



HAL
open science

Vertical Integration of CDN and Network Operator: Model and Analysis

Patrick Maillé, Gwendal Simon, Bruno Tuffin

► **To cite this version:**

Patrick Maillé, Gwendal Simon, Bruno Tuffin. Vertical Integration of CDN and Network Operator: Model and Analysis. IEEE/ACM International Symposium on Modelling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS 2016), Sep 2016, London, United Kingdom. hal-01398917

HAL Id: hal-01398917

<https://inria.hal.science/hal-01398917>

Submitted on 18 Nov 2016

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.

Vertical Integration of CDN and Network Operator: Model and Analysis

Patrick Maillé
Institut Mines-Télécom
2, rue de la Châtaigneraie
35576 Cesson Sévigné, France
patrick.maillé@telecom-bretagne.eu

Gwendal Simon
Institut Mines-Télécom
2, rue de la Châtaigneraie
35576 Cesson Sévigné, France
gwendal.simon@telecom-bretagne.eu

Bruno Tuffin
Inria
Campus Universitaire de Beaulieu
35042 Rennes, France
bruno.tuffin@inria.fr

Abstract—Content Delivery Network (CDN) services largely affect the delivery quality perceived by users. While those services were initially offered by independent entities, some large Internet Service Providers (ISPs) now develop their own CDN activities to control costs and delivery quality. But this new activity is also a new source of revenues for those *vertically integrated* ISP-CDNs, which can sell those services to content providers.

In this paper, we investigate the impact of having an ISP and a vertically-integrated CDN, on the main actors of the ecosystem (users, competing ISPs). Our approach is based on an economic model of revenues and costs, and a multilevel game-theoretic formulation of the interactions among actors. Our model incorporates the possibility for the vertically-integrated ISP to *partially* offer CDN services to competitors in order to optimize the trade-off between CDN revenue (if fully offered) and competitive advantage on subscriptions at the ISP level (if not offered to competitors). Our results highlight two counterintuitive phenomena: an ISP may prefer an independent CDN over controlling (integrating) a CDN; and from the user point of view vertical integration is preferable to an independent CDN or a no-CDN configuration. Hence, a regulator may want to elicit such CDN-ISP vertical integrations rather than prevent them.

I. INTRODUCTION

Popular multimedia service providers (e.g., YouTube, Netflix, and Twitch) are called Over-The-Top (OTT) because they deliver their content on top of the public Internet infrastructure. Measurements showed that Twitch regularly delivers more than two tera-bits of data per second [1]. To manage data delivery at such large scale, the Content Delivery Networks (CDNs) have emerged as key players in the delivery chain. A CDN exploits multiple *edge servers* in the networks of Internet Service Providers (ISPs) to deliver data on behalf of a Content Provider (CP). The type of contracts between the CDN and its client CPs, and the resulting storage space management, can significantly impact user perceived quality, CDN service prices, and CDN revenues [2]. Also, the CDN position of intermediary player between the ISP and the CP raises a number of questions regarding the impact on the competition between ISPs [3] and between CPs [4], as well as possible breaches of the net neutrality principle [5].

The interplay between all the players in the delivery chain has become trickier to analyze because of the increasingly complex Internet ecosystem [6]. This is particularly true since traditional players have developed their own CDN activities,

as is typically the case of CPs (for instance Google¹ and Amazon²), of ISPs (for instance Telefonica [7]), and of network carriers (for instance Level3³) [8]. In economics, such a configuration where one player controls a company that is also part of the delivery chain is called *vertical integration* [9, 10].

While the pricing of the CDN activity has been the subject of some research work [11, 12, 13], the complex relationships between actors and the consequences of the CDN on the ecosystem are ignored. Faratin [14] explains the raise of CDNs through an analysis of the wholesale market failure due to the lack of end-to-end coordination between ISPs and content providers; he also analyzes the best strategies of the intermediary actors using game theory as a modeling tool, but does not consider vertical integration configurations.

We focus here on one of the possible vertical integration scenarios where the CDN is owned by an ISP, and we study the impact of this vertical integration on other actors. Such an integration has been considered in the past [15, 16]. Lee et al. [15] focus on the possible cooperation among ISP-CDN entities to compete with existing CDN actors, while Hau and Berenner [16] analyze how interconnection prices between CDNs and ISPs should be set; here we rather study the impact of CDN integration on the inter-ISP competition.

The (integrated) CDN has to choose its level of investment in the infrastructure (in particular storage capacity in the edge servers) to improve the delivery quality. The integrated CDN can also differentiate the quality of the delivery depending on the ISP, favoring the integrated ISP over its competitor. Indeed, when deciding the delivery quality for requests issued by the competitor's subscribers, the CDN has to find a balance between degrading the quality, which would make the integrated ISP *comparatively* more attractive, and increasing the quality, which would make the CP *globally* more popular. In the former case, the ISP-related revenues of the integrated entity increase while the latter case results in higher CDN-related revenues.

We introduce a model, which integrates the main interplaying strategies of the players in the scenario of a delivery chain with a vertically integrated CDN. With our model, we

¹<https://cloud.google.com/compute/docs/load-balancing/http/cdn>

²<https://aws.amazon.com/cloudfront/>

³<http://www.level3.com/en/products/content-delivery-network/>

identify the most profitable strategy for each of the players (especially the ISP that owns the CDN and the competitor ISP). Furthermore, since one of our goals is to better understand whether regulation should be enforced or not, we use our model to compare the integrated configuration against traditional scenarios where the CDN is independent, and where the CP does not use any CDN.

Our model highlights diverse possible phenomena, depending on the characteristics of the user population regarding their ISP subscription choices. When users are very price-sensitive, competition among ISPs is likely to lead to low revenues from subscriptions, hence the integrated entity would favor revenues from CPs and decide to offer high quality to all users. Conversely, if user demand does not vary much with prices, prices stay high and the revenue from ISP subscriptions exceed those from CP delivery fees; the integrated CDN then chooses to differentiate traffic and thereby biases the ISP competition, a configuration suggesting some attention from regulatory bodies. Finally, from the user point of view, an integrated CDN seems to be the best option in both cases, since it leads to the largest number of users being served, even compared with an independent CDN configuration.

The remainder of the paper is organized as follows. Section II introduces the setting and the economic model we consider, including costs and revenues for the CDN and ISPs. Section III describes the strategic decisions that are to be made, and the corresponding noncooperative game in three configurations: integrated CDN, independent CDN, and no CDN. A numerical analysis for a specific cost function is carried out in Section IV, while Section V concludes and provides directions for future work.

II. MODEL: ACTORS, GAINS AND COSTS

We consider the model displayed in Figure 1, the actors being ISPs, CPs and the (unique in this work) CDN. This model does not aim at exhaustively representing all the monetary flows between the players of the delivery chain. Instead, it captures the most prominent interactions between players, especially (i) CPs paying the CDN to deliver video data to the end-users, and (ii) end-users paying their ISPs to have an access to the Internet. The key point of this model is that ISP 1 and the CDN are integrated into the same entity, which means that the revenues come from both sides. Please note that we do not deal here with the monetary agreements between the CDN and both ISPs for the delivery of data. For the sake of simplicity, we depict only two ISPs but our study applies to any number of ISPs where ISP 1 is integrated with the CDN and all other ISPs are independent network operators.

In the remainder of the section, we propose a model for the actors benefits and costs, which—as revenue-maximizing entities—determine their behavior in the interactions.

CDN revenues. The CDN earns a direct revenue from the content providers proportionally to the volume of traffic. The price per unit of volume depends on the average quality of service. For simplicity, we assume that each user induces the

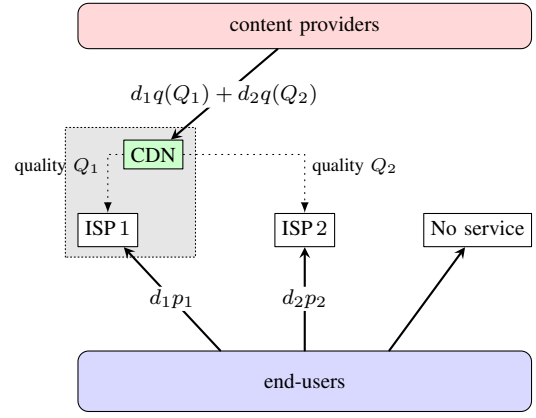


Figure 1. Representation of relations between the players of the delivery chain. The plain arrows depict monetary transactions. The dotted arrows present data delivery from the CDN to the ISPs and the gray area is the integrated CDN+ISP.

same amount of traffic so that the revenue is expressed per unit of end users. This relation between revenues and mass of users corresponds to the traditional volume-based revenues. Let d_i be the mass of end-users who are subscribers of ISP i , for $i \in \{1, 2\}$. The price asked by the CDN increases with the delivered quality level. Since we aim at studying the incentives for the CDN to differentiate services between ISPs, we may end up with different unit prices for different ISPs. We denote by $q(Q_i)$ the price per unit of users generated from ISP i 's end users, assumed to be a function of the average quality Q_i among its subscribers. This price is paid by content providers, which requested the traffic delivery service. Overall, the revenue of the CDN can thus be written as $d_1q(Q_1) + d_2q(Q_2)$. In the following, we introduce the model for the mass of users.

Quality of delivery per ISP. Content providers have to deliver their data to end-users. The default quality of experience for a subscriber of ISP i when the CPs do not use any CDN service is $Q_i = 1$. When the delivery is made from a CDN edge server, the average quality for end-users in ISP i is denoted by Q_i with $Q_i \geq 1$. Given that the CDN is integrated to ISP 1, Q_1 is set to the maximum available quality. Reaching this quality level comes at a price (e.g., in storage costs) for the CDN. But the CDN can reuse its infrastructure to deliver content to ISP 2 subscribers at no significant additional cost. This model is typical of companies like Google and Amazon, which have ultimately opened the delivery infrastructure that they first set for their own purpose [17]. Yet, the integrated entity may voluntarily downgrade the quality offered to ISP 2 subscribers. In the following, we are interested in the setting of Q_2 with respect to Q_1 with $1 \leq Q_2 \leq Q_1$. In summary, CDN-ISP 1 decides the quality Q_1 by investing into the infrastructure, and the quality Q_2 by (potentially) degrading the delivery to ISP 2 subscribers from the same infrastructure.

User preferences. Users are treated as a continuum. Follow-

ing some recent work [18], we assume that users choose which ISP (if any) to subscribe to according to a standard discrete choice model, more exactly a *logit* one, a type of model often used in economic studies. The (random) valuation for an end user for subscribing to ISP i ($i \in \{1, 2\}$) is assumed to be

$$v_i = k \ln(1/p_i) + \ln Q_i + \kappa_i,$$

to be compared to the valuation for choosing no service: $k \ln(1/p_0) + \kappa_0$, where

- κ_i ($i \in \{0, 1, 2\}$) are i.i.d. Gumbel-distributed random variables representing the variations among users in the service perception
- p_0 is a perceived cost for not subscribing to any ISP
- the valuation decreases with the service prices p_i ($i \in \{0, 1, 2\}$), k being a sensitivity parameter that we assume in the interval $(1, 2]$
- the valuation increases with the quality per download.

In our model, we use logarithmic functions, according to the Weber-Fechner law, which has been shown to properly describe the relationship between the magnitude of a physical stimulus and its perceived intensity within the human sensory system (see for example [19] for the validity of this representation in telecommunications).

From those valuations and standard discrete choice modeling computations, the probability of choosing ISP i equals

$$\frac{(1/p_i)^k Q_i}{(1/p_i)^k Q_i + (1/p_j)^k Q_j + (1/p_0)^k},$$

with $j \neq i$, which we multiply by D , the fixed base mass of users, to get the mass d_i of users subscribing to ISP i

$$d_i = D \frac{(1/p_i)^k Q_i}{(1/p_i)^k Q_i + (1/p_j)^k Q_j + (1/p_0)^k}. \quad (1)$$

Costs for the CDN. To achieve a given average delivery quality Q_1 for the users in ISP 1, the CDN has to invest in the infrastructure. Investing in capacities and servers implies a better quality for users through the access to content from closer edge servers. Since the most popular content is “cached” in edge servers and is received by users with a better quality than from origin servers, there is a bijection between storage capacity and average quality through the distribution of content popularities, so that we can express investment costs in terms of the targeted quality. Let $c(Q_1)$ be the cost incurred to provide quality Q_1 . We will study in Section IV the form of this function $c(\cdot)$ for the specific use-case of a CDN and a CP where the content popularity (i.e., the distribution of the number of downloads) follows a power-law and the investment cost grows linearly with the storage.

Global revenues. Overall, the revenue of ISP i from end-user subscriptions equals

$$R_i = d_i p_i, \quad (2)$$

and the net benefit of the CDN is

$$R_c = q(Q_1) d_1 + q(Q_2) d_2 - c(Q_1). \quad (3)$$

These revenues are the functions ISPs and CDN are willing to maximize. For the vertically-integrated CDN we need to look at $R_1 + R_c$ as the whole revenue. Some specific forms for the functions $c(\cdot)$ and $q(\cdot)$ will be considered in the next sections.

III. GAME DESCRIPTION

The different configurations (integrated CDN, independent CDN, no CDN) require analyzing a game. We consider here a Stackelberg game [20], with decisions taken at different time scales, but once and for all. We describe the order of decisions for each configuration.

A. With a Vertically-Integrated CDN

We first analyze the game when the CDN is controlled (and owned) by the same actor that controls the ISP 1. We consider actors ISP 2 and CDN-ISP 1, who both aim to maximize their revenues, R_2 and $R_1 + R_c$ respectively. ISP 2 plays with p_2 while CDN-ISP 1 plays with p_1 , Q_1 and Q_2 .

The game follows the following order of decisions:

- 1) CDN-ISP 1 decides the investment to achieve Q_1 maximizing R_1
- 2) The game on prices (p_1, p_2) is played by ISPs, still maximizing their revenues
- 3) The chosen degraded quality Q_2 is decided by CDN-ISP 1.

The game is analyzed by backward induction, i.e., each decision is assumed to be taken anticipating the consequences in the next steps. In other words, in order to solve this problem,

- 1) for any (Q_1, p_1, p_2) CDN-ISP 1 first computes $Q_2^*(Q_1, p_1, p_2)$ maximizing $R_1 + R_c$.
- 2) Then for any Q_1 , a game is played on p_1 and p_2 between ISPs maximizing their revenues using the expression $Q_2^*(Q_1, p_1, p_2)$ for Q_2 , leading to $p_1^*(Q_1)$ and $p_2^*(Q_1)$.
- 3) Finally, using $p_1^*(Q_1)$, $p_2^*(Q_1)$ and $Q_2^*(Q_1, p_1^*(Q_1), p_2^*(Q_1))$ for p_1 , p_2 and Q_2 , respectively, the value of Q_1 maximizing $R_1 + R_c$ is determined by CDN-ISP 1.

Note that prices p_i and qualities Q_i , though seeming being defined independently, are interdependent since quality Q_2 is chosen depending on prices, and prices are determined through a game, which depends on each corresponding maxima quality level Q_1 and is played anticipating Q_2 .

Unfortunately, we have not been able to prove existence and uniqueness of a Nash equilibrium in the pricing game between ISPs, but our numerical investigations have always resulted in a unique equilibrium.

On the determination of Q_2 (smallest time scale), one can remark the following trade-off: increasing Q_2 makes ISP 2 more attractive, hence induces more demand at ISP 2, as well as a higher unit revenue $q(Q_2)$ for the CDN part of the integrated CDN-ISP 1 entity. However, ISP 2 being more attractive, the CDN-ISP 1 suffers from the competition, which generates less revenues for the ISP part of CDN-ISP 1.

B. Without a CDN

In the default configuration, there is no CDN, which means that both qualities $Q_i (i \in \{1, 2\})$ have the default value, i.e. $Q_1 = Q_2 = 1$. We then end up with only the second level of the game of the previous subsection, i.e., the pricing game.

For such a configuration, the solution is known as a special case of the model in [18] where the competition between ISPs is analyzed under paid or free peering between ISPs. At the unique Nash equilibrium, when $1 < k \leq 2$ we have

$$p_1 = p_2 = \left(\frac{2-k}{k-1}\right)^{1/k} p_0,$$

and each ISP earns a revenue

$$R_1 = R_2 = D(2-k)^{1/k} \frac{(k-1)^{1-1/k}}{k} p_0.$$

C. With an Independent CDN

When the CDN is independent, it does not favor any ISP, and provides the same quality $Q_1 = Q_2 = Q$ to both of them. Here,

- 1) the quality level Q is decided first,
- 2) and then the game on p_1 and p_2 is played.

It is again analyzed by backward induction. For a given quality level Q , the unique Nash equilibrium of the pricing game is

$$p_1 = p_2 = \left(\frac{2-k}{k-1}Q\right)^{1/k} p_0.$$

from a simple adaptation of the result with no CDN.

The revenue R_c of the CDN as a function of Q is then, anticipating those Nash equilibrium prices,

$$R_c = q(Q)D \frac{2Q/p_1^k}{2Q/p_1^k + 1/p_0^k} - c(Q).$$

In the next section, we express analytically the best quality to offer based on a use-case of CDN, CP and ISP behaviors.

IV. A CASE STUDY

In this section, we compare the three cases based on numerical computations, under some extra assumptions for the forms of the CDN price (paid by CPs) and of the CDN cost (supported by the integrated entity or the independent CP), which we specify now.

A. Specific cost and price functions

1) *CDN cost for providing quality*: The infrastructure of CDN is made of some *origin servers*, which store the whole set of data related to the CPs, and multiple *edge servers*, which are located close to the end-users. The management of a CDN can be simplified as being the management of which data should be stored in which edge servers so that the number of requests from end-users that are treated by the edge servers is maximized.

We provide hereafter an example of the application of our model to a case where the CDN manages one edge-server and one origin server to serve both ISP 1 and ISP 2.

Let C be the capacity of the edge-server. In a power law popularity distribution, the x -th most popular unit of content is downloaded $Bx^{-\alpha}$ times, where B and α are parameters with usually $0.5 < \alpha < 1$. Let V be the whole volume of content of the CP. The *hit-ratio* H_1 is the proportion of requests coming from end-users in ISP 1 that are successfully treated by the edge-server:

$$H_1 = \frac{\int_0^C Bx^{-\alpha} dx}{\int_0^V Bx^{-\alpha} dx} = (C/V)^{1-\alpha}.$$

The data delivery from the edge server results in a gain of quality, denoted by λ . We set the default quality of the delivery from the origin server to 1. The average quality is

$$Q_1 = \lambda H_1 + 1(1 - H_1).$$

To achieve a given quality Q_1 for the delivery to end-users in ISP 1, the CDN should find the solution C of:

$$Q_1 = H(\lambda - 1) + 1 = (\lambda - 1)(C/V)^{1-\alpha} + 1,$$

that is: $C = V \left(\frac{Q_1 - 1}{\lambda - 1}\right)^{1/(1-\alpha)}$. With unit capacity cost β , the cost incurred to provide quality Q_1 is therefore

$$c(Q_1) = \beta V \left(\frac{Q_1 - 1}{\lambda - 1}\right)^{1/(1-\alpha)}. \quad (4)$$

2) *CDN quality-related price charged*: The price $q(\cdot)$ that the CDN charges CPs for content delivered is assumed to be of the form:

$$q(Q_i) = (Q_i - 1)^\gamma,$$

with γ a parameter reflecting the sensitivity of CPs to quality; determining its specific value is beyond the scope of this paper, we assume γ fixed here. When $Q_i = 1$, we obtain $q(Q_i) = 0$, which means that the CPs do not pay when the CDN does not provide any improved quality with respect to a simple delivery on the public Internet.

B. Preliminary remarks

Integrated case. Even at the lowest time scale, determining analytically Q_2 in terms of Q_1 , p_1 and p_2 seems intractable. We therefore resort to a numerical analysis whose results will be presented in the next subsection.

No-CDN case. That case does not depend on the cost and price functions, and is entirely treated in Subsection III-B.

Independent CDN case. We can compute the optimal value of Q as follows: when $\gamma \leq 1$ and $\alpha \in (0, 1)$, the independent CDN revenue is a strictly concave function of Q , and from the first-order condition we obtain

$$Q = 1 + \left((1-\alpha)(\lambda-1)^{\frac{1}{1-\alpha}} \frac{D\gamma}{\beta V} \frac{1}{1 + \frac{2-k}{2k-2}} \right)^{\frac{1}{1-\alpha-\gamma}}. \quad (5)$$

From that expression, we observe that the optimal quality increases with the price sensitivity k , and decreases with the total data volume V (since a given quality then implies more caching space) and with the cache space price α .

Figure 2 displays the revenue R_c of the CDN in terms of Q with the corresponding solutions at the next levels of the game, for the parameter values given in Table I. From (5), the optimal quality to propose is 2.36.

Table I
PARAMETER VALUES

population size D	100.0
volume of content of CPs V	1000.0
perceived cost for not subscribing to any ISP p_0	10.0
user sensitivity to subscription price k	1.2
quality gain when download from edge server λ	5.0
popularity skewness of content α	0.8
storage unit cost β	2.0
pricing factor for delivery quality γ	1.0

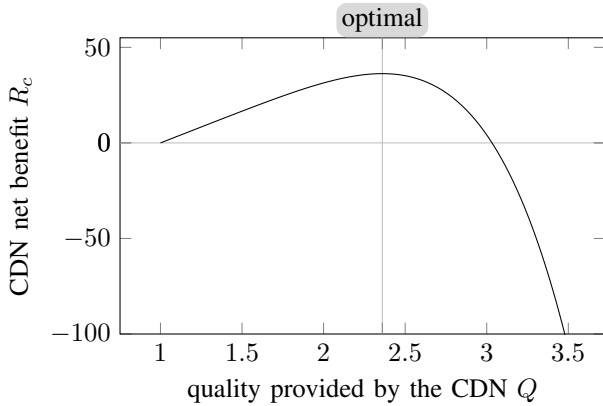


Figure 2. Revenue of the independent CDN in terms of the provided quality Q

C. Numerical analysis

In the remaining, we use the default parameters given in Table I. We will focus on two different populations of user, one being highly sensitive to ISP subscription prices ($k = 1.8$) while the other is less sensitive ($k = 1.2$). In the integrated case the optimizations, and best-response computations to get the Nash equilibrium, are performed numerically by looking at optimal values on grids.

1) *Optimal integrated-CDN investment*: Figure 3 displays the revenue $R_1 + R_c$ of the integrated entity in terms of the quality Q_1 with the corresponding solutions at the next levels of the game. The figure illustrates the two regimes that can occur: either CDN-ISP1 decides to set $Q_2 = 1$ to favor its ISP attractiveness, or it sets $Q_2 = Q_1$ to favor CDN revenues. When the population is not price-sensitive ($k = 1.2$), the first regime is preferred by CDN-ISP 1, and the optimal quality for ISP 1 requests is $Q_1 = 2.94$; on the other hand for a price-sensitive population ($k = 1.8$), the second regime is chosen (and the optimal quality is $Q_1 = 2.80$). The resulting other decisions are given in Tables II and III.

2) *Comparison between the three configurations*: Let us first compare the outputs given in Table II in the case of a population that is not price-sensitive. Regarding the delivery

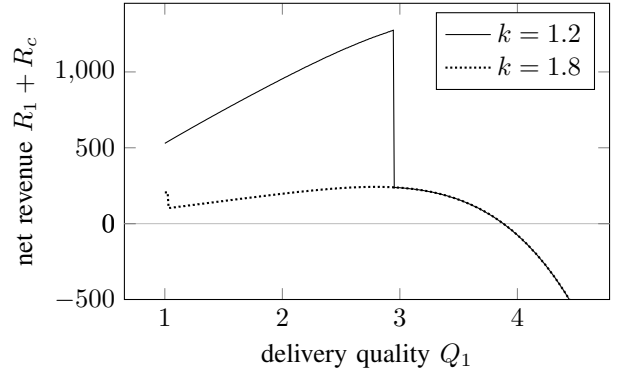


Figure 3. Revenue of the CDN-ISP 1 entity when varying quality Q_1 . Before the revenue drop, the optimal quality Q_2 is 1, while CDN-ISP 1 should set $Q_2 = Q_1$ after that drop.

quality, the integrated CDN decides here not to have ISP 2 benefit from the improved quality that it can provide (i.e., it sets $Q_2 = 1$). Compared to setting a higher Q_2 , this strategy leads to an unprofitable CDN activity. But this strategy gives ISP 1 a competitive advantage over ISP 2: because of the higher quality, ISP 1 can set a higher price than its competitor and still attract more demand.

When comparing with the configuration without a CDN, an integrated CDN increases demand at ISP 1, but does not significantly impact demand at ISP 2. Total demand is then increased. On the other hand it induces a lower price at ISP 2 which has to be more attractive on price to compensate the lower quality, hence a (slightly) smaller revenue for this ISP 2. The revenue of ISP 1 is significantly larger.

Surprisingly, when comparing an integrated to an independent CDN from the point of view of users, the integrated CDN is preferable since total demand is larger. Moreover, user choices are richer: one can select a high-quality service at a high price, or a cheaper low-quality service. However, the revenues of ISP 2 are significantly lower in the integrated configuration while those of the integrated entity are larger. Overall, a regulator may express concerns about the economic fairness of the integrated configuration, but this has to be balanced with the increase of demand it brings.

Table III presents outputs when the population is price-sensitive. The best strategy for the integrated CDN is to provide the best delivery quality to ISP 2 users and gain more thanks to the CDN activity, which is here profitable. Interestingly, the offered quality is the same for the integrated and independent cases: the goal for CDN-ISP 1 is indeed to attract as much traffic as possible, regardless of the ISP chosen by users.

The comparison between the integrated CDN configuration and the other configurations yields counter-intuitive observations. Since the population is price-sensitive, the actors are tempted to enter into a price war. Interestingly, the price war is fiercer in the integrated case than in the independent case. The main reason is that the CDN revenue in the integrated case, incentivizes CDN-ISP 1 to reduce ISP prices in order to

Table II
COMPARISON OF THE THREE CDN CONFIGURATIONS, $k = 1.2$

	price p_1	price p_2	quality Q_1	quality Q_2	CDN benefits R_c	revenues $R_1(+R_c)$	revenues R_2	demand d_1	demand d_2	overall $d_1 + d_2$
Integ CDN	68.12	30.94	2.94	1.00	-17.38	1274.82	513.93	18.97	16.61	35.58
Indep CDN	64.94	64.94	2.36	2.36	36.25	1082.31	1082.31	16.67	16.67	33.33
No CDN	31.75	31.75	1	1	0	529.13	529.13	16.67	16.67	33.33

Table III
COMPARISON OF THE THREE CDN CONFIGURATIONS, $k = 1.8$

	price p_1	price p_2	quality Q_1	quality Q_2	CDN benefits R_c	revenues $R_1(+R_c)$	revenues R_2	demand d_1	demand d_2	overall $d_1 + d_2$
Integ CDN	2.25	2.05	2.80	2.80	141.10	243.03	109.68	45.38	53.52	98.90
Indep CDN	8.10	8.10	2.74	2.74	123.52	359.92	359.92	44.44	44.44	88.89
No CDN	4.63	4.63	1	1	0	205.75	205.75	44.44	44.44	88.89

attract more traffic. It results at equilibrium in smaller revenues for *both* ISPs, for which the preferred scenario is then an independent CDN. This result explains why some powerful ISPs prefer to externalize their CDN activity. In terms of user preferences, the integrated CDN configuration is again the one that brings the largest demand, which then could be satisfactory for a regulator, something which is not obvious at first sight.

We restrict the remaining of our analysis to the impact of the price-sensitivity k on the main economic and social indicators: the quality chosen by the CDN for the users (see Figure 4), the revenues of the actors (see Figure 5), and the total demand of the population (see Figure 6).

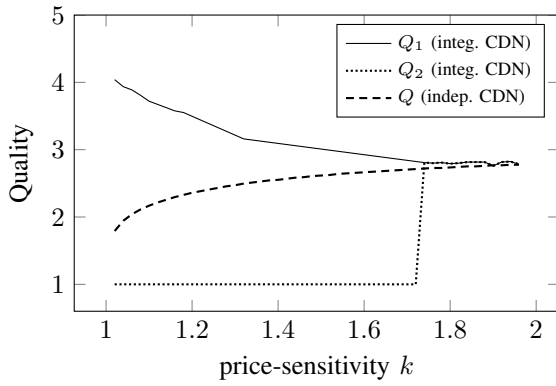


Figure 4. Optimal qualities chosen by the integrated CDN with respect to the price-sensitivity k

Figure 4 highlights the regime change of the integrated entity, which opts for favoring the CDN-related revenues for a price-sensitivity k larger than 1.7. Then, the users of both ISPs experience the same delivery quality, which corresponds approximately to what could be offered by an independent CDN. With low price sensitivities, customers of ISP 1 (resp. ISP 2) are provided a better (resp. lower) quality in the integrated case than with an independent CDN.

Figure 5 shows that for a small k , the integrated scenario

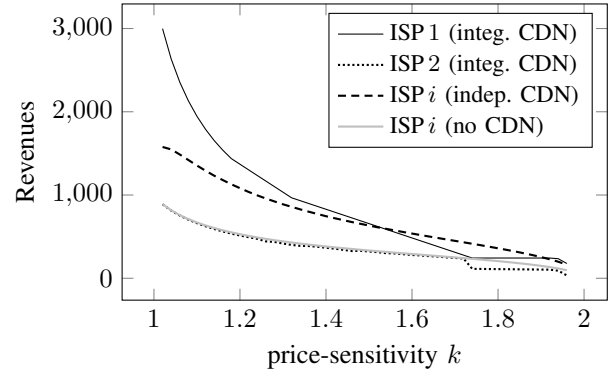


Figure 5. ISPs Revenues with respect to the price-sensitivity k

is preferred by ISP 1, but that it is not true for a price-sensitivity k larger than 1.55. More generally, the integrated entity gets significant gains only for populations that are not price-sensitive. However, the integrated entity severely impacts its ISP competitor by integrating a CDN, which brings a critical economic fairness issue.

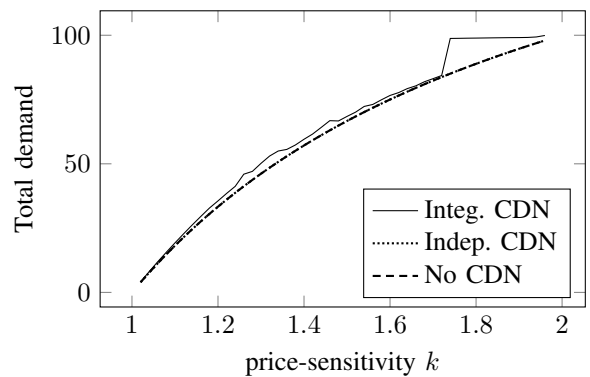


Figure 6. Total demand with respect to the price-sensitivity k

Finally, Figure 6 shows that a regulatory body that exclusively regards total demand should encourage the integra-

tion of a CDN by an ISP since the integrated configuration provides a (slightly) higher overall user demand, for any price sensitivity. This may be due to a phenomenon called double marginalization [21], where several intermediaries in the delivery chain (here, the ISPs and the CDN) want to make profit out of the market, leading to higher prices and a loss of overall attractiveness.

V. CONCLUSIONS

This paper investigates the consequences of having a CDN activity controlled by one ISP, in particular focusing on the competition among ISPs. We illustrate that when users are not too sensitive to prices, the vertically integrated entity (CDN+ISP) benefits from its position to favor its ISP consumers over the competing ISPs in order to attract more demand. Inversely, with price-sensitive users the CDN high-quality service is offered to all users, the integrated entity taking most of its revenues from content providers, which pay for quality.

Among the nonintuitive results of our model, we observe that an ISP can prefer the presence of an independent CDN rather than controlling itself a CDN service. Also, from the user point of view an integrated CDN is the best scenario, when compared with a no-CDN configuration (with worse service quality) and even an independent-CDN configuration (which leads to the so-called double-marginalization phenomenon).

This work is a first step in the study of the impact of vertically-integrated CDNs; we plan to strengthen and extend it in several directions. First the analysis is currently based on synthetic values of parameters; we plan to collect real-life data to be fitted in our model. Furthermore a key assumption is that we consider a single CDN; the case of multiple CDNs which can be concurrent needs also to be addressed and could induce a further interest for vertical integration.

ACKNOWLEDGMENT

The authors acknowledge the support of Orange. Note however that the content presented in this paper is solely the responsibility of the authors.

REFERENCES

- [1] K. Pires and G. Simon, "YouTube live and Twitch: a tour of user-generated live streaming systems," in *Proceedings of ACM MMSys*, 2015.
- [2] R. T. B. Ma and D. Towsley, "Cashing in on caching: On-demand contract design with linear pricing," in *Proc. of CoNEXT*, 2015.
- [3] P. Maillé, K. Pires, G. Simon, and B. Tuffin, "How Neutral is a CDN? An economic approach," in *Proceedings of CNSM*, 2014.
- [4] P. Maillé, G. Simon, and B. Tuffin, "Impact of revenue-driven CDN on the competition among network operators," in *Proceedings of CNSM*, 2015.
- [5] P. Maillé, G. Simon, and B. Tuffin, "Toward a net neutrality debate that conforms to the 2010s," *IEEE Communications Magazine*, vol. 54, no. 3, pp. 94–99, 2016.
- [6] P. Maillé and B. Tuffin, *Telecommunication Network Economics: From Theory to Applications*. Cambridge University Press, 2014.
- [7] N. Laoutaris, G. Smaragdakis, R. Stanojevic, P. Rodriguez, and R. Sundaram, "Delay-tolerant bulk data transfers on the internet," *IEEE/ACM Trans. Netw.*, vol. 21, no. 6, pp. 1852–1865, 2013.
- [8] Cisco, "Wholesale Content Delivery Networks: Unlocking New Revenue Streams and Content Relationships," Cisco, Tech. Rep. 710667, 2011.
- [9] K. Harrigan, *Vertical Integration, Outsourcing, and Corporate Strategy*. BeardBooks, 1983.
- [10] N. Slack, *Vertical Integration*. John Wiley & Sons, Ltd, 2015. [Online]. Available: <http://dx.doi.org/10.1002/9781118785317.weom100176>
- [11] K. Hosanagar, R. Krishnan, M. Smith, and J. Chuang, "Optimal pricing of content delivery network (CDN) services," in *Proc. of the Annual Hawaii Int. Conf. on System Sciences*, 2004.
- [12] K. Hosanagar, J. Chuang, R. Krishnan, and M. Smith, "Service adoption and pricing of content delivery network (CDN) services," *Management Science*, vol. 54, no. 9, pp. 1579–1593, 2008.
- [13] K. Hosanagar, "CDN pricing," in *Content Delivery Networks*, ser. Lecture Notes Electrical Engineering, R. Buyya, M. Pathan, and A. Vakali, Eds. Springer, 2008, vol. 9.
- [14] P. Faratin, "Economics of overlay networks: An industrial organization perspective on network economics," in *Proc. of ACM EC*, 2007.
- [15] H. Lee, L. Duan, and Y. Yi, "On the competition of CDN companies: Impact of new telco-CDNs' federation," in *Proc. of WiOpt*, 2016.
- [16] T. Hau and W. Berenner, "Vertical platform interaction on the internet: How ISPs and CDNs interact," in *Proc. of ECIS*, 2010.
- [17] L. A. Barroso and U. Hölzle, *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*. Morgan & Claypool Publishers, 2009.
- [18] P. Coucheny, P. Maillé, and B. Tuffin, "Network neutrality debate and ISP inter-relations: Traffic exchange, revenue sharing, and disconnection threat," *Netnomics*, vol. 15, no. 3, pp. 155–182, 2014.
- [19] P. Reichl, R. Schatz, and B. Tuffin, "Logarithmic laws in service quality perception : Where microeconomics meets psychophysics and quality of experience," *Telecommunication Systems*, vol. 52, no. 2, pp. 587–600, 2013.
- [20] M. Osborne and A. Rubinstein, *A Course in Game theory*. MIT Press, 1994.
- [21] J. Tirole, *The Theory of Industrial Organization*. MIT Press, 1988.