

How Neutral is a CDN? An Economic Approach

Patrick Maillé, Karine Pires, Gwendal Simon
Institut Mines-Telecom / Telecom Bretagne
Rennes, France
Email: {first}.{last}@telecom-bretagne.eu

Bruno Tuffin
Inria, France
Email: bruno.tuffin@inria.fr

Abstract—The growing importance of Content Delivery Network (CDN) in the value chain of content delivery raises concerns about the “neutrality” of these players. We propose in this paper a model to analyze the impact of revenue-oriented CDN management policies on the fairness of the competition among two content providers that use CDN services to deliver contents. We show that there exists a unique optimal revenue-maximizing policy for a CDN actor—the dimensioning and allocation of its storage capacity—that depends on prices for service/transport/storage, and on the distribution of content popularity. Using data from the analysis of traces from two major content providers (YouTube Live and justin.tv), we remark that a CDN remains a relatively neutral actor even when one of the content providers it serves tries to monopolize the CDN storage space by implementing an aggressive policy to harm its competitors.

I. INTRODUCTION

The term *Content Delivery Network* (CDN) refers to both an *infrastructure* designed to deliver content at large scale over an underlying network, and the *economic actor* providing that service. This work focuses on the economic actor.

CDN have a huge economic weight (the annual revenues of Akamai, the CDN leading company, are over two billion dollars), and a growing impact on the Internet ecosystem: *i*) CDN activities affect the traffic exchanged between network providers, and consequently their economic relationships [1], [2]; *ii*) on many aspects (per-volume charging, connectivity service) CDN actors compete with transit providers, which explains why some major transit network operators such as Level 3 have shifted a fraction of their activities to CDN; and *iii*) other actors in the value chain of content delivery have started developing a CDN activity, including Internet Service Providers (ISPs), content providers, and equipment vendors [3], [4]. This fast-moving and business-driven environment exacerbates the concerns among user and regulation communities regarding service quality and economic fairness, epitomized by the net neutrality debate [5], [6], [7], [8].

The scientific literature provides models and analyses of the interactions between content providers and ISPs in order to address network neutrality, and sometimes to propose regulation remedies [9], [10], [11], [12], [13], but the role of CDNs is barely mentioned. To the best of our knowledge, the only official report mentioning CDNs is from the Norwegian regulator [14], where it is stated that “*the ordinary use of CDN servers is not a breach of net neutrality*”. In this paper, we show that CDNs that are “normally” managed (i.e., by

rational actors) can nevertheless lead to differences in the average Quality of Experience (QoE) among content providers, which goes against neutrality principles.

More generally, the performance analysis community has barely considered the economics of CDN actors so far. Among the few notable works, we can mention [15], [16], [17] where the author(s) consider a (single) CDN and a time period of interest. The best pricing strategy is studied, but both the complex relationships between actors and their consequences on fairness and social welfare are ignored.

In this paper, we focus on the management problems faced by a CDN having to dimension and optimally use its infrastructure, sharing it among its clients (content/service providers) so as to maximize its revenue. We propose a model to analyze the behavior of a profit-maximizing CDN, and assess the impact of a CDN policy on the quality perceived by users and on the fairness among content providers. We illustrate these theoretical results with an analysis based on real data from two major service providers, which we artificially make compete for the resources of a CDN. We show that a CDN implementing a revenue-maximizing policy tends to favor incumbent content providers, but at an extent that is not dramatic, even if the said incumbent tries to take advantage of the profit-driven CDN policy by over-paying for a better service.

II. MODEL

A CDN is a *multi-tenant* infrastructure: its resources are shared among multiple Service Providers (SP). For simplicity reasons, we consider here two SPs (referred to as SP1 and SP2) but the model can be extended to more SPs. We depict the configuration/topology in Figure 1.

We distinguish two classes of CDN resources: some privileged resources that are located close to the clients in the ISP, and the remaining resources, which often correspond to the origin data-centers. Since the resources that are the most often offered by CDN are storage, we will hereafter call *cache* the privileged resources in the ISP and we will abusively use wording related to storage management. Note however that the services that are offered by today’s CDNs extend to other types of resources, typically computing. Figure 1 represents only one ISP, but multiple ISPs can be considered, each one being studied independently.

The economic flows involving the CDN are as follows:

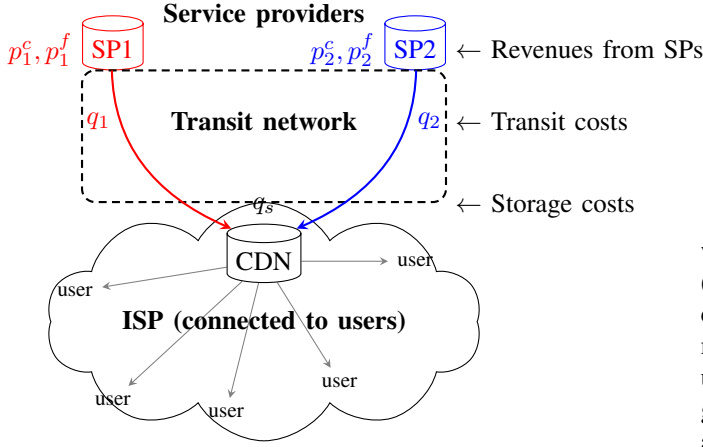


Fig. 1. Costs and revenues for a CDN located within an ISP's network.

Revenues. Each SP subscribes to the CDN service to reach its customers. The CDN charges the SPs a different price per unit of data volume delivered to users, according to whether users are served from the cache server (unit price p_i^c for SP i , $i = 1, 2$) or from the SP origin data-center, hence with lower average QoE (unit price $p_i^f < p_i^c$ for SP i).

Costs. The CDN is responsible for sending the data to users (those covered by the considered ISP). There are two cases. If the data are taken from the origin data-center of SP i , the cost is the transit cost q_i per unit of volume for the CDN (which can be low if the CDN owns the transit network, but large otherwise). Remark that those transit costs can differ among SPs (i.e., $q_1 \neq q_2$) since the path to reach the ISP of interest may differ. The second case is when the data are delivered by the CDN cache, the quality experienced by users is better, and no transit costs are incurred. On the other hand, storage in the CDN cache incurs a unit cost q_s .

If we consider those prices fixed (from long-term contracts), the decision variables of the CDN are:

- The **capacity** C of the server in the ISP
- The implemented **caching strategy**, i.e., the management of the storage space in the cache. With two service providers, the only decision variable for the CDN regards the choice of whose content to favor in the cache, summarized by the volume $C_1 \leq C$ of cached SP1 content (the volume of SP2 content cached being $C_2 = C - C_1$).

We do not deal with the extensive literature on the subject that involves time variations of the download frequency of content items. We rather consider a static problem, with content popularity values as *constant and known to the CDN operator*. Let us denote by $F_i(x), i = 1, 2$ the minimum download frequency (number of requests per time unit) for the x most popular units of content of SP i , and assume that F_i is continuous and strictly decreasing. Knowing the popularity values, the CDN stores the content of each provider that yields the largest revenues, which may result in an “unfair” strategy with respect to SPs.

The *incomes* of the CDN from SP payments equal

$$\begin{aligned} & \sum_{i=1}^2 \left(p_i^c \int_{x=0}^{C_i} F_i(x) dx + p_i^f \int_{x=C_i}^{V_i} F_i(x) dx \right) \\ &= \sum_{i=1}^2 \left(p_i^f \bar{G}_i + (p_i^c - p_i^f) G_i(C_i) \right) \end{aligned}$$

where V_i is the total volume of content proposed by SP i ($i = 1, 2$), $G_i(y) := \int_{x=0}^y F_i(x) dx$ is the cumulated user download throughput from requests for the volume y of the most popular content from SP i , and $\bar{G}_i = G_i(V_i)$ is the total user download throughput of SP i content. Without loss of generality, we ignore content with no demand, so that we can assume $F_i(x) > 0$ for all $x < V_i$ ($i=1,2$).

Storage costs equal $q_s C$. We assume $C \leq \min(V_1, V_2)$, i.e., the CDN cannot cache all the content from any SP.

For the *transit costs*, we neglect the one-shot costs for the content stored in the CDN cache: therefore transit costs only correspond to content that is not in the cache, and for each SP they are proportional to the aggregated download rate for that content. Since $C_2 = C - C_1$, the total transit costs equal

$$q_1(\bar{G}_1 - G_1(C_1)) + q_2(\bar{G}_2 - G_2(C - C_1)).$$

Overall, the net revenue of the CDN per time unit is

$$R(C, C_1) = \sum_{i=1}^2 \left(r_i^f \bar{G}_i + (p_i^c - r_i^f) G_i(C_i) \right) - q_s C \quad (1)$$

with $r_i^f := p_i^f - q_i$ for $i = 1, 2$. We limit ourselves to $r_i^f > 0 \forall i$ to ensure that the CDN makes some non-negative revenue.

Finally, let us denote the quality experienced by users by Q_c for the cached content, and by $Q_{f,i} < Q_c$ for content retrieved from SP i . The average user experienced quality is then

$$Q^{\text{tot}} = \frac{1}{\bar{G}_1 + \bar{G}_2} \sum_{i=1}^2 \left(Q_c(G_i(C_i) + Q_{f,i}(\bar{G}_i - G_i(C_i))) \right).$$

III. MAXIMIZING THE CDN REVENUE

In our analysis, we first focus on the caching strategy, i.e., determine the revenue-maximizing sharing of the storage space C (treated as fixed) among SP1 and SP2 content. Then we discuss the optimal value of the cache capacity C .

A. Whose content to cache?

We assume the total storage capacity C is fixed, and look for the best caching strategy decision (the value C_1^{opt} of C_1 maximizing the net revenue in (1), where $C_2 = C - C_1$).

By construction, each function G_i ($i = 1, 2$) is continuously differentiable (with derivative F_i), strictly increasing and strictly concave on $[0, V_i]$, hence $R(C, C_1)$ is a strictly concave function of C_1 for C fixed. The first-order optimality condition is thus sufficient, and the optimal C_1 equals:

- C if $(p_1^c - r_1^f) F_1(C) \geq (p_2^c - r_2^f) F_2(0)$
- 0 if $(p_1^c - r_1^f) F_1(0) \leq (p_2^c - r_2^f) F_2(C)$

- the unique solution in $(0, C)$ of

$$\frac{F_1(x)}{F_2(C-x)} = \frac{p_2^c - r_2^f}{p_1^c - r_1^f} \quad \text{otherwise.} \quad (2)$$

It therefore exists and is unique. Remark that for given popularity distributions, the optimal C_1 then only depends on the ratio $\frac{p_2^c - r_2^f}{p_1^c - r_1^f} = \frac{p_2^c + q_2 - p_2^f}{p_1^c + q_1 - p_1^f}$. Due to the decreasingness in x of the left-hand term in (2), the optimal C_1 decreases with the value of that ratio.

Remark also that when prices are fixed, the solution of (2) strictly increases with C : take $\tilde{C} > C$, the corresponding optimal values C_1^{opt} and \tilde{C}_1^{opt} of C_1 must satisfy

$$F_1(C_1^{\text{opt}})F_2(\tilde{C} - \tilde{C}_1^{\text{opt}}) = F_1(\tilde{C}_1^{\text{opt}})F_2(C - C_1^{\text{opt}}). \quad (3)$$

Assuming $\tilde{C}_1^{\text{opt}} \leq C_1^{\text{opt}}$ leads to $F_1(C_1^{\text{opt}}) \leq F_1(\tilde{C}_1^{\text{opt}})$ and $F_2(\tilde{C} - \tilde{C}_1^{\text{opt}}) < F_2(C - C_1^{\text{opt}})$, contradicting (3).

Example. Following the literature on content popularity [18], let us consider a Zipf or power-law distribution of the request rates among pieces of content: $F_i(x) = A_i x^{-\alpha}$ with $\alpha > 0$, for $x > x_{i,\min} > 0$. The values $x_{i,\min}$ indicate the domain of validity of the power law, and we assume they are small enough, so that $\frac{F_2(C - x_{1,\min})}{F_1(x_{1,\min})} < \frac{p_1^c - r_1^f}{p_2^c - r_2^f} < \frac{F_2(x_{2,\min})}{F_1(C - x_{2,\min})}$, and thus the optimal value of C_1 is in $(x_{1,\min}, C - x_{2,\min})$. Solving (2) then leads to C_1 being the solution x of

$$\frac{x}{C-x} = \left(\frac{A_1 p_1^c + q_1 - p_1^f}{A_2 p_2^c + q_2 - p_2^f} \right)^{1/\alpha}, \quad \text{which gives}$$

$$C_1^{\text{opt}} = \frac{C}{1 + \left(\frac{A_2 p_2^c + q_2 - p_2^f}{A_1 p_1^c + q_1 - p_1^f} \right)^{1/\alpha}}. \quad (4)$$

B. Dimensioning the cache

We can similarly determine the optimal storage capacity C for the CDN, by differentiating $R(C, C_1^{\text{opt}})$ in terms of C , with C_1^{opt} a function of C . Rewriting the conditions not to end up with (2):

$$F_1(C) \geq k F_2(0) \quad \text{or} \quad F_2(C) \geq \frac{1}{k} F_1(0)$$

with $k = (p_2^c - r_2^f)/(p_1^c - r_1^f)$, we remark that none is satisfied (since the F_i are decreasing functions) when

$$C > \max(F_2^{-1}(k F_1(0)), F_1^{-1}(F_2(0)/k)), \quad (5)$$

in which case the solution of C_1^{opt} is inside $(0, C)$. For C smaller, it may happen that the optimal value C_1^{opt} is 0 or C . Assuming (5), the envelope theorem yields

$$\begin{aligned} \frac{\partial R(C, C_1^{\text{opt}}(C))}{\partial C} &= (p_2^c - r_2^f) F_2(C - C_1^{\text{opt}}(C)) - q_s \\ &= (p_1^c - r_1^f) F_1(C_1^{\text{opt}}(C)) - q_s, \end{aligned}$$

where the last equality comes from (2). From that last expression and due to the strict increasingness of C_1^{opt} in C , the revenue is a strictly concave function of C for C sufficiently large. From (1), it is also strictly concave when $C_1^{\text{opt}} \in \{0, C\}$.

Since $R(C, C_1^{\text{opt}}(C))$ is differentiable for all C (its ‘‘interior’’ derivative when $C_1^{\text{opt}}(C)$ tends to 0 or C being equal to the derivative with a fixed C_1^{opt}), it is then strictly concave over the whole interval $[0, C]$. Hence, since the derivative of R gets negative for C sufficiently large, there exists a unique cache capacity C maximizing revenue (that is strictly positive if q_s is not too large, *i.e.*, if $q_s < \max_{i=1,2}(p_i^c - r_i^f)F_i(0)$).

Example. Considering again the case of power-law distributions, the above derivative is

$$(p_2^c - r_2^f) A_2 \left(C - \frac{C}{1 + \left(\frac{A_2 p_2^c - r_2^f}{A_1 p_1^c - r_1^f} \right)^{1/\alpha}} \right)^{-\alpha} - q_s,$$

which gives the optimal dimensioning of the storage space

$$C = \frac{((p_1^c - r_1^f) A_1)^{1/\alpha} + ((p_2^c - r_2^f) A_2)^{1/\alpha}}{q_s^{1/\alpha}}.$$

IV. ANALYSIS

Due to lack of space, we restrict the numerical analysis of our model to one situation closely linked to the problems of fairness (and neutrality) of CDNs, where we apply our model to popularity distributions obtained from real traces, and consider two SPs competing for one CDN.

We study two *user-generated live video aggregators*. These service providers offer a service such that anybody can become a *streamer*, who uploads a video stream to the aggregator, which is then in charge of preparing and delivering the video to a potentially wide population. Two main players compete: (i) an *incumbent*, namely *justin.tv*, which has been a well-established service for years, with a stable population of engaged streamers (more than five thousands simultaneously broadcasting at any time), and (ii) a *challenger*, namely YouTube Live, which has recently released this new feature to the regular YouTube service. The population of streamers of YouTube Live is one order of magnitude smaller than *justin.tv* but, at peak hours, both services have approximately the same population of viewers. In the following, SP2 refers to *justin.tv* while SP1 is YouTube Live. For both services, we extract from their public Application Programming Interface (API) the traces of the activities from January, 6th to January 31st 2014¹. In the following, we study one randomly chosen date and we abusively consider that both services use the same CDN to deliver their live streams (for such service, the resources that the CDN offers are transcoding and delivering in the access network).

Our goal is to highlight the role of CDN in three representative scenarios: (i) both service providers pay the same price for the CDN service. The transit costs are the same. In this *regular* scenario, the main question is whether the dominance of the incumbent prevents the growth of the challenger. (ii) the incumbent player deploys an aggressive strategy where it pays ten times what its competitors pays for the CDN service. It is one of the most critical question in the net neutrality

¹Dataset available at <http://is.gd/tM6xmH>

debate: can a well-established player prevents one competitor from growing? Finally (iii) the challenger is now the one that is aggressive. Regarding the parameters adopted in the scenarios, we extracted from the traces the information about the videos, meaning V_i , G_i and \bar{G}_i . Based on CDNs and Amazon pricing², the remaining parameters are: $p_i^f = 0.005$, $p_i^c = 0.5$, $q_s = 0.0000053$, $q_i = 0.94$, $Q_c = 2$, $Q_{f,i} = 0.5$ and $C = 50$. For scenarios (ii) and (iii), we defined $p_1^c = 5$ and $p_2^c = 5$ respectively.

We show in Figure 2 the QoE experienced by the end-users of both service providers with regard to the evolution of the ratio of the cache that is filled with SP1 content (recall that SP1 is the challenger YouTube). The QoE of SP1 users is represented by Q_1 (black line), while users on SP2 by Q_2 (gray line). We show with thin vertical lines the optimal values of $\frac{C_1}{C}$ for the three considered scenarios.

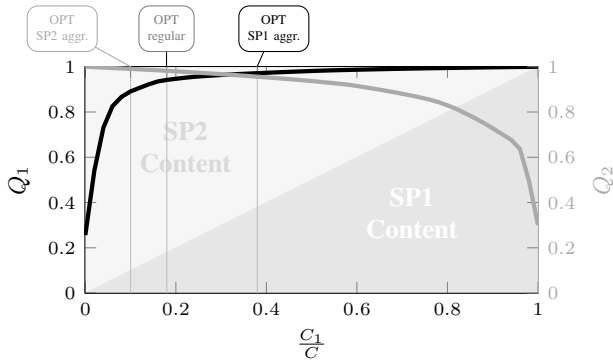


Fig. 2. Normalized quality of experience for both service providers SP1 and SP2 according to the ratio of the cache filled with content from SP1

Our main observation is that, due to the heterogeneity of video popularity, the impact of aggressive strategies is limited in all cases. By choosing to maximize its revenues, the CDN serves more content from SP2 in the regular scenario, which in turn leads to a better overall QoE for users of SP2. But the QoE remains excellent for SP1 as well (more than 0.9 of the best possible). More interestingly, both aggressive policies are not worth the price. In both cases, the CDN adjusts the ratio $\frac{C_1}{C}$ accordingly to maximize its revenues, but in both cases, the overall QoE of the competitor is not significantly affected. Even when the incumbent player pays ten times the price paid by the challenger, the users of the latter service have a QoE which is more than 0.8 of the best possible.

V. CONCLUSION

We propose in this paper a model to analyse of the policy that a profit-driven CDN should implement. This model is especially significant with regard to the multiple recent debates about network neutrality. To the best of our knowledge, this is the first attempt to model CDN from an economic standpoint with the ambition to understand the impact of CDN on the content delivery market. Our goal was to illustrate the need

for such an analysis. The theoretical model and analysis are complemented with a real-case study.

This paper opens perspectives. Fairness can be further analyzed thanks to our model. We would also like to study more generic versions of this model with multiple ISPs, players and resources within the CDN. The competition among players, and ways to regulate it toward the benefit of the whole population, are among the very first studies that we envision. We also wish to extend the analysis to more than two SPs and one CDN (and/or more broadly applicable settings). The present results apply to this specific configuration, but the analysis can be extended to any number of SPs. Note however that the outcome is case (and parameters') dependent, in particular the case of several CDNs would raise competition issues that are not considered in the paper.

REFERENCES

- [1] A. Kovaks, "Internet peering and transit," Technology Policy Institute, Tech. Rep., 2012. [Online]. Available: <http://www.techpolicyinstitute.org/files/amkinternetpeeringandtransit.pdf>
- [2] H. Saunders, "IP interconnection: trends and emerging issues," 2012. [Online]. Available: http://berec.europa.eu/files/news/ofcom_ipic.pdf
- [3] V. Bonneau, "Evolution of the CDN market," in *CDN World Summit*, 2010.
- [4] B. Schwarz, "Content delivery networks 3.0," CTOiC White paper, Tech. Rep., 2013. [Online]. Available: http://www.broadpeak.tv/upload/documentation/fichier/73-264-live_webtv_white_paper_b_schwarz_201212.pdf
- [5] P. Crocioni, "Net neutrality in Europe: Desperately seeking a market failure," *Telecommunications Policy*, vol. 35, no. 1, pp. 1 – 11, 2011. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0308596110001461>
- [6] T. Lenard and R. May (Eds.), *Net Neutrality or Net Neutering: Should Broadband Internet Services be Regulated*. Springer, 2006.
- [7] P. Maillé, P. Reichl, and B. Tuffin, "Internet governance and economics of network neutrality," in *Telecom. Eco. - Selected Results of the COST Action IS605 Econtel*, A. Hadjiantonis and B. Stiller, Eds. Springer Verlag, 2012, pp. 108–116.
- [8] T. Wu, "Network neutrality, broadband discrimination," *Journal of Telecom. and High Tech.*, 2003.
- [9] E. Altman, A. Legout, and Y. Xu, "Network non-neutrality debate: An economic analysis," in *Proc. of IFIP Networking*, 2011.
- [10] P. Coucheney, P. Maillé, and B. Tuffin, "Network Neutrality Debate and ISP Inter-Relations: Traffic Exchange, Revenue Sharing, and Disconnection Threat," Inria, Tech. Rep., 2012.
- [11] —, "Impact of Reputation-Sensitive Users and Competition Between ISPs on the Net Neutrality Debate," *IEEE Trans. on Network and Service Management*, 2013.
- [12] R. Frieden, "Internet Packet Sniffing and its Impact on the Network Neutrality Debate and the Balance of Power between Intellectual Property Creators and Consumers," SSRN, Tech. Rep., 2007, <http://is.gd/FG978B>.
- [13] P. Njoroge, A. Ozdaglar, N. Stier-Moses, and G. Weintraub, "Investment in two sided markets and the net neutrality debate," Columbia Univ., Decision, Risk and Operations Working Papers Series, Tech. Rep. DRO-2010-05, 2010.
- [14] F. Sorensen, "The norwegian model for net neutrality," <http://is.gd/HtQuuJ>.
- [15] 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.
- [16] 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.
- [17] 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.
- [18] L. A. Adamic and B. A. Huberman, "Zipf's law and the Internet," *Glottometrics*, vol. 3, pp. 143–150, 2002.

²<http://is.gd/CArKkn>, <http://aws.amazon.com/s3/pricing/>