



HAL
open science

Content Delivery Networks as a Virtual Network Function: a Win-win ISP-CDN Collaboration

Nicolas Herbaut, Daniel Négru, Yiping Chen, Pantelis Frangoudis, Adlen
Ksentini

► **To cite this version:**

Nicolas Herbaut, Daniel Négru, Yiping Chen, Pantelis Frangoudis, Adlen Ksentini. Content Delivery Networks as a Virtual Network Function: a Win-win ISP-CDN Collaboration. IEEE GLOBECOM 2016, Dec 2016, Washington, DC, United States. 10.1109/GLOCOM.2016.7841689 . hal-01422267

HAL Id: hal-01422267

<https://hal.science/hal-01422267v1>

Submitted on 25 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Content Delivery Networks as a Virtual Network Function: a Win-Win ISP-CDN Collaboration

Nicolas Herbaut, Daniel Negru
Univ. Bordeaux, LaBRI, UMR 5800
F-33400 Talence, France
{nherbaut, negru}@labri.fr

Yiping Chen
Orange Labs
Chatillon, France
yiping.chen@orange.com

Pantelis A. Frangoudis, Adlen Ksentini
IRISA lab., University of Rennes 1
Rennes, France
pantelis.frangoudis@irisa.fr, adlen.ksentini@irisa.fr

Abstract— The switch in video delivery from traditional medium to Over-The-Top (OTT) based services has been a game changer for stakeholders of the media delivery ecosystem. It has changed the value chain enabling Content Delivery Networks and Content Providers to take the lion’s share at the expense of Internet/Network Service Providers. Even if establishing collaboration between them naturally appears as the key to provide good Quality of Service to the end-users, they struggle in finding efficient and fair ways to do so. This paper directly tackles this problem and proposes a new model for content delivery actors to collaborate over a Virtualized Infrastructure, fairly balancing the revenue stream created. We list the main challenges and the new technical opportunities to solve them, among which the deployment of a distributed Network Function Virtualization (NFV) platform at the edge of the Internet Service Provider’s network where a virtual Content Delivery Network (vCDN) is proposed to be deployed. Furthermore, we apply a game-theoretic analysis to study different ISP-CDN collaboration models and identify optimality conditions for our proposed CDN-as-a-Virtual-Network-Function approach.

I. INTRODUCTION

Real-time entertainment (video and audio streaming) is the number one service on the current Internet: it accounts for 71% of the Downstream Peak Period Traffic in North America in Q1 2016¹. In Europe, adoption is not yet as large but Video-on-Demand Content Provider Netflix reached for about 10% of Peak Downstream Traffic in France, 10 months after launching the service. It has now 70 million paying users worldwide[1]. Given the current average worldwide connection speeds of 5.1 Mbps [2], video content gets provided with 720p or HD resolution. As end-users start watching Over-The-Top (OTT) content, i.e., content by-passing ISPs’ services, not only on their computer screen but also on their TV sets, it is clear that OTT content providers now compete with traditional IPTV providers. However, from an engineering standpoint, there are fundamental technical differences between the two approaches. Indeed, IPTV already offers guaranteed quality for video delivery [3], operators using switched networks [4] enabling a complete control of both network paths and equipments.

As the Internet was not originally designed for streaming high quality video, delivering this massive amount of content is

The work performed for this paper has been partially funded by the FP7 IP T-NOVA European Project (Grant Agreement N#619520) and the FUI French National Project DVD2C

¹<https://www.sandvine.com/trends/>

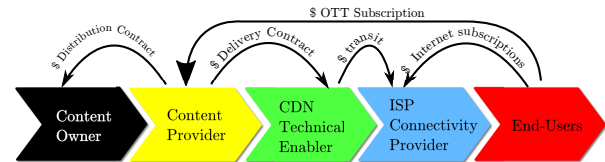


Fig. 1. Value chain for the delivery of video streaming

a challenging task. In [5], the authors identified peering points congestion, inefficient routing protocols, network unreliability and the inefficiency of existing communication protocols as factors adversely affecting such operations. One solution to mitigate those issues is to use Content Delivery Networks (CDNs). CDNs have been used to improve video streaming Quality of Experience to end-users while at the same time limiting the need for Content Providers (CP) to own an infrastructure [6]. By massively deploying servers in strategic locations, CDN Providers assign users to a close-by server, thus reducing hop count and avoiding potential congestion occurrences, while ensuring scalability and reliability.

Figure 1 displays the value chain for video content distribution. On the one hand, the Content Owner sells its content to online Content Providers (CP). On the other hand, ISPs sell plain connectivity to end-users and CPs sell them the access to OTT content. Finally, CDNs are placed between CPs and ISPs as a technology enabler. The large deployment of the CDN-based solution has induced blurred borders on the content delivery market and CPs and ISPs sometimes build their own distribution network. Over the years, companies expanded their activities out of their core business to save costs and propose new services. For example, starting as an owner of search engine technology, Google became a content provider by acquiring YouTube in 2006, then an ISP through Google Fiber in 2011 and more recently its Cloud CDN offering.

In this very competitive market where differentiators are few and the final price is the main selection criterion, the majority of ISPs are being kicked out of the video delivery value chain and are struggling to minimize the relative drop in their revenues [7] (see Figure 2²): between 2010 and 2015, ISP market capitalization increased by only 70%, while CDN

²Public quotation exported from Core US Fundamentals Data via Quandl

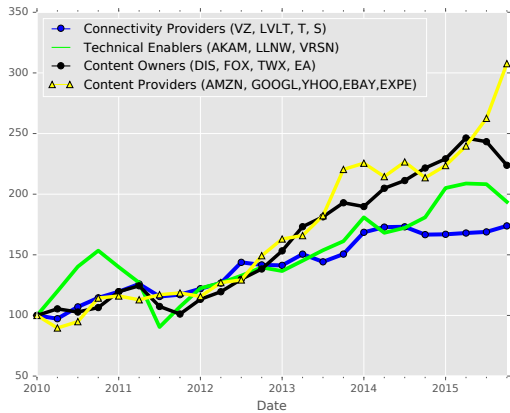


Fig. 2. Mean Market Capitalization of content delivery chain participants.

one jumped up to +100% and Content Owners and Content Providers scored between +125% and +200%.

As mentioned above, the CDN market is coveted by many big players. ISPs hope to find there a new source of revenue. However, even if their role in the delivery part is challenged, ISPs still hold a valuable asset: they manage the whole network, i.e., the complete routing process end-to-end. They can thus leverage this asset in a collaboration process with other actors (CDNs or CPs) in order to lead to strong improvements in end-user experience. CDNs and ISPs do not naturally collaborate. They have their own specific problems to solve resulting in conflicting outcomes. Authors in [8] classify these problems in two categories: **Server Selection** (SS) ones and **Traffic Engineering** (TE) ones. Traditionally, SS problems are under the CDNs’ responsibility. A CDN server can be chosen for a request if (1) it hosts the provisioned content, (2) the format of the content is compatible with the User Agent that performed the request, and (3) the server is able to serve the content with the appropriate quality. (1) and (2) are the core business of the CDNs and do not interfere with the ISP operations. The last problem (3), however, is tightly coupled with the way ISPs handle routing within their networks. User-server assignment choices may drastically suffer from undetected network bottlenecks or end-user mis-location despite their effort in inferring network characteristics [9]. TE problems are handled by ISPs through e.g., the deployment of Internal Gateway Protocols using OSPF (or IS-IS) and modifying weights in order to avoid congestion. However, ISPs do not necessarily optimize their operations to minimize *end-to-end latencies*, as required for multimedia delivery.

Therefore, technologically-speaking, this implies the creation of new platforms addressing specific collaboration challenges between ISPs, CDNs and CPs. The envisioned solution is a collaboration model around a Virtualized Infrastructure. We propose a solution where ISPs can host, provision and manage a NFV infrastructure supporting a large variety of services, while CDNs can bring their specific know-how in terms of content delivery on top of them through a CDN Virtual Network Function (VNF). This is a win-win approach,

since CDNs can expand their coverage dynamically without buying new servers, and ISPs can increase their revenues while reducing traffic between Autonomous Systems (AS). We present the key features of our architecture and the technical challenges and solutions therein in Section II. We further compare quantitatively and qualitatively in Section III our CDN-as-a-VNF approach with alternative CDN-ISP collaboration strategies using a game-theoretic model, deriving its optimality conditions under realistic market parameter values. We conclude the paper along with future challenges in Section IV.

II. A DISTRIBUTED NFV/SDN PLATFORM FOR VIRTUAL CDN DEPLOYMENT IN THE ISP’S NETWORK

In order to overcome the CDN-ISP collaboration challenges, we propose a Virtualized Infrastructure solution capable of matching the ISP’s connectivity “supply” with the CDN’s connectivity “demand”, towards a win-win approach.

We aim at instantiating a Network Function Virtualization (NFV) platform within the ISP Network capable of hosting, among others, Virtual CDN (vCDN) services as depicted in Figure 3. This platform is distributed amongst several NFV Infrastructure Points of Presence (NFVI-POP), at the edge of the network. After presenting the targets for virtualization, we will show that some key enablers for a mutually profitable cooperation can be integrated thanks to currently available technologies, especially Software Defined Networks (SDN).

A. CDN virtualization targets

In the following, we detail the scope of the supported virtualized components, in accordance with ETSI NFV use cases [10]. Even if a complete CDN system could be virtualized in the ISP network, the components that benefit the most from being located close to end-users are caching and streaming servers, as well as others, more management-related, like Monitoring or Request Routers. The CDN internal modules (e.g., ingestors, master caches, legacy routing modules, etc.) can remain in the CDN network.

Streaming VNF: This is the module in charge of storing and delivering the content to end-users. In [11], the authors address the problem of implementing legacy services with Service Chaining + NFV and show that the approach scales well. By leveraging the scalability offered by the NFV platform, the Streaming VNF is designed to accommodate fluctuating traffic by scaling out/in on demand. It also benefits from running on data centers located near the edge routers, offering good network connectivity with end-users located in its neighborhood and limiting the use of peering.

Caching and Routing Orchestrator VNF: As mentioned in the previous section, ISPs will not communicate details about their infrastructure, even if CDNs need some to optimize server selection. One way to circumvent this issue is for the ISP to present only a Virtual Network to the CDNs, which will convey information such as end-to-end delay along the route, bandwidth reservation or congestion state. Figure 3 presents the overlay exposed to the CDNs, the end-user domains connected to edge routers, a logically centralized Caching and

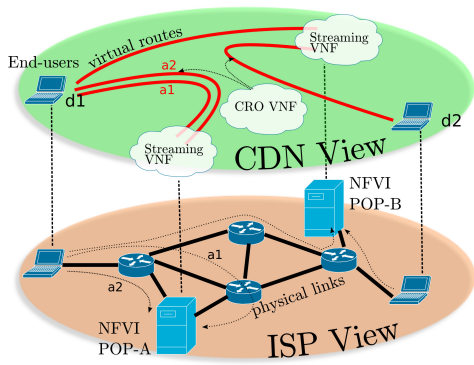


Fig. 3. ISP's Virtualized Network

Routing Orchestrator VNF (CRO VNF) and virtual routes. Each virtual route has specific properties and can be chosen by the CDN to connect a certain end-user domain with an NFVI-POP where a Streamer VNF is deployed. On top of assigning a route, the CDN is responsible for scaling out Streamer VNF instances in each NFVI-POP according to its needs. For example in Figure 3, if a bottleneck occurs in the $a1$ route of the Virtual Network, the CDN can decide through its CRO VNF to offload $a1$ to $a2$.

If the Streaming VNF in NFVI-POP A is the limiting factor, the CDN can scale out the VNF, by requiring more VNF instances to be allocated to it in the NFVI-POP A . If NFVI-POP A and routes $a1$ and $a2$ are saturated, the CDN can decide to use an alternative NFVI-POP B to absorb some traffic from the end-user domain by scaling out Streamer VNF as well as assigning alternative virtual routes to end-users. The ISP exposes an API noted "NET API" to the CDN's CRO VNF to support both route selection service, and server selection as we will describe in Section II-B2.

B. ISP platform key components for collaboration

The proposed solution relies on a vanilla NFV platform. Besides the Streaming VNF and the CRO VNF, we foresee two other necessary modules to enable a fully mutually beneficial ISP-CDN collaboration.

1) *The Marketplace*: This is a business-oriented module allowing to (1) upload the Virtualisation Deployment Unit (VDU) of the vCDN to the ISP NFV platform and (2) to negotiate a service SLA between the ISP and the CDN as presented in Figure 4.

First of all, CDNs upload their VNFs to the Marketplace triggering a notification event on the ISP side. The ISP builds a service composed of those VNFs. This new service is published on the marketplace's service catalog and is made available to the CDN Providers for configuration and launching. During service configuration, the CDN Provider expresses its requirements in an SLA³ [12] which is submitted to the ISP for pricing. After this step, the CDN Provider can accept the contract and launch the service.

³The SLA contains system requirements for the Streamer VNF and CRO VNF (CPU, RAM and disk) and virtual routes (e2e delay, bandwidth).

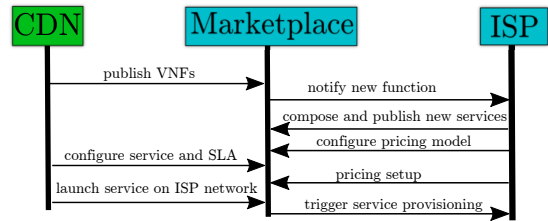


Fig. 4. CDN, ISP and Marketplace interactions

2) *Network Controller API*: This is a technical-oriented module provided by the ISP to the CDN enabling to decide the virtual route the data should take from the Streaming VNF to the end-users. The API enables functionalities to monitor virtual route states, select routes, and possibly decide on which NFVI-POP to scale out Streaming VNFs.

From a technical perspective, virtual routes can be provisioned by the ISP using MPLS. An even more flexible solution can be achieved if operators use SDN to manage the Broadband Access [13], and apply bandwidth calendaring using OpenFlow as described in [14]. Server selection can be achieved by the traditional DNS approach [6] or leveraging more advanced techniques using Virtualized Home Gateways proposed in our previous work [15].

III. COMPARISON OF THE COLLABORATION MODELS

This section outlines the existing collaboration models before formalizing them and comparing them with our proposal.

A. ISP-CDN Collaboration models

Considering that collaboration between CDN providers and ISPs is the key to provide good quality content, industry and research bodies proposed several models covered in [6], [16].

In the **Uncollaborative** scenario, CDN servers are located outside the ISP AS. When selecting a server for a particular streaming session, CDNs cannot access directly routing information (e.g., OSPF weights) and must infer congestion either by measuring TCP packet drops or by using Explicit Congestion Notifications (ECN). This leads to inefficient allocation [17], increased peering traffic for the ISP for each connected CDN, and a relatively high number of hops to reach the CDN surrogate server.

Another alternative, **Managed CDN**, (e.g. Google Global Cache) allows CDNs to deploy surrogate servers directly into the ISP network. This solution enables the ISP to reduce peering traffic and limits the number of hops between the end-user and the content servers. Although this solution brings some cost savings to the ISP and better performance, it still offers limited integration, as the ISP is unlikely to expose its complete network map to CDNs. Scalability is also an issue since surrogate servers come as physical appliances and need to be upgraded in case of traffic increase. An agreement must also be signed between the ISP and every CDN Provider.

A third approach consists of **Licensed CDNs**, (e.g. Akamai Aura Licensed CDN) where ISPs can resell CDN services to their business customers while using CDN-licensed software.

Finally, the more traditional **Telco CDN** (e.g. Level 3) solution is where the ISP assumes the whole responsibility and expertise costs to design and implement the system and takes the role of real CDN providers.

B. Study of CDN and ISP profitability

In the following section, we detail the costs and earnings related to each collaboration model in order to compare them to our proposal. Table II outlines the results.

1) *No Collaboration*: In this baseline scenario, ISPs have to assume a cost b of internal bandwidth to accommodate the multimedia streams from end-users to the server selected by the CDNs to serve the content. ISPs and CDNs must pay a transit cost t to exchange data between their respective AS. This cost could be paid to transit providers (typically a Tier-1 ISP) or can come from paid peering or possibly from settlement-free peering (in which case some collocation costs can occur). We consider that the CDNs generate P net profit from selling their service to content providers (discounting every cost except t). ISPs however make no profit in this scenario, as all the added value goes to the OTT provider.

2) *Managed CDN*: In this model, the CDN redirects a fraction $\alpha \in [0; 1]$ of its traffic to an appliance within the ISP's network. Both CDN and ISP win, as they need to cover only a fraction of the external traffic's cost $(1 - \alpha)$. The flows stay within the ISP AS and do not need to enter the CDN AS.

3) *Licensed CDN*: In this scenario, the ISP resells CDN services to its business customers. To run the service, ISPs need to maintain an infrastructure. Here, ISPs do not own the software but pay a license fee to the CDNs. The pricing l corresponds to the price that the ISP would have paid if 100% of the traffic had been served by the software under license and likewise, s is the price that the ISP would have covered if 100% of the infrastructure had been used. For the sake of simplicity, we assume that those costs are linear with respect to the share of traffic served by the system, noted β . We also assume that the ISPs pay only for the share of the infrastructure used. This assumption holds if the ISP's "CDN business unit" performs an internal re-billing for infrastructure costs or if the ISP buys computing power from a cloud provider. So in this case, the ISPs only covers βl licensing fees and only βs infrastructure cost. The important advantage of such collaboration is the reduction of cross-AS traffic by a ratio $(1 - \beta)$ and the extra earnings r for the ISP, which is also considered linear wrt to β .

4) *TelcoCDN*: This case is similar to the previous one, except that instead of licensing costs, ISPs own the software and assume the costs of engineering and maintenance d . No payment is made to the CDNs.

5) *CDN as a vNMF*: This scenario combines the advantages of Managed CDN and Licensed CDN with respect to external traffic reduction, as both α and β factors are discounted. ISPs still have to support infrastructure costs s but also still benefit from service resell. Their role is somewhat similar to a Cloud Provider, except they have more leeway to optimize the network end-to-end, from the servers to the end-users.

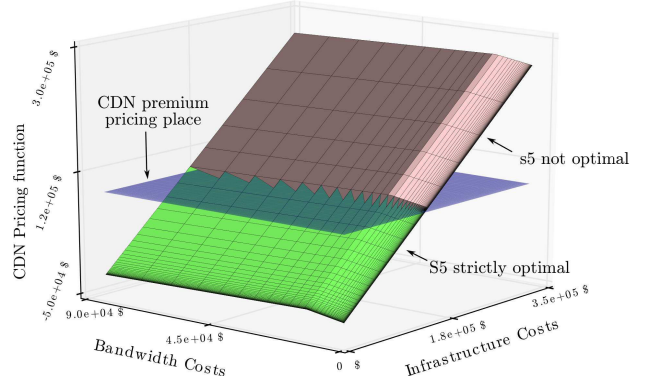


Fig. 5. Optimality zone for the CDN-as-a-VNF strategy (5)

Their pricing function noted f is used to determine the price the CDNs will pay for the hosting. The CDNs have their own pricing function g which will be used to bill their own customers for this premium service.

C. Optimality conditions for the CDN-as-a-VNF strategy

This section deals with the formalization of the optimality problem, considering Game Theory principles.

Definition 1: CDN-ISP collaboration is a Game noted G_{collab} for which each CDN's (resp. ISP's) strategy $s_i^C \in S^C$ with $i \in [0, |S^C|] = I^C$ (resp. $s_j^I \in S^I$ with $j \in [0, |S^I|] = I^I$) has payoffs p_i^C (resp. p_j^I) obtained by subtracting the corresponding costs and earnings from Table II.

Definition 2: The CDN-ISP collaboration model χ is strictly optimal if the corresponding combined strategies $(s_\chi^C, s_\chi^I) \in S^C \times S^I$ form a pure strategy equilibrium for G_{collab} .

This means that the problem can be addressed by solving a set on inequalities arising from the fact that s_χ^C and s_χ^I are strictly dominant strategies.

$$(\forall s \in S^C / \{s_\chi^C\}, s \ll s_\chi^C) \Leftrightarrow (\forall i \in I^C, p_i^C \geq p_\chi^C \Leftrightarrow i = \chi)$$

$$(\forall s \in S^I / \{s_\chi^I\}, s \ll s_\chi^I) \Leftrightarrow (\forall j \in I^I, p_j^I \geq p_\chi^I \Leftrightarrow j = \chi)$$

Definition 2 is restrictive and may not apply to all the cases. However, it provides a simple tool to study the most important factors that condition the success of the CDN-as-a-VNF strategy. We plan to study more permissive forms of optimality, where different strategies can be chosen for different use cases, in future works. Let us now determine the parameter constraints enabling strict optimality of the CDN-as-a-VNF strategy. First, we analyze the existing strategies dominations for the ISP and for the CDN.

1) *For the ISP*: The uncollaborative strategy is dominated by the Managed CDN strategy, as $\alpha > 0$ so we have $s_1^I \ll s_2^I$. We are left with the system: $s_2^I \ll s_5^I$, $s_3^I \ll s_5^I$ and $s_4^I \ll s_5^I$:

$$\begin{aligned} f(\alpha, \beta) &> \alpha s + \beta(s - r) \\ f(\alpha, \beta) &> \alpha s + \beta(s + l) \\ f(\alpha, \beta) &> \alpha s + \beta(s + d) \end{aligned} \tag{1}$$

TABLE I
NOTATIONS AND ESTIMATES

| Notation | Units | Description | Estimates for UC1 | Source used for estimates |
|--------------|-------|---|-------------------|--|
| b | \$ | internal bandwidth costs | - | - |
| t | \$ | transit costs | \$21,000\$ | bandwidth cost: http://drpeering.net , internet daily traffic at POP ⁴ |
| P | \$ | profits margin of CDN | - | - |
| α | % | share of the traffic toward CDN device placed inside ISP's AS | 80% | - |
| β | % | share of the traffic toward ISP Business Customers | 10% | - |
| s | \$ | cost of ISP nVF infrastructure | 80,000\$ | hardware requirements for swiftstack storage, based on commodity prices with a 3 year asset Depreciation |
| l | \$ | licenses paid by ISP to CDN | 24,000\$ | Based on Swiftstack pricing ⁵ |
| r | \$ | CDN business margin | 20,000\$ | with median price 0.025\$ per GB |
| d | \$ | cost for Designing and maintaining CDN software | 555000\$ | acquisition of a CDN Startup of a 20M\$ valuation |
| $f(\cdot)$ | \$ | pricing function used by the ISP to bill the CDN | - | - |
| $g(\cdot)$ | \$ | pricing function used by the CDN to bill its Customer | - | - |
| $g^*(\cdot)$ | \$ | market pricing used by CDNs for premium service | 100,000\$ | cdn77 pricing website |

TABLE II
EARNINGS AND COSTS FOR CDN AND ISP

| | ISP Earnings | ISP Costs |
|-----------------------------|------------------------------|---|
| 1 - No Collaboration | 0 | $b + t$ |
| 2 - Managed CDN | 0 | $(1 - \alpha)(b + t)$ |
| 3 - Licensed CDN | βr | $(1 - \beta)(b + t) + \beta(s + l)$ |
| 4 - Telco CDN | βr | $(1 - \beta)(b + t) + \beta(s + d)$ |
| 5 - CDN-as-a-VNF | $f(\alpha, \beta) + \beta r$ | $(1 - \alpha - \beta)(b + t) + (\alpha + \beta)s$ |
| | CDN Earnings | CDN Costs |
| 1 - No Collaboration | P | t |
| 2 - Managed CDN | P | $(1 - \alpha)t$ |
| 3 - Licensed CDN | $P + \beta l$ | $(1 - \beta)t$ |
| 4 - Telco CDN | P | $(1 - \beta)t$ |
| 5 - CDN-as-a-VNF | $P + g(\alpha, \beta)$ | $(1 - \alpha - \beta)t + f(\alpha, \beta)$ |

2) *For the CDN:* Like in the ISP case, the uncollaborative strategy for the CDN is dominated by every other strategy as $\alpha > 0$ and $\beta > 0$ since there will always be some traffic toward the CDN server located within the ISP AS. So we have $s_1^C \ll s_i^C, \forall i \in \{2, 3, 4\}$. We can also see that the Licensed CDN strategy will always dominate the Telco CDN strategy, as $\beta l > 0 : s_4^C \ll s_3^C$. Finally, the only two inequalities to consider derive from $s_2^C \ll s_5^C$ and $s_3^C \ll s_5^C$

$$\begin{aligned} g(\alpha, \beta) &> f(\alpha, \beta) + (\alpha - \beta)t \\ g(\alpha, \beta) &> f(\alpha, \beta) + (l - t)\beta \end{aligned} \quad (2)$$

On top of the optimality conditions exposed in equation (2), we add another inequality based on current market price for premium⁶ CDN delivery.

$$g(\alpha, \beta) \leq g^*(\alpha, \beta) \quad (3)$$

If ISPs pricing is above the existing market price for premium CDN, then the proposed model is not feasible as it would be outperformed by the currently deployed one. Let us now evaluate the principles of the costs estimation in the proposed CDN-as-a-VNF model.

⁶For premium CDN, more POP are used for content replication, increasing the average quality of service

3) *Costs estimation:* we considered a use case *UC1* where:

- A CDN delivers on average 45GB per user per month. Those numbers correspond to the average data consumption of a Netflix User⁷.
- A small ISP in France handles 20% of the overall traffic of this CDN (1,000,000 paying Netflix users in France).

To calculate t , the bandwidth needed to support this load corresponds to the average hourly peak bandwidth roughly amounting to 5.63%. With a bandwidth priced at 0.63\$ per Mbps/month, it sums up to $4.10^{10} \times 200000 / (30 \times 0.0563 / (60 \times 60) / (0.63 * 8.10^6)) \times \approx 21,000\$$. The cost of the ISP infrastructure has been tailored to support the entire Netflix catalog (3PB) over 6 POPs. We followed SwiftStack (an open source object storage software that can be used in a CDN use case) hardware specifications⁸ and added housing costs⁹ and maintenance cost handled by a third party¹⁰. Concerning the hardware investment, in order to determine the monthly cost, we made the assumption that it would be amortized linearly in 36 months. To estimate licensing fees l of the CDN, we used public software license prices from SwiftStack. We estimated the price of building a CDN software d by assuming that the ISP would acquire a small CDN company¹¹ for 20M\$ and amortize it over 36 months. All these estimations are summarized in Table I, along with the pointers to data.

Considering those estimations, the evaluation conducted towards the performance of our CDN-as-a-VNF model is depicted in Figure 5. We make ISP infrastructure cost s (resp. ISP bandwidth costs t) vary from $4s$ to $s/1000$ (resp. $4t$ to $t/1000$). The flat surface represents the market price g^* that the CDN currently charges for "Premium" CDN delivery. The oblique surface represents the break-even values for the CDN pricing function g . This surface has been colored green to emphasize the values where the g function is below the premium CDN fare. It denotes the area where our proposal

⁷TDG: Netflix Streaming Volume - <http://tdgresearch.com/>

⁸SwiftStack : Hardware Reference Architectures - <https://swiftstack.com>

⁹Housing Pricing - <https://www.infomaniak.ch/en/housing>

¹⁰Canonical Bootstack - <http://www.ubuntu.com/cloud/openstack>

¹¹E.g. Telstra buys video platform company Ooyala - <http://reuters.com>

| | Favor Collaboration | Avoid Lock-in | Scalable |
|-----------------|---------------------|---------------|----------|
| Uncollaborative | - - | + + | + |
| Telco CDN | + + | + + | + |
| Managed CDN | + | + | - - |
| Licensed CDN | + + | - - | - |
| CDN-as-a-VNF | + + | + + | + + |

TABLE III
COMPARISON OF DIFFERENT CDN DEPLOYMENT MODELS

outperforms the current premium CDN pricing model. The zone where the surface is colored red corresponds to a zone where the optimality for CDN as a VFN is not feasible. This occurs for high infrastructure costs, where the ISP cannot cover the infrastructure cost without making the CDN billing a price above current market price g^* . Let us now evaluate quantitatively the possible gain of the CDN-as-a-VNF approach in the case of UC1, following the assumptions made noted θ , and discuss on the potential cases of revenue sharing.

4) *Balancing the revenue*: Considering UC1, the estimated gain of the proposed CDN-as-a-VNF approach is $g^*(\theta) - g(\theta) \approx 40,000\$$, corresponding to a 60% increased profit. When this approach is used, three possible cases can occur:

- ISP stays at break-even and the CDN can increase its pricing g up to the premium CDN market price g^*
- The CDN stays at break-even and the ISP can increase its pricing f until $g(\cdot) = g^*$
- Benefits for the approach ($g^* - g(\cdot)$) can be shared between ISP and CDN.

Another interesting feature with this approach is to let market forces decide how to split the benefits of collaboration. By negotiating the balance of revenues in the feasible value space, both ISP and CDN are able to take into account their respective bargaining power to reach a fair balance. For instance, if the CDN does not have a good peering for a certain zone, it can reduce its profits to make sure that the ISP accepts its collaboration request. On the other hand, if the ISP has a lot of spare resources, it can reduce its margin to encourage the CDN to use its infrastructure.

Future works will detail more advanced game setups for ISP-CDN collaboration, not only to fairly balance revenues generated from the collaboration, but also taking into account the *Social Welfare* [18], which in our case is the end-user QoE.

D. Qualitative comparison of Collaboration models

On top of the quantitative advantages shown earlier, the CDN-as-a-VNF approach has other qualitative aspects. Table III sums them up over the different models according to the following set of features:

1) *ISP-CDN Collaboration*: The lack of collaboration between ISPs and CDN providers is as technical as it is business related. The only truly non cooperative case is the Pure Player CDN model where all traffic is exchanged between ISPs and CDNs ASs. By providing a technical way to collaborate fostered by business incentives, our solution allows matching the demand and the supply on the data delivery market, for the benefit of the end-user.

2) *Lock-in*: The possibility to be able to accommodate and make compete several actors over a given service is a good

argument in favor of CDN-as-a-VNF. Licensed CDN, on the contrary, binds durably the ISP to the CDN.

3) *Scalability*: Scalability is directly achieved when the actor is in full control of its platform like in the uncollaborative case or the Telco CDN one. Traditionally, both Managed and Licensed CDN solutions relying on physical appliances are not scalable enough to respond to the trend of increasing traffic. The VNF approach is the best option since automatic scaling out is directly supported on the platform.

IV. CONCLUSION AND FUTURE WORK

We proposed a design for an ISP NFV platform in which CDNs can run their delivery functions. We crafted technical and business solutions addressing shortcomings of current collaboration models. We translated the ISP-CDN collaboration problem as a Game and investigated the optimality conditions using sensible estimates. We are in the process of implementing a proof of concept of our approach in collaboration with French ISP Orange, into which we will measure the performance in terms of peering traffic reduction, increased quality of service for end-users and scalability, all according to real Internet traces coming from Orange's network. Furthermore, we plan to release the strict optimality condition for the collaboration and assess more complex cases where several strategies can be selected from several providers by competing clients.

REFERENCES

- [1] Netflix : Financial results. [Online]. Available: <http://ir.netflix.com/>
- [2] "State of the internet," Akamai, Q2 2015 Report. [Online]. Available: content.akamai.com/PG6575-q1-2016-soti-connectivity-report.html
- [3] ITU-T, "Y.1901 requirements for the support of iptv services, 2009."
- [4] ATIS IPTV Exploratory Group, "Report and recommendation to the TOPS council," 2005.
- [5] E. Nygren *et al.*, "The akamai network: a platform for high-performance internet applications," *ACM SIGOPS Operating Systems Review*, vol. 44, no. 3, 2010.
- [6] M. Pathan, "Cloud-based content delivery and streaming," in *Advanced Content Delivery, Streaming, and Cloud Services*. Wiley, 2014.
- [7] J. Wulf *et al.*, "Carrier activities in the CDN market - an exploratory analysis and strategic implications," in *Proc. 14th International Conference on Intelligence in Next Generation Networks (ICIN)*, 2014.
- [8] W. Jiang *et al.*, "Cooperative content distribution and traffic engineering in an ISP network," in *Proc. ACM SIGMETRICS*, 2009.
- [9] R. Krishnan *et al.*, "Moving beyond end-to-end path information to optimize CDN performance," in *Proc. ACM IMC*, 2009.
- [10] "Network functions virtualisation (NFV); use cases," ETSI, GS NFV 001-V1.1.1, 2013.
- [11] N. Bouten *et al.*, "Towards NFV-based multimedia delivery," in *Proc. IFIP/IEEE IM*, 2015.
- [12] J. Famaey *et al.*, "An SLA-driven framework for dynamic multimedia content delivery federations," in *Proc. IEEE NOMS*, 2012.
- [13] J. Rükert *et al.*, "Flexible traffic management in broadband access networks using software defined networking," in *Proc. IEEE NOMS*, 2014.
- [14] T. D. Nadeau *et al.*, *SDN: software defined networks*. O'Reilly, 2013.
- [15] N. Herbaut *et al.*, "Migrating to a NFV-based home gateway: introducing a surrogate vNF approach," in *Proc. 6th International Conference On Network of the Future (NoF)*, 2015.
- [16] V. S. Pai, "Next-generation CDN: A CoBlitz perspective," in *Advanced Content Delivery, Streaming, and Cloud Services*. Wiley, 2014.
- [17] B. Frank *et al.*, "Pushing cdn-isp collaboration to the limit," *ACM SIGCOMM Comput. Commun. Rev.*, 2013.
- [18] M. RABIN, "Incorporating fairness into game theory and economics," *The American Economic Review*, vol. 83, no. 5, pp. 1281–1302, 1993.