Impact of Content Delivery Networks on service and content innovation

Patrick Maillé, Bruno Tuffin

► To cite this version:

Patrick Maillé, Bruno Tuffin. Impact of Content Delivery Networks on service and content innovation. NetEcon 2015: the 10th Workshop on the Economics of Networks, Systems and Computation, Jun 2015, Portland, United States. pp.1 - 4. hal-01217929

HAL Id: hal-01217929

<https://hal.science/hal-01217929>

Submitted on 20 Oct 2015

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.

Impact of Content Delivery Networks on service and content innovation

Patrick Maillé
Institut Mines-Telecom,
Telecom Bretagne
patrick.maille@telecom-
bretagne.eu

Bruno Tuffin
Inria
bruno.tuffin@inria.fr

ABSTRACT

Content Delivery Networks (CDNs) are major actors of the current telecommunication ecosystem. Our goal in this paper is to study their impact on other actors of the supply chain, especially on content innovation which is a key concern in the network neutrality debate where CDNs' role seems forgotten. Our findings indicate that vertically integrating a CDN helps Internet Service Providers (ISPs) collect fees from Content Providers (CPs), hence circumventing the interdiction of side payments coming from net-neutrality rules. However, this outcome is socially much better in terms of user quality and innovation fostering than having separate actors providing the access and CDN services: in the latter case double marginalization (both ISP and CDN trying to get some value from the supply chain) leads to suboptimal investments in CDN storage capacities and higher prices for CPs, resulting in reduced innovation.

Categories and Subject Descriptors

C.2.3 [Computer Systems Organization]: Computer-Communication Networks—*Network Operations*

General Terms

Economics, Neutrality, CDNs, Innovation

1. INTRODUCTION

Content Delivery Networks (CDNs) are organizations of large distributed systems of servers positioned at the edges of the Internet network. Thanks to this placement of servers closer to users, content can be delivered with higher availability and better performances. In addition, because the most downloaded content is often chosen to be cached at the edges, the load on the Internet network is reduced.

CDNs have become key player of the Internet, serving an important proportion of content [7]. For example Akamai, one of the leaders, manages a network of around 170,000 servers worldwide and has made a revenue of more than US\$1.5 billion in 2013.

While there is a lot of research activity on CDN caching strategies at servers from a technical point of view [1], the activity on CDN economics is more limited and rather focuses on price optimization strategies for the CDN [4, 5], or on cost-minimizing CDN cache deployment strategies in

Internet AS-level topologies [3]. Very few works exist on the other hand on the economic impact of CDNs on the whole Internet ecosystem. Though, this economic impact is worth studying, especially in the era of the vivid network neutrality debate. Among the focuses of this debate are the payments from distant content providers (CPs) to ISPs (the so-called *side payments*), with service differentiation—if not blocking—when those payments are not settled. This threat has triggered a lot of reactions and discussions about how the network should be ruled, and the FCC took a strong stance in February 2015 to prevent such practices. But surprisingly, CDNs have not been included in the debate, while they ease the pressure on ISPs (one of the claims of ISPs being that they could not afford upgrading the capacity of their network without side payments from CPs whose flows saturate their links) and could on the other hand themselves introduce some differentiation, which may look unfair and against the neutrality principles (although not against recent rules). This issue is discussed in [6]. In [8], the impact of economic strategies of a CDN is investigated and illustrated to indeed lead to differences of treatment among CPs.

One of the main arguments of neutrality proponents is that side payments from ISPs would reduce content innovation because of an increased economic strain. We particularly aim in this paper at investigating the impact of CDN on content innovation (for example, content was considered fixed in [8]). The questions we would like to answer are: 1) What is the impact of CDNs on content innovation with respect to a no-CDN situation? 2) What is the impact on other actors (ISPs, demand from users, quality)? 3) Should ISPs propose their own CDN service and would that replace the side payments without breaking the neutrality rules? 4) What is the best solution for content providers and users, an integrated CDN or an independent one not colluding with ISPs?

The next sections propose a model to answer those questions, where we analyze and compare three cases: no CDN, a vertically integrated CDN (or in other words, a CDN collaborating with the ISP), and a non-cooperative interaction between a CDN and an ISP. The non-cooperative CDN being an intermediary between CPs and the ISP, the case when it has no storage capacity is also considered for comparison purposes.

The rest of the paper is organized as follows. Section 2 presents the model that we will use, Section 3 describes the outcomes of the three different situations and the order of decisions, and Section 4 tries to answer the above questions by analyzing the respective outcomes.

2. MODEL

We consider the model displayed in Figure 1 representing the supply chain and relations between the main actors of the Internet: content/application providers, CDNs, ISPs, and end users.

Figure 1: Representation of relations between users, ISPs, CDN and CPs

Users and content providers are treated as a continuum, meaning that no individual user (or individual piece of content) has an impact on the system (i.e., they are assumed infinitesimal).

Figure 1 also displays the payments: CPs pay the CDN for delivering content (subscription price q per unit of content), users pay ISPs for access (subscription price p per unit of users). In this formulation, the price paid by CPs to CDNs is independent of the content being cached or not; indeed, CPs are not clearly aware of what is cached, they pay the CDN for delivering their content, but the CDN has a positive externality in caching content, since increasing quality and as an indirect consequence the number of customers. In turn, this may also induce more CPs subscribing to the service. Another induced assumption in our model is that ISPs let CDNs freely install servers at the edge of their network, a situation observed in practice for some ISPs with CDNs reducing the load on ISPs' architecture and that content providers only deal with CDNs.

The next subsections characterize the behavior of each of the four classes of actors described in Figure 1.

2.1 Users

We consider a standard discrete choice model, more exactly a *logit* one, for both (infinitesimal) end users and CPs, having the choice between subscribing to the service or not. The choice is then between two options, made of a valuation for each option plus a Gumbel random variable (r.v.). More on this type of user modeling can be found in [2]; it is adapted here to our context.

For an end user, the (random) valuation for subscribing to the service is assumed to be

$$v_u = k \ln(1/p) + \ln Q + \ln n + \kappa_s,$$

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

- κ_s and κ_0 are i.i.d. Gumbel random variables representing variations among users in the service perception;

- the valuation decreases with the service price p , k being a sensitivity parameter that we assume strictly larger than 1;
- it also increases with the quality Q per download and the mass n of content. Indeed, better quality and more choices of content induce an increased level of satisfaction.

The logarithm function that we use comes from the Weber-Fechner law describing the relationship between the magnitude of a physical stimulus and its perceived intensity within the human sensory system, that was also observed to be valid for quality perception of network services [10].

The probability of choosing to subscribe is then (see [2] for details) $\frac{(1/p)^k Q n}{(1/p)^k Q n + (1/p_0)^k}$, that we multiply by D , the fixed base mass of users, to get the mass d of end users actually subscribing:

$$d = D \frac{(1/p)^k Q n}{(1/p)^k Q n + (1/p_0)^k}. \quad (1)$$

2.2 Content/Applications

Like for users, we consider a discrete-choice model for CPs; the valuation for subscribing to the (CDN) service is here $v_c = l \ln(1/q) + \ln d + \kappa_c$, that is compared to the valuation of no subscription: $l \ln(1/q_0) + \kappa_0$. Again, κ_c and κ_0 are independent Gumbel random variables still representing the variations among content providers. The valuation decreases with price q , l being a sensitivity parameter (assumed strictly above 1), and increases with the mass d of end users because more end users induces more business opportunities. There is no direct dependence on quality here (the quality is experienced by users), but an indirect one because a better quality means more customers.

If the base mass of CPs is N , the actual mass of CPs in action is

$$n = N \frac{(1/q)^l d}{(1/q)^l d + (1/q_0)^l}. \quad (2)$$

It is this quantity n that we will consider as an indicator for innovation.

2.3 ISP's utility

We consider a single ISP. Its utility is its revenue:

$$R_i = dp. \quad (3)$$

Operational costs are assumed negligible, and the infrastructure already existing. The goal of the ISP will be to determine the price p maximizing R_i .

2.4 CDN revenue and quality of service

The CDN is also playing on its price q in order to maximize its revenue R_c , defined by

$$R_c = nq - \ell C, \quad (4)$$

where ℓ is the price per unit of server capacity C . Costs of transits to caches are assumed negligible.

The CDN activity induces a better quality perceived by users (as represented in their valuation). The quality of service per download seen by users is a function of capacity C : we assume the quality is 1 for the cached content, and $\delta < 1$ for the uncached one. Let F be the cumulative

distribution function (cdf) of popularity of content and $\bar{F}$ the complementary cdf. Thus $\bar{F}(x)$ gives the proportion of content with popularity (number of downloads per unit of time) at least x . $\bar{F}^{-1}(y)$ is then the number of downloads per unit of time of the y^{th} most popular unit of content. But the most popular contents are cached, representing a proportion $\min(1, C/n)$, the average quality per download is thus

$$Q = 1 \cdot P + \delta \cdot (1 - P) \text{ with } P = \frac{\int_{y=0}^{C/n} \bar{F}^{-1}(y) dy}{\int_0^1 \bar{F}^{-1}(y) dy}, \quad (5)$$

where P is the probability that a given request is served from the cache. In our numerical section, we will take F as an exponential cdf with rate η , so that $F^{-1}(y) = (-y \ln y + y)/\eta$ and $Q = (\delta + (1 - \delta) \frac{C}{n} (1 - \ln \frac{C}{n}))$, which is actually independent of the parameter η .

3. ANALYSIS

Let us now specify the order of decisions. We assume that (i) At the largest time scale, the CDN and ISP fix the prices and caching capacity; (ii) at the smallest time scale the mass of content/applications and users are decided.

This is solved by backward induction: CDN and ISP make their choice anticipating the impact on content and users.

3.1 Determining the masses of content and users for fixed prices

At the smallest time scale, we have the system of equations on n and d (with p , C and q fixed) given by (1) and (2), where Q depends on n as in (5).

This leads to the fixed-point equation in n

$$n = N \frac{(1/q)^l D \frac{(1/p)^k Q n}{(1/p)^k Q n + (1/p_0)^k}}{(1/q)^l D \frac{(1/p)^k Q n}{(1/p)^k Q n + (1/p_0)^k} + (1/q_0)^l} \quad (6)$$

with Q again given in (5). For this equation, the left-hand side is increasing starting from 0. The right-hand side is also starting from 0, so that $(n = 0, d = 0)$ is a (trivial) solution. Indeed, if no user, no interest to produce content, and reciprocally.

If we exclude the solution at $n = 0$, this fixed-point equation seems intractable analytically, but it can easily be solved numerically. Figure 2 draws the right- and left-hand sides of the equation for $k = 2$, $l = 2$, $p_0 = 3$, $q_0 = 1$, $p = 1$, $q = 2$, $D = 10$, $N = 10$, $C = 3$, $\delta = 0.6$ so that $Q = 1$ when $n < C$, $Q = (0.6 + (1 - 0.6)(\frac{-C}{n} \ln \frac{C}{n} + \frac{C}{n}))$ when $n \geq C$ (those values will be used throughout the paper). We get here an intersection point different from $(0, 0)$ at $n = 7.1081$.

Remark that here the equilibrium $(0, 0)$ is not stable in the sense that a small increase in n or d would put the values n or d above the equilibrium, inducing further increases of the values; this can be seen from the derivative of the red (solid) curve being larger than 1 at that equilibrium¹. On the other hand, $(n = 7.1081, d = 9.83193)$ is locally stable: small variations of d and n lead back to that point, because the dashed curve has a derivative smaller than 1.

The values of n and d being determined that way for any combination of p , q and C , the next section considers the determination of the latter values at the larger time scale.

¹computing the derivative of this curve, it can be shown that $(0, 0)$ is a stable equilibrium if $ND < 1$

Figure 2: Determination of the solution for n from (6) when prices and capacity C are fixed.

Figure 3: Best-response functions when $C = 0$

3.2 Selection of prices and caching capacity

At the largest time scale, prices p and q and capacity C can be decided knowing the resulting values of n and d . We consider three situations: 1) $C = 0$ (no caching), q being decided by the ISP; 2) the CDN and ISP cooperate as in the case of a vertical integration of services; 3) the CDN and ISP play a non-cooperative game on their decisions, so that we look for a Nash equilibrium. For that latter case, we will consider both the case without caching from the CDN (who is then just an intermediary) and the case where the CDN decides the caching capacity.

3.2.1 Without CDN

In this case $C = 0$ and q is actually decided by the ISP, that is, the CPs need to connect to the ISP. We look for $\max_{p,q \geq 0} (dp + nq)$.

With the above values for parameters, a maximum revenue $\max_{p,q \geq 0} (dp + nq) = 42.12$ is at $(p = 5.92, q = 1.09)$.

3.2.2 Cooperation between the CDN and the ISP

This happens when the CDN and the ISP are vertically integrated. We look for $\max_{p,q,C \geq 0} (dp + nq - \ell C)$.

Still with the above parameter values, we get a maximum value of 46.43 obtained at $(p = 7.47, q = 1.06, C = 3.31)$.

3.2.3 Pricing game between the CDN and the ISP

First, assume that the CDN is just acting as an access service to CPs, but does not cache anything, i.e., plays on q but with $C = 0$. Figure 3 displays the best-response functions. We clearly see a unique Nash equilibrium (intersection point of the two best-response curves) at $(p = 4.74, q = 2.03)$.

Now assume that the CDN decides its storage capacity C at the same time as its price q , and also at the same time the ISP sets p . We again have a game, but now the CDN has

two decision variables. For our example parameter set, the equilibrium is for $p = 4.92$, $q = 2.04$, $C = 0.126$, obtained after performing successive best-responses from alternating actors (the so-called fictitious play [9], that leads to a stable equilibrium). That equilibrium yields revenue 28.64 to the ISP and 11.77 to the CDN.

4. IMPACT

Figure 4 compares the outcomes for our model when the user sensitivity to price k varies, the other values being those taken so far. We can make the following observations:

Figure 4: Effect of user price sensitivity k

1. Comparing the case without caching with that of an integrated CDN, despite the cost of capacity caching brings an increase of about 10% of revenue. Caching increases the subscription price of users, and decreases that of CPs. Innovation (mass of content) is slightly larger when caching, but it is not the case of the number users despite the increased

quality, due to the larger subscription fee.

2. Comparing the case where the CDN is just an access provider to CPs (with $C = 0$) to the case where this is handled by the ISP itself, quite surprisingly the integrated case leads to much more content (around +40% for our values) than competition. The effect on user demand is smaller, and not always in the same direction.

3. An independent and non-cooperative CDN leads to a much smaller caching capacity and a worse QoE than the integrated case.

4. The best situation in terms of content mass among all scenarios is that of an integrated CDN for our (above) values, but can be the no-CDN situation if users are very sensitive to prices. The difference is quite drastic with respect to an independent CDN.

5. To discuss the difference between caching (and asking subscription from CPs) and side payments, the latter correspond to the “no CDN” scenario, and the former to the “integrated CDN” case. Caching (integrated CDN) increases the provider’s revenue and reduces the load on the network architecture from the ISP point of view. In this sense, it allows to circumvent the side payment proposal. But as we said, this also increases quality and mass of content at a limited cost in terms of users’ subscription, meaning that regulatory bodies and user associations can hardly complain about the situation, unless using other types of arguments.

5. REFERENCES

- [1] H. Atzmon, R. Friedman, and R. Vitenberg. Replacement policies for a distributed object caching service. In *Proc. of CoopIS*, pages 661–674, 2002.
- [2] P. Coucheney, P. Maillé, and B. Tuffin. Network neutrality debate and ISP inter-relations: Traffic exchange, revenue sharing, and disconnection threat. *Netnomics*, 1(3):155–182, 2014.
- [3] S. Hasan, S. Gorinsky, D. C., and R. K. Sitaraman. Trade-offs in optimizing the cache deployments of CDNs. In *Proceedings of IEEE INFOCOM, Toronto, Canada*, pages 460–468, 2014.
- [4] K. Hosanagar, J. Chuang, R. Krishnan, and M. Smith. Service adoption and pricing of content delivery network (CDN) services. *Management Science*, 54(9):1579–1593, 2008.
- [5] 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.
- [6] P. Maillé, G. Simon, and B. Tuffin. Toward a Net Neutrality Debate that Conforms to the 2010s. <https://hal.inria.fr/hal-01127958>, 2015.
- [7] P. Maillé and B. Tuffin. *Telecommunication Network Economics: From Theory to Applications*. Cambridge University Press, 2014.
- [8] P. Maillé, B. Tuffin, G. Simon, and K. Pires. How Neutral is a CDN? An Economic Approach. In *Proc. of IEEE CNSM*, Rio de Janeiro, Brazil, Nov. 2014.
- [9] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behaviour*, 14:124–143, 1996.
- [10] P. Reichl, R. Schatz, and B. Tuffin. Logarithmic laws in service quality perception : Where microeconomics meets psychophysics and quality of experience. *Telecommunication Systems*, 52(2):587–600, 2013.