Leveraging Information Centric Networking in Over-The-Top Services

Damien Saucez, Chadi Barakat, Thierry Turletti

Leveraging Information Centric Networking in Over-The-Top Services

Damien Saucez[1], Chadi Barakat[1], Thierry Turletti[1]

INRIA – Sophia Antipolis – France

Abstract. The ubiquity of broadband Internet and the proliferation of connected devices like laptops, tablets or TV result in a high demand of multimedia content such as high definition video on demand (VOD) for which the Internet has been poorly designed with the Internet Protocol (IP). Information-Centric Networking and more precisely Content Centric Networking (CCN) overtake the limitation of IP by considering content as the essential element of the network instead of the topology. CCN and its content caching capabilities is particularly adapted to Over-The-Top (OTT) services like YouTube or Netflix that distribute high-definition multimedia content to millions of consumers, independently of their location. However, bringing content as the most important component of the network implies fundamental changes in the Internet and the transition to a fully CCN Internet might take a long time. Despite this transition period where CCN and IP will co-exist, we show that OTT service providers and consumers have strong incentives for migrating to CCN. We also propose a transition mechanism that leverages caching and enable loosely collaboration.

1 Introduction

The Internet usage has dramatically changed since its creation and a large part of the traffic is now generated by the consumption of multimedia content [1] like video on demand (VOD) or high quality calls. However, the TCP/IP model used today binds content to topological location limiting the mobility of contents, devices, and users between devices. In order to tackle the limitation of IP, Information Centric Networking and Content-Centric Networking (CCN) in particular [6–9] have been proposed to completely remove the notion of location in data networks. In CCN, when a user wants to consume content, it simply has to announce the name of the content it is interested to consume to its neighbor routers. If the content is not available at a neighbor, this last forwards the interest to its neighbors and so on until the content is found. The content then back tracks the trail followed by the interest until the initial requester is reached and intermediate router can cache the content to speed up further content retrieval.

In this paper, we show that *Over-The-Top* (OTT) service providers like Netflix [2], YouTube [3], Hulu [4], or the Xbox LIVE [5] that produce and distribute content to a large basis of consumers can take advantage of using CCN and its loosely collaborative caching mechanism. The idea of CCN and ICN in general is appealing however their deployment require fundamental changes in the Internet [10] and it is then inevitable that CCN will co-exist with IP. Despite the fact that CCN would be only partially deployed in the Internet, we show that it is possible to design a mechanism to interconnect CCN islands over the IP Internet that leverages the use of CCN in OTT even though CCN is not deployed in the Internet core. We design our interconnection mechanism with an overlay optimized for performance and a semi-centralized naming resolution infrastructure that dynamically discover the caches and the content that they store.

To achieve these goals, we first give the necessary background on CCN in Sec. 2. Afterwards, in Sec. 3, we give incentives for OTT consumer networks and OTT service providers to migrate to CCN and determine the conditions that make the adoption profitable for both of them. We then present, in Sec. 4, an overlay-based technique allowing efficient interconnection of OTT consumers and providers over the Internet. The particularity of our overlay is that it is designed to operate in a loosely collaborative environment where consumers might not cache the contents they consume. Our technique relies on a dynamic overlay that associates names and CCN islands and that is optimized for fast transfers. The technique to associate names and CCN islands is described in Sec. 4.3. Finally, Sec. 5 concludes this work.

2 Content-Centric Networking in a Nutshell

Content-Centric Networking (CCN) proposed by Jacobson et al. in [6] rethink the behavior of network while removing any notion of location or topology to only keep the notions of content and content names. In CCN, contents can be anywhere in the network potentially moving or replicated. When a CCN client wants a particular content, it sends an *Interest message* to its neighboring routers. When a CCN router receives such an interest, it forwards it to its neighbors that can potentially provide the content identified by the name contained in the Interest message. Because there is no information about its origin in an Interest message, the router maintains an entry in its *Pending Interest Table* (PIT). The entry maps the name in the interest to the list of *faces* (i.e., interfaces) on which the interest has been received. When the *publisher* of the content receives an interest, it replies with a *data message*. The data message contains the name of the interest and its corresponding *ContentObject*. The data message is sent directly to the face on which the interest has been received. The data message then follows the trail defined by the

PITs until it arrives at the origin of the interest message. As there is no notion of location in CCN, the content can be at any place. Therefore, every router can maintain a cache to copy and store the ContentObjects from the data messages that transit it and any subsequent request for the same content will then receive the content from the copy instead of the content publisher.

3 Incentives for OTT to use CCN

Before trying to integrate CCN with Over-The-Top services, it is important to determine whether there are incentives for such integration. As stated in Sec. 2 or in [6], CCN, compared to IP, provides better security and performance. This last point is very interesting for OTT service providers that deliver multimedia content where performance is a key factor for the adoption of the service by consumers. With CCN, the content can be retrieved from the caches in the different CCN islands, instead of always being delivered by the content publisher. As a result, content retrieval is faster for the consumer and the operational cost of the publisher is reduced. Moreover, as the content is cached by the consumers and because the consumer can provide the content to other consumers, the overall performance increases with the number of consumers instead of decreasing as it is the case in IP today where the content is delivered by the hosting server. This property is particularly interesting because it dampens the effect of flash crowds which are normally very costly for OTT service providers as they have to provision their servers and networks to support them. Using CCN with caching at the consumers has then a direct impact on the profit earned by the OTT service provider as its costs are reduced. However, to benefit from the caching capabilities of consumers, the producer must propose real incentives to its consumers to *collaborate* and cache the content. To understand how incentives can be provided, it is necessary to remember that content in OTT is provided either freely to the consumer or in exchange of a fee. When the content is provided freely, the incomes for the publisher are ensured by advertisements dispersed in the content (e.g., banner, commercial interruptions. . .). A consumer has incentives to collaborate with the system if it receives some sort of discount, expressed in advertisement reduction or fee reduction. On the one hand, the discount has a cost for the publisher as its revenues will be reduced. On the other hand, the collaboration from its consumers reduces its operational costs. Hence, the publisher must determine the optimal discount, such that it maximizes its profit. The situation for the consumer is the exact opposite: its costs are increasing because it is providing content to other consumers but its revenues also increase as it receives a discount on its expenses. We determine in Sec. 3.1 the conditions to be respected by the system to increase the profit of both publishers and consumers when caching is used in CCN. As long as these conditions are respected, deploying CCN at the OTT is beneficial for both producers and consumers.

3.1 Profitability Conditions

To generalize and formalize the discussion started above, let's define $\mathcal{C}$, the set of consumers of a given content, $\mathcal{C}_c \subseteq \mathcal{C}$ the set of collaborative consumers (i.e., those caching content), and $\mathcal{C}_n \subseteq \mathcal{C}$ the set of consumers that do not collaborate (i.e., those that not cache content). We have,

$$\mathcal{C}_n \cup \mathcal{C}_c = \mathcal{C} \ \wedge \ \mathcal{C}_n \cap \mathcal{C}_c = \emptyset. \tag{1}$$

Let $\mathrm{F} : \mathcal{C} \mapsto \mathbb{R}$ be the fee requested to the consumer, $\mathrm{C} : \mathcal{C} \mapsto \mathbb{R}$ be the cost of providing the content to a consumer from the publisher, and $\mathrm{D} : \mathcal{C} \mapsto \mathbb{R}$ be the discount given to the consumers for caching and serving the content on behalf of the publisher. The general formulation of the profit p for a given content at the publisher is given by

Symbol	Description
$\mathcal{C}$	Set of consumers
$\mathcal{C}_c$	Set of consumers that collaborate
$\mathcal{C}_n$	Set of consumers that do not collaborate
$\mathrm{C}(c)$	Cost supported by the producer to provide the content to consumer c
$\mathrm{C}^*(c)$	Cost supported by consumer c to retrieving the content from the publisher
$\hat{c}$	Approximation of a unique cost for the producer
$\mathrm{C}(c, n)$	Cost supported by consumer c to provide the cached content to consumer n
$\mathrm{C}^*(n, c)$	Cost supported by consumer n to retrieve the content from consumer n
$\mathrm{D}(c)$	Discount paid by the publisher to consumer c
$\hat{d}$	Approximation of a unique discount paid by the producer
$\mathrm{E}(c)$	Expenses supported by consumer c to retrieve the content
$\mathrm{F}(c)$	Fee requested by the producer to consumer c
p	Profit of the content producer from providing the content
$\mathrm{R}(n, c)$	Demand function of the consumer n to retrieve a cached copy from consumer c

Table 1. Summary of the variable and function symbols used in Sec. 3.1

$$p = \sum_{x \in C} \mathrm{F}(x) - \mathrm{C}(x). \tag{2}$$

When discount is provided to collaborative consumers, this profit becomes:

$$p = \sum_{c \in \mathcal{C}_c} \mathrm{F}(c) - \mathrm{D}(c) - \mathrm{C}(c) + \sum_{n \in \mathcal{C}_n} \mathrm{F}(n) - \mathrm{C}(n). \tag{3}$$

When the number of collaborative consumers is high and their caches are well dimensioned, we can assume a balanced scenario where the publisher only provides content to collaborative consumers and pays the discount to them, whereas non-collaborative consumers retrieve content directly from caches inside the network. In somehow, we ignore the requests from non-collaborative coming directly to the publisher as they do not enter in the calculation of the optimal discount to propose. Eq. (3) then becomes

$$p = \sum_{c \in \mathcal{C}_c} \mathrm{F}(c) - \mathrm{D}(c) - \mathrm{C}(c) + \sum_{n \in \mathcal{C}_n} \mathrm{F}(n). \tag{4}$$

In this balanced scenario, the publisher gains at providing discount to its collaborative consumers as long as the profit generated with discount (Eq. (4)) is higher than the profit generated without any form of discount (Eq. (2)), as shown below.

$$\sum_{c\in\mathcal{C}_c} \mathrm{F}(c) - \mathrm{D}(c) - \mathrm{C}(c) + \sum_{n\in\mathcal{C}_n} \mathrm{F}(n) \geq \sum_{x\in C} \mathrm{F}(x) - \mathrm{C}(x), \tag{5}$$

$$\Leftrightarrow \sum_{n\in\mathcal{C}_n} \mathrm{C}(n) \geq \sum_{c\in\mathcal{C}_c} \mathrm{D}(c), \tag{6}$$

where Eq. (5) is simplified in Eq. (6) because of the conditions in Eq. (1). In other words, the publisher gains at providing discount to some consumers if the total cost of distributing the content to consumers that do not collaborate would be higher than the total amount of discount given to the consumers that collaborate. Eq. (6) defines the conditions under which collaboration from consumers is profitable for the producer. To determine the optimal discount function $\mathrm{D}(x)$, we first have to determine the conditions necessary for a consumer to collaborate, which decides on the sizes of the two sets $\mathcal{C}_c$ and $\mathcal{C}_n$ In order to determine whether a consumer will collaborate or not, we must determine its costs structure and compare it to the discount proposed by the publisher. As above, we consider a balanced scenario where consumer that do not collaborate always consume cached copies of the content as these copies provide better quality to them for a lower cost [11]. In addition, to obtain an upper bound on the cost per collaborative consumer, we add the further assumption for that. The expenses for a non collaborative consumer n to retrieve the content are the minimum over all cached copied and can be written as:

$$\mathrm{E}(n) = \mathrm{F}(n) + \arg\min_{c \in \mathcal{C}_c} \mathrm{C}^*(n, c),\ \forall n \in \mathcal{C}_n, \tag{7}$$

where $\mathrm{C}^*(n, c) : \mathcal{C} \times \mathcal{C} \mapsto \mathbb{R}$ is the cost for consumer n to retrieve content from consumer c.

Consumers that collaborate have to support the cost of maintaining and distributing copies. Therefore, if $\mathrm{C}(c, n) : \mathcal{C} \times \mathcal{C} \mapsto \mathbb{R}$ gives the cost for the collaborative consumer c to provide the cached content to the non collaborative consumer n, then the cost of the collaboration for c is given by

$$\sum_{n \in \mathcal{C}_n} \mathrm{R}(n, c) \cdot \mathrm{C}(c, n), \tag{8}$$

where $\mathrm{R}(n, c) : \mathcal{C} \times \mathcal{C} \mapsto \{0, 1\}$ is a binary function that determines whether consumer n demands to retrieve content cached at consumer c or not. The total expenses for the collaborative consumer c is thus

$$\mathrm{E}(c) = \mathrm{F}(c) + \mathrm{C}^*(c) + \sum_{n \in \mathcal{C}_n} \mathrm{R}(n, c) \cdot \mathrm{C}(c, n) - \mathrm{D}(c),\ \forall c \in \mathcal{C}_c. \tag{9}$$

To accept to collaborate in the system, the expenses payed by an undecided consumer u must be lower if it collaborates than if it does not. From Eq. (8) and Eq. (7) we thus have:

$$\mathrm{F}(u)+\mathrm{C}^*(u)+\sum_{n\in\mathcal{C}_n}\mathrm{R}(n,u)\cdot\mathrm{C}(u,n)-\mathrm{D}(u)\leq \mathrm{F}(u)+\arg\min_{c\in\mathcal{C}_c}\mathrm{C}^*(u,c). \quad (10)$$

In other words, a consumer should collaborate if the cost overhead caused by the collaboration compared to the cost it should have supported without cooperation (i.e., cost when collaboration minus cost when no collaboration) is compensated by the discount, or:

$$\mathrm{D}(u)\geq\left(\mathrm{C}^*(u)+\sum_{n\in\mathcal{C}_n}\mathrm{R}(n,u)\cdot\mathrm{C}(u,n)\right)-\arg\min_{c\in\mathcal{C}_c}\mathrm{C}^*(u,c). \quad (11)$$

A producer that wants to maximize its profit by promoting collaboration with the help of discount has to solve the following optimization problem defined by the combination of Eq. (6) and Eq. (11):

$$\sum_{n\in\mathcal{C}_n}\mathrm{C}(n)\geq\sum_{c\in\mathcal{C}_c}\mathrm{C}^*(c)+\sum_{n\in\mathcal{C}_n}\mathrm{R}(n,c)\cdot\mathrm{C}(c,n)-\arg\min_{i\in\mathcal{C}_c}\mathrm{C}^*(c,i), \quad (12)$$

$$\mathcal{C}_n\cup\mathcal{C}_c=\mathcal{C}, \quad (13)$$

$$\mathcal{C}_n\cap\mathcal{C}_c=\emptyset. \quad (14)$$

Unfortunately, the producer cannot determine an exact solution for this optimization problem as it knows neither the cost structure of its consumers nor their demands. As a first approximation, the producer can offer the same discount $\hat{d}$ for all the consumers that collaborate and assume that the cost of providing all the consumers, when no collaboration is performed, is equal and of value $\hat{c}$ (e.g., the average cost). In this case, Eq. (6) is approximated by

$$\hat{d} \leq \frac{|\mathcal{C}| - |\mathcal{C}_c|}{|\mathcal{C}_c|}\hat{c}. \tag{15}$$

Eq. (15) shows that the discount is a function of the proportion of consumers that collaborate in the system with the more the number of consumers collaborating, the less the discount can be. This equation is summarized in Fig. 1 for different proportions of collaborative consumers. For a given proportion of collaborative consumers, the profit for the producer is increased as long as the choice of the discount is below the curve. The exact value that maximizes the profit depends on Eq. (11) as the discount is the incentive for consumers to enter in the collaborative system. As this demand is unknown by the operator and depends of the consumer, it can use a conservative approach by first proposing a discount laying in its own area of acceptable discounts. For that amount of discount, the operator will observe a certain amount of consumers entering the collaborative system. If the proportion of collaborative consumers is too high (resp. too low), then the operator reduces (increases) the discount to reduce (resp. increase) the proportion of collaborative consumers. By continuously adapting its discount, the producer can construct the consumers demand function and progressively offer personalized discounts.

Interestingly, as shown by Fig. 1, negative costs are possible and imply negative discounts. A negative value of cost corresponds to the situation where the producer earns values when using its network. This situation typically corresponds to the case where the OTT service provider is at the provider end of of a customer-provider link [12]. In this case, the reduction of traffic caused by the collaboration must be compensated by the increase of the fee requested to the consumers (i.e., a negative discount) which can be justified by the offer of a premium service.

4 Loosely Collaborative Inter-connection of CCN Networks

We have seen in Sec. 3 that Over-The-Top service providers and consumers have incentives in using CCN for participating in the OTT service. We can thus imagine that CCN is deployed at the content producer network and in each consumer network. However, as appealing is content-centric networking, its deployment requires fundamental changes in the Internet and it is very likely that the Internet core will not shift to the

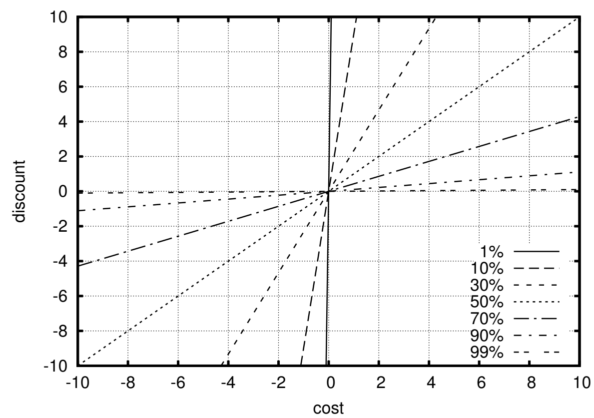

Fig. 1. Evolution of the optimum discount with the cost ($\hat{c}$) for different percentage of the number of consumer collaborating. The producer increases its profit as long as the discount (y-axis) is below the curve for a given cost (x-axis).

new paradigm rapidly, if ever [10]. We will thus see the emergence of CCN networks at the OTT consumer and producer networks, forming *CCN islands* composed of devices supporting CCN but inter-connected together via the traditional IP Internet core. To ensure this co-existence with IP, the CCN islands will form an IP overlay. The construction of such IP overlays has been extensively studied for peer-to-peer systems [13, 14] and mostly focus on reliability and scalability. However, OTT services provide multimedia content making so the performance criterion passing before reliability and the scalability. In addition, we have seen in Sec. 3 that interest of migrating to CCN for an OTT service is in the caching capabilities of CCN. Finally, the adoption of the new model assumes collaboration from the consumers but while strong collaboration is possible inside a network operated by the administrative instances, the OTT consumers are spread on the Internet under distinctive administrations and only a loosely collaboration is possible (e.g., some consumer can cheat or have policies that they cannot disclose). To achieve our goals, we propose an overlay that (*i*) is closely embedded on the IP network in order to minimize path stretch, (*ii*) authorizes traffic deflection to consumer caches as long as performance of consuming contents from caches are better than those of consuming content directly from the publisher, and (*iii*) dynamically learns the cached contents and the consumer CCN islands while operated in loosely collaborative networks.

4.1 IP Overlay for CCN Interconnection

The network in a CCN island can be composed of an arbitrary number of CCN nodes that communicate together by natively using the CCN protocol. Each island is connected to the Internet with a dual-stacked gateway which role is to allow CCN islands to seamlessly communicate over the IP Internet core. The role of a gateway is three-fold: (*i*) associate content to possible CCN islands, (*ii*) encapsulate CCN messages over IP and forward them on the Internet core, and (*iii*) decapsulate packets received from other CCN islands and forward them to the CCN island it is attached to. Fig. 2 summarizes the overlay behavior in normal conditions.

The association between contents and their possible locations is achieved by the *Naming Resolution Infrastructure* (NRI) that dynamically binds content names to the IP address of the gateways of the CCN islands that potentially cache the content. The NRI binds each name with a list of gateway addresses, each address annotated with a priority and a weight. The gateway with the lowest priority value address should be selected because it will offer the best performance. In addition, CCN provides the support of multipath transfers, then if several choices are possible (i.e., several addresses with the same lowest priority value), they can be used in parallel proportionally to their weight (i.e., an address with a high weight is more likely to be selected than an address with a low weight). The presence of several addresses with different priorities improves the robustness of the binding as it is possible that the best gateways determined by the NRI might not be reachable from the NRI's client. A detailed description of the NRI is given in Sec. 4.3.

On the one hand, when a gateway receives an Interest message from its CCN island, it first determines the list of remote gateways that cache the content by sending a resolution querying to the NRI. The gateway selects the best remote gateways according to priorities and weights found in the binding retrieved from the NRI. The interest message is then encapsulated into an IP packet which source is the address of the local gateway and the destination is the address of the remote gateway. To maximize the chance of traversing the middleboxes in the Internet while minimizing the connection time, the CCN message is not encapsulated directly in IP or TCP but in UDP [15]. The encapsulated packet is then sent directly to the Internet core. This encapsulation scheme constitutes a minimalist overlay that minimizes path stretch as packets are sent to the gateways via the direct IP path which avoid detour that can degrade performance. Moreover, our overlay does not require any routing mechanism as it relies directly on its underlay avoiding so routing inconsistencies but also convergence problem. It is worth noticing that the gateway can cache the binding it learns from the NRI for future use.

On the other hand, when a gateway receives an encapsulated CCN message, it decapsulates and forwards it into the CCN network it is attached to. Care must be taken when messages are decapsulated. Indeed, by principle, CCN messages do not contain any explicit information about the node that initially emitted interest for the content. To overtake this constraint, CCN uses symmetric paths between the Interest message and its corresponding Data message. When a Data message is received, it is forwarded to the faces on which an Interest messages for the same content name have been received. Alas, the face that is connected to the Internet at the gateway is an IP interface and thus implies that CCN messages must be encapsulated in IP with the destination address corresponding to the address of the gateway that sent the corresponding Interest message. As this address is not present in the Data message, additional state must be kept by the gateway that received the Interest. To comply with CCN and keep the system its genericity, we do not add this state into CCN but instead use the notion of *virtual IP face*. A virtual IP face is seen as a normal face by CCN. However, virtual IP faces are dynamically instantiated when an Interest message is decapsulated and destroyed whenever the corresponding data message has been encapsulated and sent back to the initial gateway. The necessary state is maintained directly by the virtual IP face that remembers the parameters of the encapsulating IP packet (i.e., IP addresses and UDP ports). Hence, the virtual face can encapsulate and send Data message it receives to the appropriate gateway. On the contrary, when an encapsulated Data message is received, the gateway only has to decapsulate it and forward it to its attached CCN network that maintain PIT entries for the content provided in the message.

The overlay that we propose is a tradeoff between the constraints of IP, the location independence of CCN, and the performance requirements of OTT services. On the one hand, IP strongly depends on the topological location because of the IP addresses, but, on the other hand, CCN is decoupled to location as content can be anywhere in the network. For that very reason and because OTT services require good network performance, the interconnection of CCN islands is done by the mean of a dynamically constructed overlay directly embedded in the IP infrastructure. Hence, the locations remain invisible in the CCN networks but CCN messages are sent optimally between the islands according to the IP path.

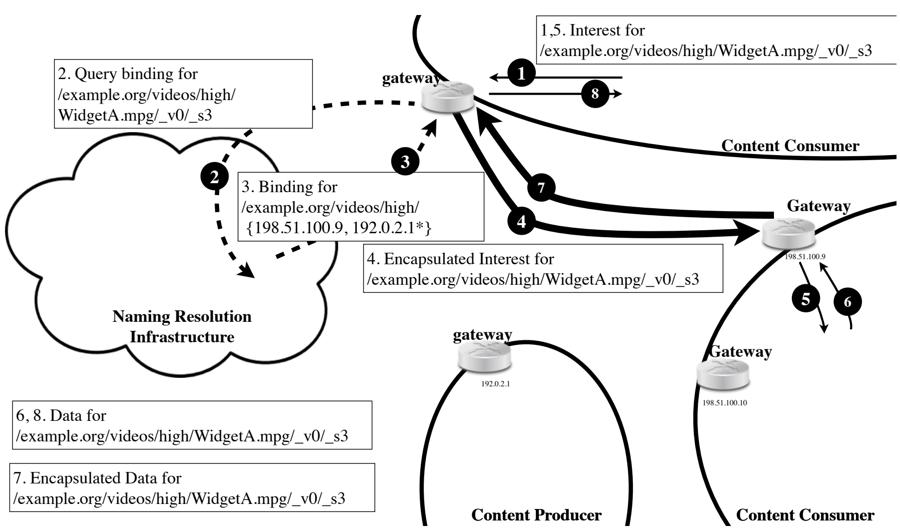

Fig. 2. Example of the transfer of CCN messages with the overlay

4.2 Performance Based traffic Deflection

The role of the Name Resolution Infrastructure is to associate content names to the address of the gateways of CCN islands that potentially cache that contents. As discussed in Sec. 4.3, these binding are constructed such that caches offering the best performance are preferred. However, the conjunction of OTT services and caching forms a loosely collaborative system where the content might not be cached in the selected island even if the binding presumed so. With a traditional IP approach, if the content is not available at a cache, the consumer must try itself another cache as each of its request is targeted to a particular cache with the IP address. On the contrary, with the location independence of CCN, if an Interest message arrives at a CCN island that does not cache the content, it can seamlessly be recursively re-directed to another CCN island that potentially caches the content. Unfortunately, the performance decreases with the number of hops [11] and OTT services depend on network performance. For that reason and to ensure correctness of Interest forwarding in the overlay (i.e., if the content exists, the Interest message will eventually arrive at a CCN island that stores the content), when a CCN island does not store the content requested by another CCN island, it redirects the interest message to the CCN island publisher where it is certain that the content is available. In other words, an Interest message is *deflected* only once to a cache. To allow this, Interest messages are annotated when they enter the island and when a gateway that connects the CCN network to the overlay receives such annotated Interest message, it forwards it to the publisher island. The address of the publisher island gateway is learned with the bindings where publisher gateway addresses are flagged. Fig. 3 summarizes the behavior of the overlay when re-direction is required.

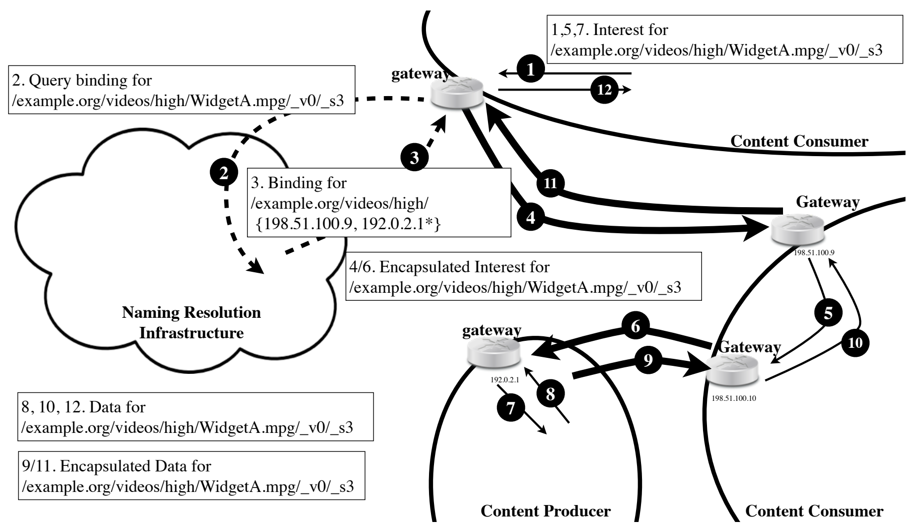

Fig. 3. Example of the transfer of CCN messages with the overlay and re-direction to the content producer

4.3 Name Resolution Infrastructure

The very purpose of the Naming Resolution Infrastructure (NRI) is to provide bindings between content names and CCN islands that cache them. However, as the content can be cached at any position in CCN, the producer of the content will not observe the CCN islands that consume cached content. However, to leverage the use of the cache in our overlay, the bindings must be constructed with the knowledge of all the CCN islands that consume the content. To that reason, we construct a semi-centralized NRI where every query for a particular content is transmitted to the publisher of the content via a decentralized hierarchical resolution infrastructure. In the NRI, the publisher learns all its consumers and their caching infrastructure. The publisher can then construct the binding to leverage the use of caches while following its own policies. The hierarchical infrastructure used to convey the binding queries from consumers to the publisher is built directly on IP and follows the same principle as the Domain Name Service (DNS) and could even be implemented directly with it [16]. We copied model DNS for our NRI because its structure is semi-centralized and years of operation have proven its robustness, scalability, and flexibility. In CCN, the namespace is a prefix tree where the name associated to a node is a prefix of the name associated to any of its children [6]. We can therefore build a hierarchy of *naming resolution servers*, each server being authoritative for the name assigned to it, and the name of a server is a prefix of the name assigned to the children of the servers. A resolution server is then in charge of maintaining the bindings for the names it is authoritative and the list of its children with their assigned name. To resolve a name, an NRI client sends a *resolution query* to a server at the root of the naming tree. The server returns in a *referral message* the address of the children that are authoritative for the longest prefix matching the queried name. The client then sends a query to one of these children that will, at its turn return the address of its longest-matching children, and so on until the server that is authoritative for the perfectly matching queried name is reached. When a server receives a query for a name it is authoritative for, it replies the client with the binding for the name of interest. As the server that provides the binding is the same for any client that is willing to consume the content identified by this name and that this server is operated by the publisher, the publisher is able to construct a map of the CCN islands that consume its content. This consumer map is used to construct the bindings dynamically and then leverage the caching capabilities of CCN. Fig. 4.3 gives an example of name resolution in the NRI.

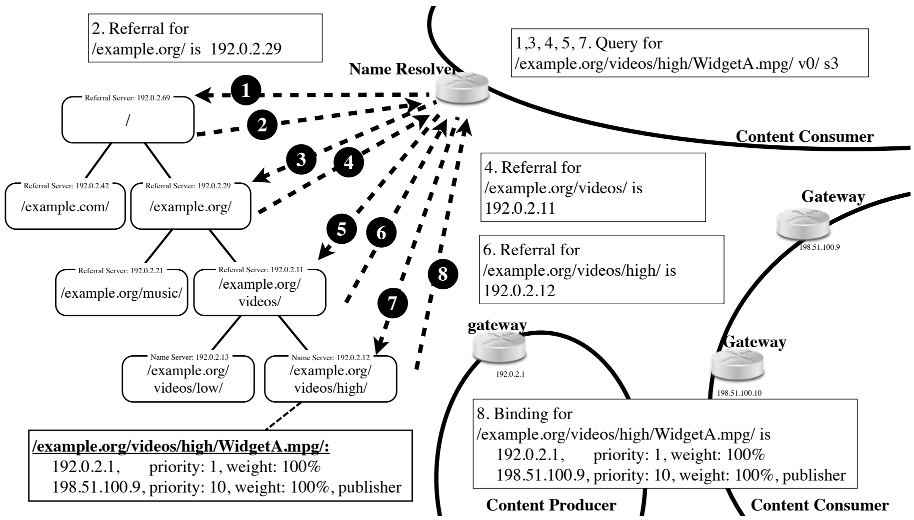

Fig. 4. Naming hierarchy traversal

The binding construction is based on the consumer map the producer builds. To optimize the use of cache and retrieval performance, the optimal binding is built by the producer for each consumer. To do so, the map of consumers is consolidated with BGP information and the *Application-Layer Traffic Optimization* (ALTO) service [17]. With this abstracted map of the network, the publisher determines the topological distance between the islands and group consumers by their topological location as nodes inside a prefix present similar performance [18]. It is worth noticing that our service is loosely collaborative and an island that consumes the content might not cache it. Because of this particularity, the producer should build a statistical caching behavior of its consumers. To do so, the producer can monitor its consumers with the out-of-band mechanism it maintains with its consumers to customize (e.g., with advertisements) the content it provides them even if the static content is cached.

5 Conclusion

In this paper, we show that CCN presents strong incentives for Over-The-Top (OTT) service providers and consumers like YouTube or Netflix. However, the transition to CCN might be long because of the fundamental changes it causes in the network. Therefore, we present a transition mechanism based on a dynamic overlay optimized for performance and a naming resolution infrastructure that allows OTT services to leverage the use of CCN caches in a global Internet where CCN is only partially deployed. The principle of our solution is to place a gateway speaking CCN an IP at the border of each CCN island and to dynamically build bindings between the name of contents cached in the island and the address of the gateway of the island. The binding is then used to dynamically construct an overlay to inter-connect the CCN islands and then leverage the use of caches.

References

1. Labovitz, C., Iekel-Johnson, S., McPherson, D., Oberheide, J., Jahanian, F., Karir, M.: ATLAS Internet Observatory 2009 Annual Report. Technical report, Arbor Networks, the University of Michigan and Merit Network (2009)
2. Netflix (2012) www.netflix.com.
3. Youtube (2012) www.youtube.com.
4. Hulu (2012) www.hulu.com.
5. Xbox live (2012) www.xbox.com/live.
6. Jacobson, V., Smetters, D.K., Thornton, J.D., Plass, M.F., Briggs, N.H., Braynard, R.L.: Networking named content. In: Proceedings of the 5th international conference on Emerging networking experiments and technologies. CoNEXT '09, New York, NY, USA, ACM (2009) 1–12
7. Koponen, T., Chawla, M., Chun, B.G., Ermolinskiy, A., Kim, K.H., Shenker, S., Stoica, I.: A data-oriented (and beyond) network architecture. In: Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications. SIGCOMM '07, New York, NY, USA, ACM (2007) 181–192
8. Ahlgren, B., D'Ambrosio, M., Marchisio, M., Marsh, I., Dannewitz, C., Ohlman, B., Pentikousis, K., Strandberg, O., Rembarz, R., Vercellone, V.: Design considerations for a network of information. In: Proceedings of the 2008 ACM CoNEXT Conference. CoNEXT '08, New York, NY, USA, ACM (2008) 66:1–66:6
9. Zhang, L., Estrin, D., Burke, J., Jacobson, V., Thornton, J., Smetters, D.K., Zhang, B., Tsudik, G., claffy, k., Krioukov, Dmitriand Massey, D., Papadopoulos, C., Abdelzaher, T., Wang, L., Crowley, P., Yeh, E.: Named Data Networking (NDN) Project (October 2010)

10. Rexford, J., Dovrolis, C.: Future Internet Architecture: Clean-Slate Versus Evolutionary Research. Communications of the ACM **53**(9) (September 2010) 36–40
11. Aggarwal, V., Feldmann, A., Scheideler, C.: Can isps and p2p users cooperate for improved performance? SIGCOMM Comput. Commun. Rev. **37** (July 2007) 29–40
12. Gao, L., Rexford, J.: Stable internet routing without global coordination. In: Proceedings of the 2000 ACM SIGMETRICS international conference on Measurement and modeling of computer systems. SIGMETRICS '00, New York, NY, USA, ACM (2000) 307–317
13. Androutsellis-Theotokis, S., Spinellis, D.: A survey of peer-to-peer content distribution technologies. ACM Comput. Surv. **36** (December 2004) 335–371
14. Buford, J.F., Yu, H., Lua, E.K.: P2P Networking and Applications. (December 2008)
15. Honda, M., Nishida, Y., Raiciu, C., Greenhalgh, A., Handley, M., Tokuda, H.: Is it still possible to extend TCP? In Thiran, P., Willinger, W., eds.: IMC 2011, 11th Internet Measurement Conference, New York, NY, USA, ACM SIGCOMM, ACM SIGMETRICS, USENIX, ACM (November 2011)
16. Mockapetris, P.: Domain names - concepts and facilities. RFC 1034 (Standard) (November 1987) Updated by RFCs 1101, 1183, 1348, 1876, 1982, 2065, 2181, 2308, 2535, 4033, 4034, 4035, 4343, 4035, 4592, 5936.
17. Seedorf, J., Burger, E.: Application-Layer Traffic Optimization (ALTO) Problem Statement. RFC 5693 (Informational) (October 2009)
18. Cai, X., Heidemann, J.: Understanding block-level address usage in the visible internet. In: Proceedings of the ACM SIGCOMM 2010 conference on SIGCOMM. SIGCOMM '10, New York, NY, USA, ACM (2010) 99–110