implemented for transmitting content. For example, video streaming services such as YouTube serve of HTTP and RTSP, another protocol derived from HTTP, to deliver the content. CCN permits to deliver content at the network layer and completely suppress the overhead caused by HTTP.

CCN based Micro CDNs may provide interesting features for ISPs and networks in general. However, there are still many parameters and studies to be carried out. In this paper, we focus one specific parameter of Micro CDNs: the Gateways and its impact on the caching features.

3 SINGLE GATEWAY OR MULTIPLE GATEWAYS

A Gateway is a network node interfaced to connect to other networks that use different communication protocols.

In our case, the Gateways connect the Micro CDN with the Internet. As Internet uses the TCP/IP communication model and the Micro CDN uses CCN, both networks are interfaced with Gateways. In a Micro CDN based on CCN, the content is not expected to be produced in the CCN network. The CCN network retrieves the information from the Internet. As such, every time a content cannot be found in a cache, it is demanded through a Gateway and subsequently retrieved from the Internet.

Interfacing two networks through Gateways is a challenging problem. In particular, mapping a CCN name into an IP packet is a complicated task and it depends on naming conventions (Shang et al., 2013). However, this is not the only interesting point. As CCN is a network of caches, different manners of implementing Gateways will affect drastically the performance of the Micro CDN. As a consequence, it is essential to determine the effects produced by using alternative Gateway designs to interconnect the CCN and Internet networks.

In the Figure 1, we have explicitly said similar requests are always retrieved from the same gateway. We assume that the CCN network has a routing protocol associated. In this case, it is Open Short Path First for Named data (OSPFN) (L. Wang, A. M. Hoque, C. Yi, A. Alyyan, and B. Zhang, 2012; Moy, 1998) OSPF is the default routing protocol for CCN. OSPF finds the shortest path towards the requested node. As such, similar requests are always to follow the same routing protocol and will be guided through the same nodes.

There exists an alternative manner to explain this problem: using catalogs instead of gateways. A catalog is a storage entity where all the demanded content is previously saved. Then, if we consider an autonomous CCN network without connection with external networks, the catalogs can be interpreted as the Gateways of the network. Single or Multiple Gateways can be translated into Single or Multiple catalogs. Thus, every request is answered with one of the catalogs. The problem becomes into assessing if having one single catalog or multiple catalogs distributed across the network are better choices.

For the analysis on the impact of Gateways into caching features, we consider two types of Gateway designs that are presented in the Figure 1: Single and Multiple Gateways. The Single Gateways are repre-

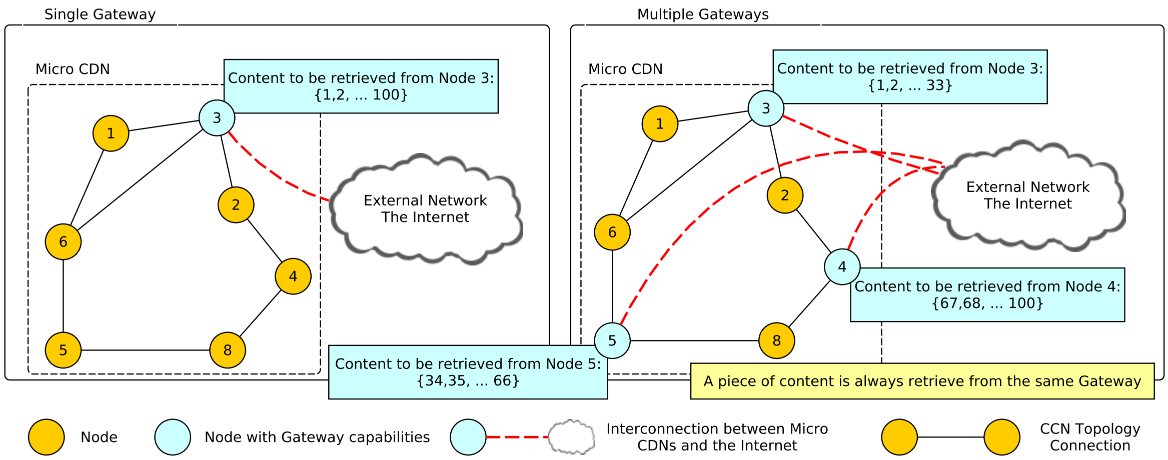

Figure 1: Micro CDN connected to the Internet using Single and Multiple Gateways.

sented on the left side of the Figure while the Multiple Gateways are on the right side. In both cases, a Micro CDN network using CCN is connected to some external network, the Internet. In the case of Single Gateways, the Micro CDN has an internal topology composed of many CCN nodes. However, only one of these nodes is a Gateway and is connected to the external network. This node is responsible for retrieving content from the Internet, when the content can not be found in the CCN network. In the case of Multiple Gateways, the Micro CDN has multiple CCN nodes that have access to the external nodes. As in the previous case, these nodes are responsible for retrieving content from the external network.

In the rest of the paper, we validate the following hypothesis:

1. Multiple Gateways achieve better caching performance than a Single Gateway.
2. Multiple Gateways distribute the load better across the network than Single Gateways.

4 SIMULATION ENVIRONMENT

In order to assess the hypothesis under a realistic and comprehensive simulation environment, we used SocialCCNSim, a discrete-event simulation tool written in Python, available at (https://github.com/mesarpe/socialccnsim,). This simulation tool supports all the parameters and scenarios we have presented throughout the paper.

Regarding the topology, we consider Micro CDNs as being part of the ISP infrastructure. The topological structure of the Internet is not expected to change and will be just like today's Internet. As a consequence, ISP topologies are the best candidates for evaluating our strategies. We choose four ISP level topologies and one tree topology: Abilene, Dtelecom, GEANT and Tiger (Rossi and Rossini, 2011), and a binary tree of 4-levels. We can see the topologies in the Figure 2.

In our scenario, requests are issued from 8 nodes in the network. These nodes are selected randomly from the network. Every node follows the Poisson process law with a frequency of 5Hz.

The content popularity model is a function that establishes the popularity of every piece of content, i.e., how often every single piece of content is requested. The content popularity is commonly modeled with a probability distribution function such as Zipf or MZipf (Rossi and Rossini, 2011). In the literature, the (M)Zipf α parameter ranges from 0.6 to 2.5. For instance, the catalog of the PirateBay is modeled with $\alpha = 0.75$, DailyMotion with $\alpha = 0.88$, while the VoD popularity in China exhibits a α parameter ranging from 0.65 to 1.0 (Fricker et al., 2012). In our case, we consider a broad range of α values from 0.65 to 2.0. Sometimes we refer to as popularity model with $\alpha = 1.1$

With regards to the configuration of the caches, we have selected Leave Copy Everywhere (LCE) as the caching strategy to organize the data stored in the network. Every time a *Data* message is transmitted in a CCN network, LCE leaves a copy in the cache of the CCN node. Least Recently Used was used as replacement policy, to determine what elements must be evicted from the caches.

The performance of caches is analyzed using the *Cache Hit* metric. When we analyze an individual cache, we report a Hit operation if an element is found in a cache otherwise a Miss operation. In a network of caches as CCN, we measure efficiency of caches

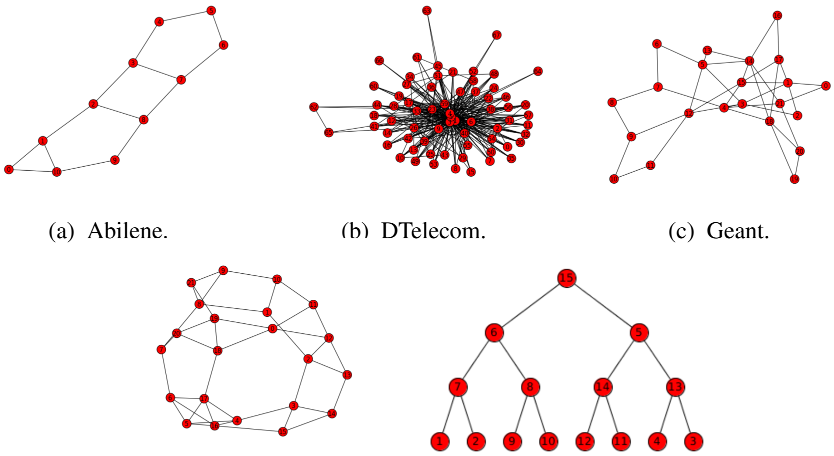

(a) Abilene. (b) DTelecom. (c) Geant.

(d) Tiger. (e) Tree.

Figure 2: Topologies used across the simulations: 4 ISP-level topologies and a 4-level binary tree.

with Cache Hit shown in Equation 1: $hits_i$ refers to the number of *Interest* messages answered by the cache of node i, while $miss_i$ the number of unanswered *Interest* messages. $|N|$ refers to the number of nodes in the topology.

$$CacheHit = \frac{\sum_{i=1}^{|N|} hits_i}{(\sum_{i=1}^{|N|} hits_i) + (\sum_{i=1}^{|N|} miss_i)} \quad (1)$$

Table 1: Simulation environment.

Parameters	
Number of requested content	10^6
Cache size Ratio	$\{10^{-6}; 10^{-5}; 10^{-4}; 10^{-3}\}$
Popularity Model	MZipf($\alpha = \{0.65; 1.1; 1.5; 2.0\}, \beta = 0$)
Topologies	Abilene, Tree, Geant Tiger, DTelecom
Caching Strategy	LCE
Replacement Policy	LRU
Request Model	Poisson Process
Request Placement	Uniform Probability Law
Gateways configurations	Single (random), Multiple Disjoint
Repetitions	3
Simulated Time	86,400 seconds
Routing Algorithm	Open Shortest Path First

5 RESULTS

Hypothesis #1: Multiple Gateways achieve better caching performance than Single Gateways

Our intuition is that Multiple Gateways can achieve better performance in terms of Cache Hit than Single Gateways. To assess this hypothesis, we resort to compare performance of Cache Hit in several scenarios. We vary topologies, popularity scenarios and cache sizes. We expect to find that caching performance of the overall network is always better with Multiple Gateways than with Single Gateways.

In the Figure 3, we present the results of our comparison of Multiple Gateways and Single Gateways after 480 experiments. In the x-axis, we show the Cache Size Ratio while the y-axis shows the Cache Hit results. The blue line represents the Disjoint Multiple Gateways while the red line refers to the Single Gateway. Every line of charts corresponds to one topology while every row corresponds a different popularity model.

With a first look on the charts, we can see that the blue line always surpasses the red line: it means that the Cache Hit Ratio for Multiple Gateways is always better than the ratio for Single Gateways. It does not matter which topology we consider or what popularity scenario we evaluate, the condition always stand. With 480 experiments, we can conclude that Multiple Gateways have always better Cache Hit performance than single gateways.

There exist certain cases where the difference of performance between Single and Multiple gateways increase. This is the case for a popularity model with ($\alpha = 1.5$) and for the DTelecom topology also stands when $\alpha = 1.1$. Our intuition is that the connection de-

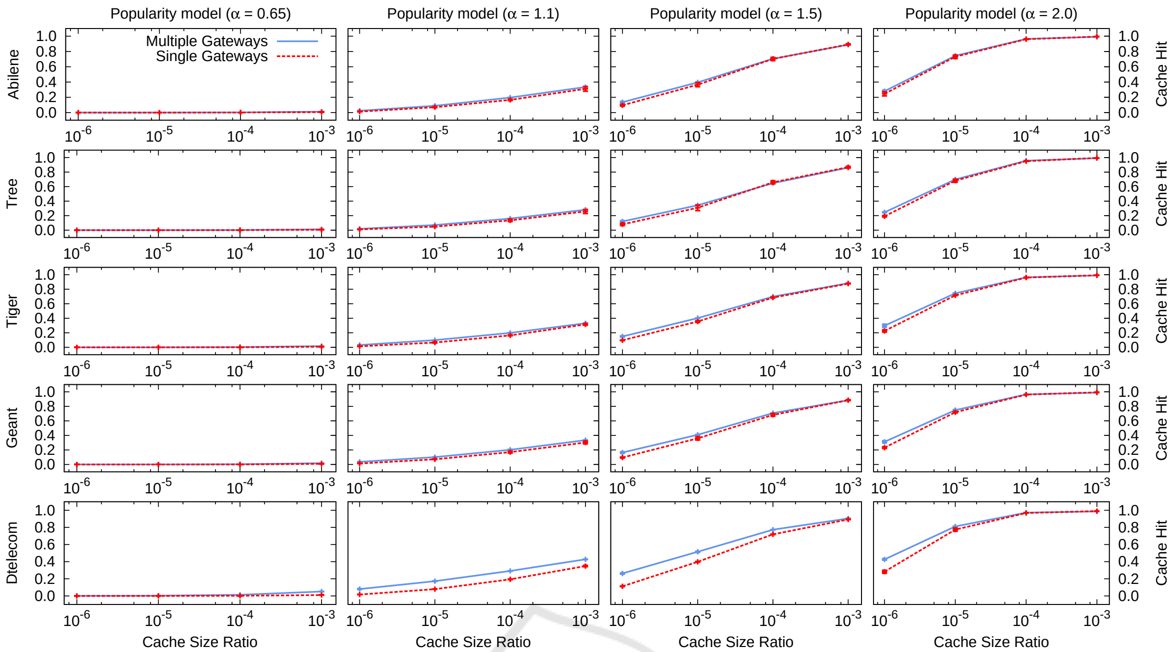

Figure 3: Single vs Multiple Gateways.

gree of the topology has an influence on the caching performance. The connection degree stands for the average number of edges connected to every node in the topology. For instance, Abilene and the Tree topology shares a connection degree of 2.54; Tiger and Geant hold a connection degree between 3 and 4 while Telecom is bigger than 10. If we consider the topologies shown, the gap between Single and Multiple Gateways is not important. Once the connection degree of the topology increases, the gap of performance increases. As we can see analyzing the figure with a top-bottom approach, the gap performances get increased as we switch from less connected topologies towards more connected topologies. The reason for this behavior can be found on the number of paths towards the different gateways: as more alternatives of connections appear, alternative paths can be created to reach different gateways. Thus, gateways become less active and their caches are crossed by less requests.

Hypothesis #2: Multiple Gateways distribute the load better across the network than Single Gateways

To prove that, we collect all the paths generated during the experiences. From these paths, we count all the processed *Interest* messages by the nodes and average them over the total number of processed *Interest* messages. For example, let us imagine that we have two requests that cross through nodes 1,2,3,4 and 4,1,5. As we can see there are 2 requests for content 1 and 4, one request for content 2, 3 and 5. In total, there are 7 requests. For instance, $\frac{2}{7}$ of the requests are for content 1. The aim of this experiment is to show that Multiple Gateways generate a more uniform distribution of the requests than a Single Gateway.

In the Figure 4, we show the results of our experiment. The Figure presents two pie charts: one describes the percentage of processed *Interest* messages with Single Gateways while the second Multiple Gateways. Every pie chart is divided into 11 pieces. Every slice of the pie represents the processed *Interest* messages by every node. After, the slices that surpassed the sum of the average of processed messages plus its standard deviation were separated a bit from the core of the pie, to highlight heavier load.

From the pie charts, we can observe that with Single Gateways, most of the *Interest* messages are managed by only a few nodes. With more precise number, 50% of the *Interest* messages are processed by two nodes of the network. With Multiple Gateways, the processing of *Interest* messages is distributed across all the nodes of the network. It means that even when some nodes are heavy loaded than others, there exists a fair distribution of the charge of the network. Thus, we can conclude that Multiple Gateways are useful to balance the load of processed *Interest* messages.

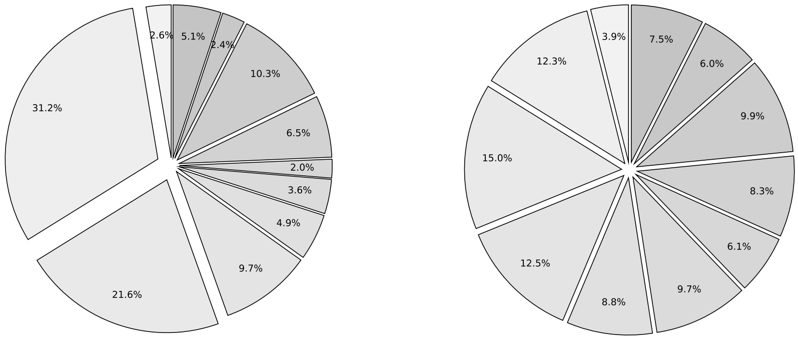

(a) Single Gateways. (b) Multiple Gateways.

Figure 4: Processed Interests per node with Single vs. Multiple Gateways. The pie chart represents the percentage of processed messages by every node. Colors are used to simplify the reading.

6 RELATED WORK

CCN has attracted considerable attention for their caching features. The data structure used for replacing content in the CS is called Replacement Policy (RP). These RPs are largely used in operating systems such as Least Recently Used (LRU), Most Frequently Used (MFU) or First-In First-Out (FIFO). Rosensweig et al. (Rosensweig et al., 2013) have shown that, in the long term, the replacement policies can be grouped into equivalence classes. This means that their results are prone to be similar. Then, the coordination of multiple CS has been studied with the use of caching strategies: Leave Copy Everywhere, ProbCache, Leave Copy Down, MPC (Bernardini et al., b), SACS (Bernardini et al., 2014) or Cache "Less for More". These caching strategies have shown interesting results and are summarized and compared in (Bernardini et al., a).

Rossi et al. (Rossi and Rossini, 2012) study performance of heterogeneous and homogeneous cache sizes. Although heterogeneous caches achieve better results than homogeneous caches, the gain is modest and it incorporates high complexity for managing and maintaining heterogeneous cache sizes.

Another hot topic is where to deploy the caches. In case every Internet router adds caching capabilities, enormous cache sizes (100 Petabytes) will be needed to achieve acceptable rates of performance (Fricker et al., 2012). Due to prohibitive cost of adding caches everywhere and of the required sizes, researchers started discussing about locating caches at edge network (or ISP facilities). Fayazbakhsh et al. (Fayazbakhsh et al., 2013) affirm that placing the caches at edge network is a best choice because of two reasons: the same caching results can be achieved and it is easier to manage only caches at the edges of the network than at global scale. In this sense, Imbrenda et al. (Imbrenda et al.,) analyzed traffic on a real ISP infrastructure and concluded that with a negligible amount of memory of 100MB, the load in the access network can be reduced by 25%. As a consequence, the research on the field points out that CCN must be implemented at the edges of the network and this is the reason to implement CCN into Micro CDNs networks.

7 CONCLUSION

ISPs are facing certain difficulties with the video on demand services. CCN networks appear as an emerging technology to address these issues. Results on the state of the art highlight that caching infrastructures may be helpful to reduce the traffic of the ISP network. Indeed, we have revisited the notion of Micro CDNs deploying the CCN technology. it is essential to study interconnection of these Micro CDNs with the Internet. In this paper, we study the impact on the caching of this interconnection. Caching features will play a major role in the implementation of CCN networks. We have evaluated two alternatives of building gateways: Single Gateways and Multiple Gateways. The multiple disjoint gateways have shown to improve the performance of the network in terms of caching performance, but also in balancing the load of the network. we can obtain a gain of up to 10% in the efficiency of the caches by just doing a small adaptation in the configuration of the network and also distributing better the charge across the network. When a 1 or 10% of gain in terms of cache efficiency may translate into saved money for the exchanged traffic between networks.

The impact on caching of interconnection of CCN networks with the Internet is an important subject that aboards many other subjects. We are interested in revisiting this subjects in the early future. For instance, it is essential to determine the impact of different caching strategies and different routing algorithms on the construction of Gateways. It is important to create models to calculate the potential gains by using Single and Multiple Gateways. Sometimes, the implementation of Single or Multiple Gateways may not be oriented by performance matters but for costs. If this is the case, we expect that models are developed to measure the trade-off of performance against costs.

ACKNOWLEDGMENTS

This work has been partially funded by the Austrian Research Promotion Agency (FFG) and the ENIAC Joint Undertaking, grant 843738 eRamp.

REFERENCES

Bernardini, C., Silverston, T., and Festor, O. A Comparison of Caching Strategies for Content Centric Networking. In *IEEE Globecom 2015*.

Bernardini, C., Silverston, T., and Festor, O. MPC: Popularity-based Caching Strategy for Content Centric Networks. In *IEEE ICC 2013*.

Bernardini, C., Silverston, T., and Festor, O. (2014). Socially-aware caching strategy for content centric networking. In *Networking Conference, 2014 IFIP*, pages 1–9. IEEE.

Cisco (2013). Cisco visual networking index: Global mobile data traffic forecast update, 2013-2018. Technical report, Cisco.

Fayazbakhsh, S. K., Lin, Y., Tootoonchian, A., Ghodsi, A., Koponen, T., Maggs, B., Ng, K., Sekar, V., and Shenker, S. (2013). Less pain, most of the gain: Incrementally deployable icn. In *ACM SIGCOMM 2013*.

Fricker, C., Robert, P., Roberts, J., and Sbihi, N. (2012). Impact of traffic mix on caching performance in a content-centric network. *Computing Research Repository*, abs/1202.0108.

http://blog.netflix.com/2014/04/the-case-against-isp-tolls.html. Netflix. the case against isp tolls. http://blog.netflix.com/2014/04/the-case-against-isp-tolls.html. last seen: 10-16-2015.

https://github.com/mesarpe/socialccnsim. https://github.com/mesarpe/socialccnsim. last seen: 25-03-2014.

http://www.cnet.com/news/france-orders-internet-provider-to-stop-blocking-google-ads/. Cnet. france orders internet provider to stop blocking google ads. http://www.cnet.com/news/france-orders-internet-provider-to-stop-blocking-google-ads/.

Imbrenda, C., Muscariello, L., and Rossi, D. Analyzing cacheable traffic in isp access networks for micro cdn applications via content-centric networking. In *ACM ICN 2014*.

L. Wang, A. M. Hoque, C. Yi, A. Alyyan, and B. Zhang (2012). OSPFN: An OSPF based routing protocol for Named Data Networking. Technical Report NDN-0003, NDN Consortium.

Moy, J. (1998). OSPF Version 2. RFC 2328 (Internet Standard). Updated by RFCs 5709, 6549, 6845, 6860.

Rosensweig, E. J., Menasch, D. S., and Kurose, J. (2013). On the steady-state of cache networks. In *IEEE INFOCOM 2013*, pages 863–871. IEEE.

Rossi, D. and Rossini, G. (2011). Caching performance of content centric networks under multi-path routing (and more). Technical report, Telecom ParisTech.

Rossi, D. and Rossini, G. (2012). On sizing CCN content stores by exploiting topological information. In *IEEE INFOCOM 2012 NOMEN Workshop*, pages 280–285. IEEE.

Shang, W., Thompson, J., Cherkaoui, M., Burke, J., and Zhang, L. (2013). NDN.JS: A javascript client library for Named Data Networking. In *IEEE INFOCOMM 2013 NOMEN Workshop*.