

# Negotiating premium peering prices: A quantitative model with applications

Costas Courcoubetis, Singapore University of Technology and Design  
Laszlo Gyarmati, Qatar Computing Research Institute  
Nikolaos Laoutaris, Telefonica Research and Development  
Pablo Rodriguez, Telefonica Research and Development  
Kostas Sdrolias, Athens University of Economics and Business

We have developed a novel methodology for deriving bandwidth prices for premium direct peering between Access ISPs (A-ISPs) and Content and Service Providers (CSPs) that want to deliver content and services in premium quality. Our methodology establishes a direct link between service profitability, *e.g.*, from advertising, user- and subscriber-loyalty, interconnection costs, and finally bandwidth price for peering. Unlike existing work in both the networking and economics literature, our resulting computational model, built around Nash bargaining, can be used for deriving quantitative results comparable to actual market prices. We analyze the US market and derive prices for video, that compare favorably with existing prices for transit and paid peering. We also observe that the fair prices returned by the model for high-profit/low-volume services such as search, are orders of magnitude higher than current bandwidth prices. This implies that resolving existing (fierce) interconnection tussles may require per service, instead of wholesale, peering between A-ISPs and CSPs. Our model can be used for deriving initial benchmark prices for such negotiations.

Categories and Subject Descriptors: 500 [Networks]: Network economics

General Terms: Economics, Experimentation, Measurement

Additional Key Words and Phrases: Interconnection economics, pricing, premium service delivery, tussle analysis

## ACM Reference Format:

Costas Courcoubetis, Laszlo Gyarmati, Nikolaos Laoutaris, Pablo Rodriguez, Kostas Sdrolias, 2016. Negotiating premium peering prices: A quantitative model with applications. *ACM Trans. Internet Technol.* V, N, Article A (January YYYY), 21 pages.

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

## 1. INTRODUCTION

There are many “tussles” [Clark et al. 2005] affecting the Internet, ranging across regulation, privacy, network interconnection, and pricing. The economics of peering [Norton 2012], is among the thorniest, but yet, least understood ones. The term peering refers to the interconnection between networks for the purpose of exchanging traffic directly between them. Classic *unpaid peering* played a crucial role in the evolution of

---

Author’s addresses: C. Courcoubetis, Engineering Systems and Design, Singapore University of Technology and Design, Singapore 487372, Singapore; email: [costas@sutd.edu.sg](mailto:costas@sutd.edu.sg); L. Gyarmati, Qatar Foundation, Qatar Computing Research Center, P.O. Box 5825, Doha, Qatar; email: [lgyarmati@qf.org.qa](mailto:lgyarmati@qf.org.qa); N. Laoutaris and P. Rodriguez, Telefonica Research and Development, Plaza de Ernest Lluch i Martin, 5 08019, Barcelona, Spain; email: {nikolaos.laoutaris,pablo.rodriguezrodriguez}@telefonica.com; K. Sdrolias, Department of Informatics, Athens University of Economics and Business, 76 Patission St., Athens 10434, Greece; email: [sdrolias@aueb.gr](mailto:sdrolias@aueb.gr).

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

© YYYY ACM 1533-5399/YYYY/01-ARTA \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

the Internet. It was usually set up between local access ISPs (or regional tier-2 networks) of similar size for the purpose of avoiding charges and longer paths through upstream “transit” providers. The rationale behind such interconnection agreements was that since the involved networks exchanged comparable traffic volumes, there was no issue of who should pay whom.

This has changed recently with the establishment of peering between dissimilar networks, namely Access ISPs (A-ISPs) and Content and Service Providers (CSPs). Whereas classic unpaid peering between A-ISPs was targeting transit cost reduction, the peering between A-ISPs and CSPs is primarily driven by the need to offer improved quality, *i.e.*, low latency and high bandwidth to traffic and revenue generating services such as video, search, online social networks, and gaming. Such a *premium peering* may include a well provisioned interconnection at various Internet Exchange Points (IXPs), CDN hosting of the CSP’s content inside the A-ISP, and even packet level prioritization to avoid congestion at the access part of the network and especially on the last mile. Comcast and Netflix recently signed such a premium agreement<sup>1</sup>. Although there are several (and in many cases controversial) studies attempting to estimate the exact bandwidth price derived by that deal<sup>2</sup>, the involved parties have not yet published any official report regarding the details of their agreement.

**The problem.** Classic unpaid peering was justified on the basis of traffic symmetry, which no longer exists since CSPs inject into A-ISPs several orders of magnitude more traffic (*e.g.*, video) than they receive from them (effectively just HTTP requests) [Dhamdhere and Dovrolis 2010; Labovitz et al. 2010]. This has opened the door to fierce conflicts between A-ISPs and CSPs about who should pay whom and at what rate. The recent disagreement between Free and Google is a good example [The New York Times 2013]. The main battleground for such conflicts is NANOG.<sup>3</sup> Therein peering coordinators have been arguing about payments and have tried to relate them to the question of “*who benefits the most from the premium peering?*”. They have focused mainly on benefits from reduced transit costs [Norton 2012] and have largely left untouched the question of “*who can monetize better the superior traffic delivery quality?*”.

The research literature on networks and economics has developed models of the problem that capture both transit cost savings and monetization of QoE but has derived mostly *qualitative* conclusions. Indeed, as it will be elaborated later (Sect. 6) such models miss many of the important details of the conflict, as prices are derived based on a bilateral basis rather than in a competitive market. Moreover, they are not driven by real data, and hence, cannot derive *quantitative* results, *i.e.*, actual prices for premium peering. Consequently, such models cannot be validated, nor can they be used for deriving actual prices to be used in peering negotiations.

Motivated by the above, we propose a first-of-its-kind framework capable to analyze premium peerings quantitatively. Our model establishes a direct link between service profitability, distribution costs, and ultimately bandwidth price. It suggests that the way forward and away from existing tussles is to perform per service peering. As with other economic models of the Internet, several questions emerge. Can the model capture the essential dynamics of the Internet ecosystem without becoming computationally intractable? Can the basic parameters of the model be extracted from real world publicly available data? Last, can the results derived by the model be validated against reality?

<sup>1</sup><http://arstechnica.com/business/2014/02/netflix-is-paying-comcast-for-direct-connection-to-network-wsj-reports/>, gives emphasis to the improved QoE after the Comcast-Netflix direct interconnection deal.

<sup>2</sup><http://blog.streamingmedia.com/2014/02/heres-comcast-netflix-deal-structured-numbers.html>

<sup>3</sup>[www.nanog.org](http://www.nanog.org)

**Our contributions.** Our main contribution is a methodology leading to a computational model with the properties mentioned above that can be used by A-ISPs and CSPs as a benchmark to derive initial prices for their negotiations around paid peering. We have implemented our model and released a web-site that allows for the analysis of arbitrary markets beyond the one examined in Sect. 5. Specific contributions can be broken down as follows:

- **Modeling:** Unlike prior work on peering that includes transit ISPs [Ma et al. 2010], we focus on the bilateral relationship between A-ISPs and CSPs. In fact, both the A-ISPs and CSPs markets have consolidated in the last decade: the 10 largest telecom operators in the world provide access to more than 2B subscribers [TeleInfo 2012], whereas the 10 largest CSPs amount to more than 63% of the time US users spend online [Nielsen 012a]. Thus, it suffices to establish a small set of agreements for *premium peering* directly between leading A-ISPs and CSPs, guaranteeing that most of the content meets most of the eyeballs through highly provisioned, congestion free, low-delay paths. Furthermore, the content and service discussed here can be replicated in CDNs and clouds close to the end users, from where it is both cheaper and faster to deliver. Combining the “replicability” of content and services with consolidation, seems to indicate that *the bulk of content will flow directly between A-ISPs and CSPs* whereas transit ISPs will handle long-haul communication (email, voip, etc.), long-tail content distribution (smaller A-ISPs, CSPs), and backup paths (for both small and large A-ISPs, CSPs). We therefore focus on the relationship between A-ISPs and CSPs whose synergy is essential for high quality content distribution. In the context of this relationship, we model both costs (e.g., for the extra ports needed for direct peering) and new profits due to improved QoE that translates into increased engagement time of existing customers, and positive (incoming) churn of new customers taken from competitors. Such positive, incoming, churn benefits both the A-ISP and the CSP, whereas the increased engagement time can typically be monetized only by CSPs (since A-ISP offer typically flat contracts for wired broadband [Laoutaris and Rodriguez 2008]). Capturing ISP- and CSP-side churn poses major modeling challenges. One of our main modeling contribution is a churn model across different services and ISPs that is reasonably close to reality while remaining computationally feasible (Sect. 2). We allocate the total surplus by solving a Nash bargaining problem<sup>4</sup>, which outputs fair side-payments that can, in turn, be translated into per bit prices for premium peering.

- **Data driven approach:** We were able to populate most of the parameters of the proposed model with real-world data. A methodological contribution of this work is that populating economic models of the Internet with real data is difficult but not infeasible, at least to a certain level of accuracy that may still produce usable quantitative results. For example we were able to introduce in the model real data for cost of peering, ISP and CSP market shares and profit margins, per service engagement time and sensitivity to QoE, etc. (Sect. 4). Of course, some parameters are more difficult to set than others, e.g., loyalty factors, so we had for some cases to analyse alternative scenarios with respect to such parameters to check the robustness of our results.

- **Practical takeaways:** We run our model for the US market and analyze the results. Our findings show that:

- *CSPs have more ways to monetize improved QoE and thus in most cases they are the ones that pay for premium peering.* Indeed, assuming a similar level of customer loyalty on the A-ISP and CSP markets, the CSP always has the extra benefit of being able to monetize further the premium peering by increasing the engagement time of its existing users and then turning that into additional revenue, through advertising

<sup>4</sup>Nash Bargaining and its solution concept, is widely used in economics as a model for sharing fairly common assets, costs or profits, see for example [Cole et al. 2013].

or paid subscriptions. On the other hand, in view of flat broadband contracts, A-ISPs currently do not have any means to monetize improved QoE.

– *Large ISPs typically receive payments from CSPs for premium peering whereas smaller ones may offer it for free or even pay for it.* The reason is that the only benefit for ISPs relates to churn. An ISP that connects to the content and services of a popular CSP through a premium peering connection can use the superior QoE to attract customers from competing ISPs. This creates new revenues for the ISP but their extent depends on the market size already captured. A large ISP that already holds a sizable part of the market has a smaller pool from which it could attract new customers. If the ISP is small, however, then a premium peering allows it to increase its revenues substantially by attracting a large number of new customers from competitors that do not currently offer premium peering. In that case it is to the best-interest of the small ISP to offer premium peering for free or even pay the CSP for it.

– *Customer loyalty and the ability to translate improved QoE into CSP-side revenues affect the volume and even the directionality of payments.* The previous observations assume a balanced level of loyalty between ISPs and CSPs. If any of the two sides has a substantially higher “lock-in” on its customers then this can have a profound impact on the price of premium peering, and even change the directionality of payments. Lock-in’s on the ISP can be due to brand power or fixed term contracts with disconnection penalties. For CSPs, lock-in can result from the uniqueness of a service or content, the effort already put in learning and customizing it, or from externalities (e.g., in social networks and gaming). Whatever the reason, if users on one side of the market are more loyal than on the opposite side, then they will reduce the revenues from incoming churn and induce the Nash bargaining solution to compensate by transferring payments from the other side of the market. The final outcome depends jointly on the aforementioned imbalance and the ability of the CSP to monetize QoE.

– *Timing matters and first movers have the advantage.* We observe that the post-peering profits for an A-ISP are highly sensitive to what other pairs have established before. The reason is that the first premium peering deals have a larger potential to attract new customers from competitors, whereas late deals are more of a defensive move for retaining new customers rather than bringing new ones. The magnitude of the economic impact is substantial (20% in our case study).

– *Balancing the interconnection economics will require per service peering.* Our case study of the US market indicates that a fair monthly payment from Google to Comcast for premium peering across all services would cost somewhere in the vicinity of \$9.5M and this seems to be stable across much of the parameter space. Breaking down this aggregate payment into per service payments and dividing by the generated bandwidth of each service we were able to derive a \$ value per Gbps per month for each service. It turns out that the fair bandwidth price predicted by our model for video is between \$0.5K–\$1K per Gbps per month which is remarkably close to real transit and paid peering prices from the US market. On the other hand, the Gbps per month \$ value predicted for search, is orders of magnitude higher than current bandwidth prices. This effectively means that current prices are fair for high bandwidth but low revenue per bit services like video, but far from fair for high revenue services such as search. Trying to price fairly amorphous bandwidth that includes such strikingly different constituents might be one of the main reasons for the inability to reach a consensus between A-ISPs and CSPs about what constitutes a fair price. A practical takeaway from our work is that per service peering might help in resolving the tussle.

In the next sections we build our model for pricing bilateral peering contracts. We start first with modeling the dynamics of customers’s churn as a result of a premium peering deal between an A-ISP and a CSP.

## 2. CHURN MODEL

In this section, we propose a formal model that captures both A-ISP and CSP churn and also considers the fact of how important a service is for end users. Before presenting the details of the model, we next describe the essence of our churn model. We assume that customers switch A-ISPs and CSPs in a probabilistic fashion (*i.e.*, flip coins to decide) according to rules we will disclose later on. Customer churn occurs each time the state of the network connectivity is changed. This involves premium peering to be established (or removed) between an A-ISP and a CSP. In such a case, all the services offered by the given CSP to the customers of the A-ISP, change quality from best-effort to premium (or the reverse). We compute the customer transitions<sup>5</sup> in two steps. In the first phase, customers decide to switch A-ISPs because other A-ISPs offer some service, for which they really care, in premium quality. In the second phase, customers within the same A-ISP switch CSPs for the services they are interested in and which are now offered with higher quality by a different CSP. Once these two transitions take place the ecosystem is in equilibrium and will be perturbed again when the network connectivity changes. We proceed by defining precisely our churn model for customers.

### 2.1. Definitions

Let  $ISP$  denote the set of A-ISPs on the market while let  $CSP$  be the set of CSPs. We arbitrarily order the  $K$  different services of the service set  $S$  as  $(1, 2, \dots, K)$  (*i.e.*, service 1 is video, etc.). As a first approximation, to describe the service preferences of a customer, we define its ‘generic’ type  $T$  to be the vector of  $(T_1, T_2, \dots, T_K)$ , where the  $k$ th component  $T_k \in CSP$  defines the CSP providing the  $k$ th service to the customer of the specific type. Let  $\mathcal{T}$  be the set of possible generic types. If a customer is of type  $T$ , we know the CSPs she is using for his services.

We now refine our type definition to include the identity of the most valued service for customers of this generic type (*e.g.*, video is the most valued service of the customer of that generic type). The key idea is to refine the type as little as possible, while capturing the most important aspects of customer differentiation within the same type. The main reason for this is complexity. By doing so we reduce the set of possible customer types, while keeping the ‘first-order’ information for motivating customer transitions between A-ISPs and CSPs. Our new complete definition of a customer type becomes  $(i, (s, T))$  where  $i \in ISP$ ,  $s \in \{1, \dots, K\}$  denotes the most valuable service in this subtype and  $T \in \mathcal{T}$ . The *state* of the type  $(i, (s, T))$  corresponds to the number  $N(i, (s, T))$  of customers of that type at ISP  $i$ .

### 2.2. Transition of customers

We define now the transitions between the states of customer types. Such transitions occur (i) because customers find their preferred service being offered by the same CSP at premium quality at a new A-ISP, or (ii) because at their current A-ISP the quality of their preferred service became available at a higher level by an other CSP. Hence, the transitions of the customers are motivated by *changes of quality only regarding the customers’ preferred service*. In all cases, a customer’s preferred service category remains the same (*e.g.*, if video is the customer’s most valued service, this remains the case during all the transitions), whereas the customer may change its A-ISP or CSP. Therefore, in case of (i), a customer is keeping the same  $(s, T)$  and just switches its A-ISPs. In case (ii), the customer changes type within the same A-ISP, by only changing the CSP of her preferred service. By defining the corresponding transition probabilities, the transitions of the various types are straightforward since we keep

<sup>5</sup>Throughout the paper, we use the terms transition, switch, and churn interchangeably.

track the number of customers of each type per A-ISP. We next describe the customer churn if a premium peering is established.

**Phase 1: churn across A-ISPs.** If A-ISP  $j$  and CSP  $x$  establish a premium peering deal, the customers transition from type  $(i, (s, T))$  to type  $(j, (s, T))$ . Let  $n_1 = N(i, (s, T))$  and  $n_2 = N(j, (s, T))$  denote the number of customers that A-ISP  $i$  and  $j$  have in the given types, respectively. We determine the transition of the customers along the following rules:

- (1) The transition of customers occurs only if  $T_s = x$  and A-ISP  $i$ , had no premium peering with CSP  $x$ .
- (2) The probability of transition is  $\gamma = (1 - \beta(i))h(s)$ , where  $\beta(i)$  is the stickiness (*i.e.*, loyalty) to A-ISP  $i$  and  $h(s)$  is the probability of customers who mainly care about service  $s$  (*e.g.*, video) to switch ISPs because the quality of  $s$  is improved.
- (3) As a result of the state transition, the number of churning customers is  $\Delta = n_1\gamma$ , hence the new states become  $N(i, (s, T)) = n_1 - \Delta$  and  $N(j, (s, T)) = n_2 + \Delta$ , respectively.

**Phase 2: churn across CSPs.** If A-ISP  $i$  and CSP  $x$  establish a premium connection, the customers churn from type  $(i, (s, T))$  to type  $(i, (s, T'))$ , where  $T'_s = x$  and  $T_s \neq y$  hold. If the starting states of these types are  $N(i, (s, T)) = n_1$  and  $N(i, (s, T')) = n_2$ , we compute the volume of churning customers as follows:

- (1) The transition of customers occurs only if CSP  $T_s$  has no existing premium peering agreement with A-ISP  $i$ .
- (2) The probability of the churn is  $\gamma = (1 - \theta(T_s))g(s)$  where  $\theta(T_s)$  is the stickiness to CSP  $T_s$  and  $g(s)$  is the probability for customers that mainly care about service  $s$  (*e.g.*, video) to switch CSPs because the quality of  $s$  improves.
- (3) The number of transitioning customers is  $\Delta = n_1\gamma$ ; the new states become  $N(i, (s, T)) = n_1 - \Delta$  and  $N(i, (s, T')) = n_2 + \Delta$ , respectively.

### 3. COMPUTING THE PRICE OF PREMIUM PEERING

We now introduce an economic framework for computing fair peering prices. Our analysis captures the revenues and the expenditures of an ISP and a CSP both before and after they establish a premium peering agreement, taking also into account customer churn. Initially, the traffic flows between A-ISP  $i$  and the CSP  $x$  throughout a transit ISP (T-ISP), while under a premium peering this traffic will flow directly between them. To analyse the economic gains from premium peering we need to define the revenues and expenditures of the actors before and after the agreement. We use the “ notation to distinguish the various quantities after peering from the respective values before peering is established .

#### 3.1. Profits before peering

For A-ISP  $i$ , let  $u(i)$  denote the profit per customer,  $n(i) = \sum_{s,T} N(i, (s, T))$  the total number customers, and  $n(i)_{\xi=x} = \sum_{s,T|T_\xi=x} N(i, (s, T))$  the number of customers that subscribe for service  $\xi$  to CSP  $x$ .

CSPs have two main sources of income: subscription fees and advertisement-based revenue. If  $R(s)$  = profit per customer subscribing to service  $s$  (assumed uniform across CSPs), the total subscription related profit of CSP  $x$  from the customer base of ISP  $i$  is  $\sum_{\xi \in S} R(\xi)n(i)_{\xi=x}$ .

The advertisement profit of a CSP is a product of  $\tau(s)$ , the engagement time (*e.g.*, minutes) of the customers, *i.e.*, the time the customers spend using service  $s$ ,  $\rho(s)$ , the rate of the clicks<sup>6</sup> of the advertisements in case of service  $s$ , and  $a(s)$ , the profit per click in case of service  $s$ .

Without a peering agreement, A-ISP  $i$  and CSP  $x$  have to pay for transit, being charged the product of the traffic volume and the unit price of the transit service. Let  $\varphi(s)$ ,  $t(i)$ ,  $t(x)$  denote the average traffic rate required per user engagement time of service  $s$  and the unit cost of transit (*e.g.*, per Mbps) for A-ISP  $i$  and CSP  $x$ , respectively. Then, summing up all the traffic originating from the services of CSP  $x$  and terminating at A-ISP  $i$ , the total expenditure of A-ISP  $i$  is:

$$c_t(i) = t(i) \sum_{\xi \in S} \varphi(\xi) \tau(\xi) n(i)_{\xi=x},$$

and similarly for CSP  $x$  the transit cost is

$$c_t(x) = t(x) \sum_{\xi \in S} \varphi(\xi) \tau(\xi) n(i)_{\xi=x}.$$

Based on the above, the profits of A-ISP  $i$  and CSP  $x$  before the peering agreement are respectively  $V_i = n(i)u(i) - c_t(i)$  and  $V_x = \sum_{\xi \in S} (R(\xi) + a(\xi) \rho(\xi) \tau(\xi)) n(i)_{\xi=x} - c_t(x)$ . These expressions contain solely the monetary aspect of the relation between the two actors  $i, x$ , *i.e.*, they do not include the revenues and expenditures related to the other A-ISPs and CSPs present in the market since these are not affected by premium peering.

### 3.2. Profits after peering

If the A-ISP and the CSP agree to engage in a premium peering, their profits change due to three reasons. First, the cost of transmitting the data between the parties changes. They do not incur anymore transit cost but need to pay for the new costs of peering  $c_p(i)$  and  $c_p(x)$  respectively, for A-ISP  $i$  and CSP  $x$ <sup>7</sup>.

Second, customers are becoming more engaged with the services of the CSP if the services are accessible in premium quality. The enhanced quality may affect all the parameters of the CSPs' revenue streams: the profit per customer paying subscription fees ( $\hat{R}(s)$ ), the engagement time of the customers ( $\hat{\tau}(s)$ ), the rate of clicks ( $\hat{\rho}(s)$ ), and the profit per click ( $\hat{a}(s)$ ). Moreover, enhanced quality services may impose additional traffic loads ( $\hat{\varphi}(s)$ ). On the A-ISP side, the profit per customer could also change due to the improved quality ( $\hat{u}(i)$ ).

Third, the improved quality acts as a driving force for the churn of the customers, as we discussed it in details in Section 2 for calculating the new values of  $\hat{N}(i, (s, T))$ . Our post-peering values of the profits for A-ISP  $i$  and CSP  $x$  become  $\hat{V}_i = \hat{n}(i)\hat{u}(i) - c_p(i)$  and  $\hat{V}_x = \sum_{\xi \in S} (\hat{R}(\xi) + \hat{a}(\xi) \hat{\rho}(\xi) \hat{\tau}(\xi)) \hat{n}(i)_{\xi=x} - c_p(x)$  respectively.

<sup>6</sup>or impressions depending on the scheme of the advertisement campaign

<sup>7</sup>The peering costs of the actors are not necessarily identical although both parties handle the same amount of traffic. For example, one of them may face a collocation cost to build into the peering facility (*e.g.*, into a public IXP) where the other party is already present. The cost of the peering depends on several factors like the volume of the traffic, collocation needs, and whether the peering agreement includes a CDN service provided by the A-ISP.

### 3.3. Nash bargaining solution

Based on the profits of the actors before and after the premium peering, we quantify the total monetary benefit of the peering agreement as

$$U = \hat{V}_i - V_i + \hat{V}_x - V_x. \quad (1)$$

If  $U$  is negative, then trivially there is no sense to make such an agreement, and if positive, as result of their cooperation, it should be shared among them. We use the Nash Bargaining solution [Fudenberg and Tirole 1991] as a concept to define fair profit allocations when the possible actions of the players are to establish the peering connection at a specific level of QoS or not peer at all.

We assume that each player has an equal urge to settle the negotiations and that arbitrary monetary transfers are possible (i.e., that after peering the players can pay each other any amount that they have agreed upon). This assumption to allow for an arbitrary payment among the actors simplifies our solution concept tremendously; the solution is to peer if there are positive total profits, and compute the monetary transfers to ensure that each player obtains equal extra profits computed on top of her fall back solution profits (not peer at all, if negotiations break). Formally, we like to find the fair new *total* profits  $z_i$  and  $z_x$  for the two actors by solving

$$\max_{z_i + z_x = U} (z_i - V_i) (z_x - V_x). \quad (2)$$

Once we obtain the net fair profits, we compute the respective payment from CSP to ISP as  $\hat{V}_x - z_x$ , which can be positive or negative. The interpretation is that if under a fair arbitration, CSP  $x$  should end up with a net profit of  $z_x$  that is less than her realised profit  $\hat{V}_x$  after peering (with no side payments to the ISP), he should turn back to the ISP the part of the profit that is above the fair amount. The same results values but with an opposite sign follow if we consider the case of the payment from the ISP to the CSP. In our case the fair payment  $w_x$  is

$$w_x = \frac{1}{2} \left[ (\hat{V}_x - V_x) - (\hat{V}_i - V_i) \right]. \quad (3)$$

The solution definition guarantees that if the peering results to a positive surplus  $U \geq 0$ , none of participants end up with less profit than they would realize before the establishment of the premium peering.

Note that each provider's costs are included in the calculations of the profits. In our definition of the Nash bargaining solution we assumed equal bargaining power, which is a reasonable choice to serve as a benchmark. It is justified by the fact that bargaining power is related to the urge to settle the bargaining process into an agreement, and assumed to be the same for both participants.

## 4. PARAMETERIZATION OF THE MODEL

In this section, we apply the proposed model quantitatively to investigate the derived premium peering prices in the U.S. Internet ecosystem.

Based on real datasets, we present the assumptions we make on the parameters of the model. Initially, we distinguish two types of assumptions: modeling and parameterization assumptions. The modeling assumptions reduce the parameter state size by claiming sameness of some metrics, while for the parameterization assumptions we use real-world datasets in order to assign reasonable values to the model's parameters. Furthermore, we classify the various assumptions into four different sets: Services assumptions (Section 4.1), CSPs assumptions (Section 4.2), A-ISPs assumptions (Section 4.3) and General market assumptions (Section 4.4). In Table I, we summarize the considered assumptions, before proceeding to the justification of our parameterization.



Parameter-Subsection	
Services Parameters	Probability of A-ISP churn-4.1.1
	Probability of CSP churn-4.1.1
	Before peering traffic rates-4.1.2
	Post peering traffic rates-4.1.2
	Before peering engagement times-4.1.2
	Post peering engagement times-4.1.2
CSPs Parameters	Relative Importance-4.1.2
	CSP stickiness assumptions-4.2.1
	Post peering advertising rates-4.2.1
	Market shares-4.2.2
A-ISPs parameters	Estimation of CSP stickiness-4.2.2
	A-ISP stickiness assumptions-4.3.1
	Number of subscribers and annual profits-4.3.2
General parameters	Estimation of A-ISP stickiness-4.3.2
	Market conditions-4.4.1
	Transit and peering costs-4.4.2

Table I: Model parameters.

#### 4.1. Services

We assume that a premium peering deal includes at most, the following four types of services: video, online social network, search, and gaming. This assumption is motivated by the fact that the aforementioned services cover most of the time an average end-user spends online, while capturing more than 70% of the online advertisement-based revenues in 2013 [IAB 2013], under the assumption that Display/Banner and Rich Media ad-formats, fall in any of these types of services.

Furthermore, we distinguish the video services in two sub-types: user-centric video, and commercial based video services. User-centric video is delivered by CSPs whose profits are based on advertisements (e.g., Youtube), while commercial based video services are financially based on monthly subscription fees (e.g., Netflix). Another difference between these two types is their the duration (and hence the end-users' engagement times), since for user-centric content is typically smaller in comparison with commercial based video. Finally, the service quality for user-centric content highly varies, compared to the commercial based that offers typically higher quality video services.

*4.1.1. Modelling assumptions*. Initially we assume that end-users' favorite type of service falls in any of the aforementioned service categories, e.g., can be either search or video or osn or gaming. Furthermore, we attempt to capture the inherent characteristics of each service that differentiate itself from all the other services, in respect with the customer lock-in on both sides of the market. Thus, we assume that the probabilities of churn are not equal for the various services, but increasing in a linear manner. **Probability of A-ISP churn.** The probability of the A-ISP churn ( $h(s)$ ), increases in the following order of the services: search, video, OSN, gaming. In case of gaming and OSN, the related network effects increase the churn probability, while the power of a video service is related to the variety of its available content. Finally, the search engines can be replaced in the easiest way as far as the usability aspect is concerned, as long as the search results are relevant. As explained before, the increase is linear:  $h(search) + 3\mu = h(video) + 2\mu = h(osn) + \mu = h(gaming) = h$ . We use  $h = 0.4$  and  $\mu = 0.1$  as parameters.

**Probability of CSP churn.** Similar trends are considered in case of the probability of the CSP churn ( $g(s)$ ). However, in this case the order of the services reverses, e.g., replacing the search engine has the highest probability. Here we assume again a linear relation between the services:  $g(gaming) + 3\nu = g(osn) + 2\nu = g(video) + \nu = g(search) = g$ . We use  $g = 0.4$  and  $\nu = 0.1$  as parameters.

Service	Relative importance
User-centric Video	15.50%
Online social networks	19.96%
Search	8.82%
Gaming	9.78%
Commercial-based video	45.94%

Table II: Importance of services based on spent time

4.1.2. *Parameterization assumptions*. We use real-world datasets in order to assign reasonable values to the model's parameters. Hence, we apply the forthcoming parameterization to analyze the peering decisions of A- ISPs and CSPs.

**Importance of services.** As Nielsen [Nielsen 2010] points out, in 2010 end-users spent 6.01 minutes for online gaming services per day. Assuming that online gaming hours up 6% year-over-year [VG24/7 2013], it appears that current online gaming engagement times are 7.45 minutes per day.

We use the minutes end-users spend on Netflix, Youtube, Google search and Facebook as proxies for estimating the daily engagement times for commercial-based video, user-centric video, search and online social network services respectively. Based on [Brain 2014], it turns out that on average, end-users spend 11.8 minutes per day on Youtube and 15.2 minutes per day on Facebook. In the case of online movies, it has been estimated that subscribers of Netflix spend around 92 minutes per day on online views [Dave 2014]. As the market coverage of Netflix in the U.S. market is around 38% [Nielsen 2013], we conclude that on average, users spend 34.96 minutes per day on online movies. Finally for online search, we assume that end-users spend around 6.72 minutes of their online time per day (see for example [hashtags.org 2013]).

Given that the time the customers spend online describes accurately the importance of the service categories, we present in Table II, the relative importance of each service. **Before peering traffic rates.** As major online video providers (such as Netflix and Amazon), use Adaptive Bit Rate streaming<sup>8</sup>, we assume that the average download traffic rates for commercial-based video, are on average 22.5 MB/min (Recommended broadband connection speed). Furthermore, based on a press release<sup>9</sup> we assume that on average end-users watch user-centric video services on 480p encoding, which translates to 7.5 MB per minute of engagement time.

For computing the traffic rates of online social network services we use the report of Sandvine [Sandvine 2014]. The median monthly consumption of north american fixed accesses was 19.4 GB, and during the peak hours Facebook accounted for the 1.99% of the total traffic, and hence we assume that the average traffic for this type of service is 0.84 MB/min per user.

The traffic of online gaming is 0.12% of the traffic of video according to [Cisco 2012], and assuming similar traffic volume for search as the gaming one, it is derived that the traffic rates for online gaming are 0.051 MB/min, and 0.054 MB/min for search.

**Post-peering traffic rates.** We assume that the post peering quality of both the user-centric and commercial-based video increases from 480p to 720p, resulting 2.0-time larger traffic rate:  $\hat{\varphi}(\text{user} - \text{centric video}) = 2.0\varphi(\text{user} - \text{centric video})$ , and  $\hat{\varphi}(\text{commercial} - \text{based video}) = 2.0\varphi(\text{commercial} - \text{based video})$ . All the other traffic rates remain the same, as for these services, the premium quality means dominantly reduced latency:  $\hat{\varphi}(s) = \varphi(s)$ ,  $\forall s \in \{\text{OSN}, \text{search}, \text{gaming}\}$ .

<sup>8</sup>[http://scenic.princeton.edu/network20q/wiki/index.php?title=Adaptive\\_Bitrate\\_Streaming](http://scenic.princeton.edu/network20q/wiki/index.php?title=Adaptive_Bitrate_Streaming)

<sup>9</sup><http://sonnati.wordpress.com/2011/09/15/bandwidth-is-running-out-lets-save-the-bandwidth>

CSP	video [comScore 2014a]	OSN [socialfresh.com 2012]	search [comScore 2014b]	gaming*
Google	39.39%	29.6%	69.16%	20%
Microsoft	8.89%	0%	18.85%	20%
Yahoo	13.19%	4.06%	10.56%	20%
Facebook	22.01%	68.95%	0%	20%
AOL	16.52%	0%	1.43%	20%

Table III: Customer distribution among the services of the ad-powered CSPs (\*assumption). Video briefly refers to user-centric video services.

CSP	online movies [Nielsen 2013] *
Netflix	38%
Amazon Prime	13%
Hulu Plus	6%

Table IV: Customer distribution among the services of the subscription-based CSPs

Content provider	Profit, USA (M\$)
Google	294.29
AOL	2.42
Microsoft	37.15
Yahoo	23.39
Facebook	68.02

Table V: Estimated quarterly profits of ad-powered CSPs

## 4.2. CSPs

We consider two types of CSPs: Ad-powered CSPs, which do not offer subscription-fee based services, *i.e.*, their revenues are generated solely via advertisements:  $R(s) = 0$ ,  $\hat{R}(s) = 0$ ,  $\forall s \in S$ , and subscription-based CSPs which charge end-users a fixed fee per month, although they may also earn additional advertisement-based profits.

### 4.2.1. Modeling assumptions .

**CSP stickiness.** We assume that the stickiness of the CSPs is identical:  $\theta(x) = \theta$ ,  $\forall x \in CSP$ .

**Post peering advertising rates.** For any type of CSP, we assume that premium quality does not impact the click through rates and the per click revenues of the services:  $\rho(s) = \hat{\rho}(s)$ ,  $a(s) = \hat{a}(s)$ ,  $\forall s \in S$ .

Finally, a CSP  $x$  is able to capture churning customers only if it had already customers for the given  $\xi$  service:  $\sum_{i,s,T_\xi=x} N(i, (s, T)) > 0$ .

### 4.2.2. Parameterization assumptions .

**Market shares for ad-powered CSPs.** We rely on the reports of comScore on the popularity of video and search services [comScore 2014a] and [comScore 2014b], while for osn services we use a report published by socialfresh.com [socialfresh.com 2012]. More specifically for Google, we sum the popularity of Google+ and Blogger because Blogger is owned by Google. Owing to the lack of specific reports, we assume a uniform distribution of customers in case of games taking into account the fact that all the CSPs have massive gaming services. We distribute the end users among the CSPs for each service category, proportionally to the time they spend using the given service, and summarize our findings in Table III.

**Market shares for subscription-based CSPs.** We use a report from Nielsen published in 2013, to distribute U.S end-users across the considered subscription-based CSPs [Nielsen 2013]. Netflix holds 38% of the market, Amazon Prime 13% and Hulu Plus 6% (see Table IV).

**Profit per engagement time.** Based on the 2014-Q2 financial reports, the net income for Google [Google 2014], Facebook [Facebook 2014], Yahoo [Inc. 2014] and AOL [AOL 2014], were 3422 M\$, 791 M\$, 272 M\$ and 28.2M\$ respectively. For Microsoft its net income was 432 M\$, obtained by its 2013-Q2 financial report [Microsoft 2013], which

Service	USA (\$/min)
Video	0.00092
Online social networks	0.00181
Search	0.01002
Gaming	0.00092

Table VI: Estimated profit per engagement time

Access ISP	Number of subscribers
Comcast	19,025,000
Time Warner	11,306,000
Cox	4,590,000
Charter	3,917,000
Cablevision	3,060,000
others	3,872,800

Table VII: Access ISPs and their market shares in the US, Q3 2012 [Group 2012]

is the most recent financial report that explicitly denotes the economic growth of the Online Services Division.

These numbers correspond to a worldwide advertising audience. Due to the lack of official data corresponding the U.S profits of the considered CSPs, we assume that the volume of their profits is proportional to the number of customers to whom the content is shown. Hence, the profit realized in a country is proportional to the percentage of country’s Internet users, *i.e.*, we assume uniform profit sharing across countries. By combining an Internet Live Stats report [AFP 2011] on the worldwide number of Internet users with the U.S. Internet users in 2012 (around 254,295,536), we estimate the quarterly profits of CSPs in the US from the global profit with a ratio of 0.086 (Table V).

Turning to the subscription-based CSPs, we consider the following providers: Netflix, Amazon Prime and Hulu Plus. Netflix charges its existing clients for \$7.99 per month, while it has recently announced to raise its monthly fees for new streaming subscribers to \$8.99 [variety.com 2014]. Amazon Prime members pay \$99 per year, or \$8.25 per month, while Hulu Plus costs \$7.99 per month. Furthermore as Hulu Plus earns additional profits by showing advertisements along with its content, we assume an average CPM of \$27.61 [ReelSeo 2013], while Netflix and Amazon do not show ads.

We use the aforementioned computed profits of the ad-powered CSPs (Table V) to estimate the value of the  $\rho(s)a(s)$  product, which is the profit per engagement time for service  $s$ . First, based on the report of IAB [IAB 2013], search accounted for 43%, while digital video for 7% of the advertisement spendings. We assume 80–20% distribution of revenues between OSNs and gaming from the 22% of displayed and rich media ads. Second, we assume that all the CSPs derive the same profit per engagement time given a specific service. Hence, we summarize the profits of the considered CSPs and then share the aggregate profit among the services proportional to the percentages reported in the IAB study. Finally, we divide this profit per service value with the aggregate time the users spend in the service category and the total number of subscribers the A-ISPs have. We present the derived  $\rho(s)a(s)$  values in Table VI.

**Estimation of stickiness for ad-powered CSPs.** We estimate the stickiness of the ad-powered CSPs using historical data on their market shares. For user-centric video, we use data of Nielsen [Nielsen 012b] on the top U.S. online video sites (2009–13), while for search, we use the reports of ComScore [comScore 2012] on the U.S. search engine rankings (2008–13).

For each CSP, we compute the ratio of the maximum and the minimum of their market shares. This ratio is a worst-case estimation of the stickiness of the CSP. We use  $\theta = 0.36$ , AOL’s ratio with respect to search, as the lower bound of the CSPs’ loyalty, while  $\bar{\theta} = 0.80$  as upper bound (the ratio of Microsoft in case of video).

### 4.3. A-ISPs

We use publicly available reports on the US Internet access market to capture the number of end users of each A-ISP (Table VII).

#### 4.3.1. Modelling assumptions .

**A-ISP stickiness.** We distribute the different types of customers across the A-ISPs uniformly, while there is no correlation between the usage patterns of the different services. Hence, all the A-ISPs have initially the same type distribution—but with different size of customer base. In any case, the stickiness of the A-ISPs is identical:  $\beta(i) = \beta, \forall i \in ISP$ .

#### 4.3.2. Parameterization assumptions .

**A-ISPs' number of subscribers and annual profits.** We derive the profit per customer for A-ISPs as the difference between the price of a subscription and the cost of provisioning it. A report of ITU gives the average price of broadband access [ITU 2011]: \$20.0 for the US (in 2010). The average provisioning cost of a broadband connection was 46.45% of its price in 2010 [CMT - Comision del Mercado de las Telecomunicaciones 2010]. From the above data we compute the annual profits of the A-ISPs, assuming that end-users pay their monthly Internet access fee, in order to use any of the vast array of services offered in the current Internet. As the services offered by the considered CSPs are obviously only a subset of the online services, we deduct that there is a portion of the total annual A-ISPs' profits generated due to the presence of these CSPs and their menu of services. The more significant is the service, based on the engagement time end-users spent on that service (see section 4.1.2 for details), the higher is the portion of the perceived A-ISPs' profits that is attributable to that service. For example, as osn services are more significant than search (based on Table II), it appears that a larger portion of the total annual A-ISPs' profits is generated due to the presence of osn services in comparison with the profits due to search services.

**Estimation of A-ISPs' stickiness.** We estimate the stickiness of A-ISPs using Leichtman Research Group's historical data on the US broadband market (2009–13) [Group 2012]. For each A-ISP, we compute the ratio of the maximum and the minimum of their market shares based on the historical data. This ratio is a worst-case estimation of the stickiness of the A-ISP. We use  $\underline{\beta} = 0.77$ , Cablevision's ratio, as the lower bound of the A-ISPs' loyalty, while  $\bar{\beta} = 0.95$  as upper bound (the ratio of Time Warner).

### 4.4. General

#### 4.4.1. Modelling assumptions .

**Market conditions.** Initially, we assume that the market lacks premium peering agreements:  $q(i, x) = 0, \forall i \in ISP, \forall x \in CSP$ , and that both A-ISPs and CSPs face the same transit cost:  $t(i) = t(x) = t, \forall i \in ISP, \forall x \in CSP$ .

#### 4.4.2. Parameterization assumptions .

**Peering and transit costs.** Based on [fiercewireless.com 2013], we use 1K \$ per Gbps/month as transit cost. We estimate the cost of peering based on the publicly advertised prices of an Internet eXchange Point (IXP) [ESpanix 2012]. Specifically, the yearly fees are: yearly fee of \$2700, and fees \$4000 (FE), \$4500 (GB) and \$14000 (10G) per port per year. We assume that in case of a premium peering agreement, at least one additional port should be allocated at each parties (*i.e.*, direct peering).

## 5. QUANTITATIVE EVALUATION

In this section we feed data from Sect. 4 into our model, in order to investigate qualitatively the volume and the direction of payments in the U.S. Internet market. Initially, we concentrate on the Comcast-Google pair (largest A-ISP and ad-powered CSP, respectively), assuming that they are the first to introduce premium peering into the U.S. market. Our experiments reveal that three parameters have the most pronounced impact on the prices of premium peering: Namely, (i) the stickiness of the A-ISPs ( $\beta$ ),

(ii) the stickiness of the CSPs ( $\theta$ ), and (iii) the magnitude of the increase of the engagement times ( $a\rho$ ), affect the characteristics of the ecosystem the most. We demonstrate this by implementing several market scenarios, starting from simple ones and gradually introducing more parameters.

### 5.1. The effect of engagement time increase & customers churn

We initially assume that only one service is included in the bilateral premium peering agreement, either search or user-centric video. This permits us to derive interesting conclusions around individual services. Later on we generalize and consider all services in parallel.

**Zero churn on both sides of the market.** If customers are fully loyal to both their A-ISP and CSP, then a premium peering between Comcast and Google is not going to attract any customers from their respective competitors. The only benefit of the premium peering will be an increase of engagement time for their existing customers. While flat monthly payments do not permit Comcast to monetize this increase of engagement time, Google is in position to earn additional profits, and this induces the Nash bargaining solution to transfer some of the benefit over to Comcast. The exact amount depends on the magnitude of increase of the engagement time. We examine two cases, a conservative one, in which the engagement time for video increases by 7.48% and for search by 0.2% (as per the reports of [Krishnan and Sitaraman 2012; ZDNet 2006; Kohavi et al. 2009]), and an optimistic one, in which the engagement time doubles as a result of the establishment of premium peering. Of course the interesting case here is the second, since there is little sense in discussing payments if the premium peering has a small impact on end user QoE and engagement time.

Under the conservative estimation, the monthly payment from Google to Comcast for video is \$11814, while for search is \$2673, reflecting the fact that the assumed increase of engagement time is much higher for user-centric video. In the optimistic case, the monthly payment for video is \$0.18M and for search \$1.17M. Notice that in this case, search yields higher payments than user-centric video. Having equalized the increase of engagement time for both services (doubled), search is carrying a much higher \$ revenue per engagement time increase than video (\$0.01 per extra minute for search versus \$0.00092 per extra minute for user-centric video according to Table VI).

**Customer churn on one side of the market.** Next we demonstrate the effect of customer churn on the direction and the volume of payments.

We start by considering non-zero customer churn on the content side of the market, while keeping the access side churn free. This means that by establishing a premium peering, Google can attract a part of its competitors' customers, whereas Comcast cannot. Having two sources of extra benefit, induces the Nash solution to transfer a larger amount of money from Google to Comcast (for both search and video), compared to the previous case in which the increased engagement time was the only extra benefit. The exact amount depends on the joint effect of the engagement time increase, and loyalty of end-users towards their CSP parameters. For example, when the CSPs' loyalty is equal to 0.4, the monthly payment for user-centric video services from Google to Comcast is \$0.32M, while for CSP's loyalty=0.8, it is down to \$0.23M. In the second case, Google can attract fewer customers from its competitors and thus its benefit, and consequent payment to Comcast is smaller (both amounts are derived under the optimistic case of engagement times increase).

Next we invert the picture and consider non-zero churn on the ISP side, and zero churn on the content market side. Now Comcast is able to gain a portion of its competitors' customers, while Google benefits only through the increase of engagement time of its existing customers. The direction of payment now is not strictly from CSP to ISP but it can be the other way around if the incoming churn of the ISP is more

beneficial than the increase of engagement time for the CSP. For example, under the conservative case for the increase of engagement time mentioned before [Krishnan and Sitaraman 2012; ZDNet 2006; Kohavi et al. 2009], Comcast has to pay Google, for all but the highest values of A-ISP loyalty as depicted in Fig. 1.




Fig. 1: Monthly payments transferred between Comcast and Google, for zero CSP churn, while customer churn may occur on the access side. Comcast has to pay Google, unless the A-ISP's loyalty is very high.

In the optimistic case, however, where the engagement time doubles, Google's profits mask the profits of Comcast from incoming churn and thus the payments are from the former to the latter. For example, the monthly payment for user-centric video from Google to Comcast is \$0.10M for A-ISP loyalty equal to 0.4. For ISP loyalty equal to 0.8 the payment increases to \$0.15M, since in this case Comcast attracts fewer customers from competitors and thus the Nash solution assigns it a bigger chunk of Google's benefits from increased engagement.

**Customer churn on both sides of the market.** We now assume that both Google and Comcast are in position to attract new customers as a result of their premium peering agreement. We still consider a single service, either search or user-centric video.

CSP's loyalty	M\$(search)	M\$(user-centric video)
0.0	1.77	0.34
0.2	1.69	0.29
0.4	1.59	0.25
0.6	1.48	0.20
0.8	1.38	0.15
1.0	1.27	0.11

Table VIII: Monthly payments from Google to Comcast for A-ISP's loyalty equal to 0.5. Customers' loyalty towards their CSP impacts the realized payments.

The direction and volume of payments depend on the interplay of all three parameters (engagement time increase and two loyalty parameters). Fixing A-ISP loyalty to 0.5, Table VIII shows that for an optimistic end-users' engagement times increase, payments are from Google to Comcast but decrease as the loyalty to CSP increases. The reason is that Google attracts customers of its competitors, alongside a significant increase of its advertising revenues. These benefits outweigh the benefits of Comcast from incoming churn, and thus the Nash bargaining imposes payments from Google to Comcast. However, as the number of customers attracted by Google decreases, due to increased CSPs' loyalty, the monthly payments to Comcast decrease as well.

Figure 2 covers the case of pessimistic increase of engagement time. Now the extra advertising profits for Google are small and this reduces the volume of payments. At a certain point of CSP loyalty, the benefits from incoming churn in the ISP and CSP counter-balance each other, leading to small payments. A free-peering agreement make sense in this case.




Fig. 2: Monthly payments for search and user-centric video for low engagement time increase. In most cases Google pays Comcast, although a free-peering agreement could be proved fair for fully CSPs' loyalty.

**The general case: churn in both sides and multiple types of services.** Having examined progressively the effects of engagement time increase and churn on individual services, we next move into the more general and realistic case, in which multiple types service (search, video, osn and gaming) are impacted by the premium peering.

Qualitatively, all our previous observations remain valid also under this more general case. The volume of payments, however, is now higher in absolute value since it is the result of a summation of multiple side-payments for each one of the considered services. Instead of doing a parameter space exploration with respect to loyalty, we will now attempt to derive payments under loyalty levels extracted from the real US market.

We compute the monthly payments between Comcast and Google for the four combinations of the extracted loyalty bounds (Sect. 4). We use  $\underline{\theta} = 0.36$ , as a lower bound of the CSPs' loyalty, and  $\bar{\theta} = 0.80$  as an upper bound. Likewise, we assume  $\underline{\beta} = 0.77$ , and  $\bar{\beta} = 0.95$  as lower and upper A-ISPs' loyalty bounds respectively.

The results of Table IX indicate monthly payments from Google towards Comcast in the order of \$8.92M–\$10.1M under all combinations of loyalty for the optimistic increase of engagement time. In Sect. 5.4 we translate these payments into equivalent bandwidth prices.

An important observation on these payments is that they are actually lower bounds on the amount that the CSP should pay. The reason is that our parameterization in Sect. 4 has considered only interconnection costs for the A-ISP but no further costs from the increased traffic inside the access and metro networks. Using a CDN co-located with customers (*e.g.*, at DSLAMs) can effectively serve the traffic with little impact on the access (beyond the last mile). In that case, both the A-ISP and the CSP would have close to zero IXP costs, whereas the A-ISP would have to incur the costs of provisioning the CDN infrastructure<sup>10</sup>. We have executed this scenario using a CDN cost of \$4K per Gbps per month.<sup>11</sup> Table X shows that the payments from Google to Comcast are somewhat higher for this superior service ( \$9.6M–\$10.8M per month).

	$\underline{\theta}$	$\bar{\theta}$
$\bar{\beta}$	10.1 \$	9.49 \$
$\underline{\beta}$	9.58 \$	8.92 \$

Table IX: Million \$ paid by Google to Comcast (per month), without CDN services

	$\underline{\theta}$	$\bar{\theta}$
$\bar{\beta}$	10.8 \$	10.1 \$
$\underline{\beta}$	10.3 \$	9.6 \$

Table X: Million \$ paid by Google to Comcast (per month), with CDN services

<sup>10</sup>See for example the deployment of Telco CDNs, established mainly by major ISPs , including AT&T, BT and Telefonica.

<sup>11</sup><http://blog.streamingmedia.com/2012/09/cdn-pricing-stable-survey-data-shows-pricing-down-15-this-year.html>






Fig. 3: Volume of monthly fees for Comcast-Google and Cablevision-Google pairs, against multiple values of CSP's loyalty. The large A-ISP always gets the money, while the direction of the money flow between Google and Cablevision strongly depends on the A-ISP's loyalty.

## 5.2. Impact of A-ISP's market share

Next we investigate the impact of the size of an A-ISP's customer base on its paid peering relationship with CSPs. Hence, we compare the previous results involving Comcast and Google with a new set of results between Cablevision and Google. Comcast is the biggest player in the US residential broadband market, while Cablevision is the smallest. In both cases we assume that no previous premium peering agreements exist in the local market.

As Comcast already holds most of the market, a smaller A-ISP has much more to gain from a premium peering relationship with an important CSP. Therefore, whereas before payments went consistently from Google to Comcast, in the case of Cablevision, the ISP can pay the CSP (for low ISP loyalty that permits for maximum incoming churn) and even if it gets paid (under high ISP loyalty that permits less incoming churn), the volume of payment is significantly smaller than the corresponding one for Comcast under similar market conditions. Figure 3 presents concrete \$ values for the above observations for our US case study.

Notice that our results verify real market trends that have large A-ISPs receiving payments from CSPs for direct peering whereas small A-ISPs offering it for free.

Throughout our experiments we have assumed equal transit prices for all the A-ISPs and CSPs, as discussed in Section 4.4.1. One could claim that small A-ISPs tend to buy transit more expensive than larger ones, due to smaller market power. For instance, transit prices may change by a small multiplicative factor between small and large ISPs. Still the difference in the remaining market to capture between a small ISP large one is in orders of magnitude at least, thus masking the effects of buying at bulk discounts on transit costs.

## 5.3. Timing matters

We implement a scenario where Comcast is the last A-ISP to agree with Google, and compare it with the case in which Comcast and Google are the first to establish a premium peering agreement. If Comcast acts in the first place (aggressive peering), then it is able to attract a significant portion of its competitors' customers, while late agreements (conservative peering) is more of a defensive measure for retaining its customers, as its competitors have already exploited the privileges of premium peering deals. If this is the case, Comcast's customer base has been decreased due to the migration of some of its end-users to another, more aggressive, access provider. Hence the post peering profits of Comcast will be smaller in comparison with the case in which it acts proactively, as depicted in Fig. 4.

Nevertheless, Google is able to increase its customers' population by attracting a portion of Comcast's end-users, which used to interact with another CSP, before the premium peering deal. Hence the direction of payments is still from the CSP to the

A-ISP. Interestingly enough, the amount of payments from Google to Comcast for the conservative peering policy is higher in comparison with the aggressive one. This is justified by the fact that as Comcast delays the premium peering deal and hence it is not in position to attract new customers, Google has to pay higher premium peering fees in order to offer the appropriate incentives to Comcast to accept the deal (see Fig. 5).




Fig. 4: Comcast's monthly profits for an aggressive vs a conservative A-ISP's peering policy. The post peering profits are decreasing as Comcast becomes more reluctant to establish the premium peering deal.




Fig. 5: Monthly charges from Comcast to Google for an aggressive vs a conservative A-ISP's peering policy.

#### 5.4. Comparison with real bandwidth prices

Up to now, we have used as metric the monthly payment (in \$) from one party to the other. Next we will compute an equivalent price per Gbps per month for each one of the considered services. The transformation involves just a division of the actual payment reported before, by the extra volume of data transmitted because of the premium peering. The purpose of this study is to see how our payments correlate with actual bandwidth prices in the market.

Based on our most realistic scenario, and assuming that the pair Comcast-Google is the first to establish a premium peering deal, we demonstrate in Table XI the equivalent price per Gbps per month for each service (for a significant engagement times' increase).

CSP's loyalty	user-centric video	search	osn	gaming
0.0	0.96	1355	69	0.22
0.2	0.87	1276	62	0.19
0.4	0.77	1196	54	0.17
0.6	0.67	1116	47	0.14
0.8	0.56	1035	38	0.11
1.0	0.45	954	30	0.08

Table XI: K \$ per Gbps/month transferred between Comcast and Google for A-ISP's loyalty equal to 0.95.

Our first important observation is that *the premium peering prices that we compute for video are reasonably close to real bandwidth price in the market*. Indeed, assuming that paid peering prices are in many cases set to less than the half the price of the corresponding transit prices<sup>12</sup>, and based on current market reports, which claim that transit prices in 2013 have dropped below 1 K \$ per Gbps/month [fiercewireless.com 2013], we observe that our results are in the same order of magnitude with both the current transit and paid peering prices.

<sup>12</sup>See for example: <http://drpeering.net>

The fact that video premium prices match the above real bandwidth prices is very important for the validity of our model, as video dominates traffic on the Internet. Therefore, if the model and its parameterization are sufficiently representative of reality, then the predicted prices for video should correlate with the real prices for amorphous (*i.e.*, independent of service) bandwidth on the market, which seems to be the case according to our results.

The next column of Table XI, which corresponds to search, leads us to a second important observation. *The predicted fair prices for search are several orders of magnitude higher than the corresponding prices for video and the real transit and paid peering prices of the month.* This might appear surprising but it is actually consistent with everything else. Video has high volume and low supporting revenue stream, whereas search has low volume and a high supporting revenue stream. A direct consequence of this observation is that *currently CSPs are paying ISPs only for the bandwidth heavy and low profit service (video), but get to enjoy the delivery of their high profit service (search) almost for free.*

The aforementioned conclusions substantiate our claim that per service peering might provide a transcendent framework on how a fair peering price could be determined. In fact, there is an ongoing work on enabling "application-specific peering", by deploying Software Defined Networking (SDN) at the various Internet eXchange Points (IXPs) [Gupta et al. 2014]. This architecture is aiming to provide each ISP (either A-ISP or CSP) the flexibility to directly express more sophisticated interconnection strategies, such as the per service peering. Hence, our work provides another perspective on the emergence of application-specific peering, by presenting a quantitative analysis of the endogenous asymmetries of the various types of services.

Finally, based on our framework, we have developed *PeeringCalc*, a website that permits users to evaluate our model, on their own case-study (URL: <http://peeringcalc.nes.aueb.gr>).

## 6. RELATED WORK

[Wu et al. 2011] have performed a two-sided analysis using a Stackelberg model to examine the balance of power between an access and a "content-heavy" ISP. [Ma et al. 2010] have examined multi-lateral settlement of inter-connection using the Shapley value to compute a fair split of costs and profits between access ISPs, transit ISPs, and content providers. All these works differ from ours in the level of detail that they capture. Instead of looking at amorphous bandwidth, we model per service traffic costs, increase of engagement time, and most notably, churn. Our model is then parameterized with real data and used to derive quantitative results that can be compared to real bandwidth prices from the market. The aforementioned works derive only qualitative conclusions without attempting to perform any kind of quantitative validation.

[Lodhi et al. 2012] designed a framework to analyze network formation decisions in the Internet. [Dhamdhere and Dovrolis 2010] focused on the flattening of the Internet structure. These works too do not consider the extra revenues from peering nor they attempt to derive fair prices for it.

## 7. CONCLUSIONS AND FUTURE WORK

In this work we have attempted to create a model that connects directly all the ingredients of the tussle: the profits from incoming churn, the profitability of improved QoE, per service characteristics, including loyalty, and the costs of interconnection (peering, CDN, *etc.*). Unlike all previous work in the area our model derives actual bandwidth prices that can be contrasted with real bandwidth prices on the market and thus serve as a benchmark for actual negotiations and/or regulation.

Our model cannot obviously capture every detail of reality but we have put special effort in consulting with industry insiders to make sure that we capture at least the basic dynamics of the tussle. The fact that we output prices that follow intuition qualitatively, and, in some (important) cases match real bandwidth prices from the market, is in our opinion an indication that the model is useful and moving to the right direction. Our basic model of bilateral bargaining takes a snapshot of current markets state, including customer lock-in and market externalities by considering various parameters such as A-ISP and CSP churn probabilities and A-ISP and CSP loyalty. The derived prices are relevant for this particular moment. Of course, our results will not remain relevant once the market evolves. In such case, one needs only to recompute the experiments under the new market conditions.

There are of course many improvements that we would like to pursue as well as additional case studies to be conducted. For example, we would like to encode more detailed cost structures to the model and additional data around them. Our definition of loyalty can obviously benefit from additional experimentation in the area and similarly for the engagement time parameter.

## REFERENCES

- AFP. 2011. Number of Internet users worldwide reaches 2 bln: UN.
- AOL. 2014. Q2 2014 Quarterly Earnings Summary.
- Statistic Brain. 2014. Social Networking Statistics.
- Cisco. 2012. Visual Networking Index, 2011-2016.
- D.D. Clark, J. Wroclawski, K.R. Sollins, and R. Braden. 2005. Tussle in Cyberspace: Defining Tomorrow's Internet. *IEEE/ACM Trans. Networking* 13, 3 (June 2005), 462–475.
- CMT - Comision del Mercado de las Telecomunicaciones. June 2010. Local loop unbundling prices.
- R. Cole, V. Gkatzelis, and G. Goel. 2013. Mechanism Design for Fair Division: Allocating Divisible Items without Payments. In *Proceedings of ACM EC*.
- comScore. 2012. comScore Releases November 2012 U.S.
- comScore. 2014a. comScore Releases March 2014 U.S. Online Video Rankings.
- comScore. 2014b. comScore Releases March 2014 U.S. Search Engine Rankings.
- Pareesh Dave. 2014. Netflix users watching 92 minutes a day, but YouTube remains king.
- A. Dhamdhere and C. Dovrolis. 2010. The internet is flat: Modeling the transition from a hierarchy to a peering mesh. In *Proceedings of CoNEXT'10*.
- Espanix. 2012. Peering prices.
- Facebook. 2014. Q1 2014 Quarterly Earnings Summary.
- fierewireless.com. 2013. Transit Prices Fall Sharply.
- D. Fudenberg and J. Tirole. 1991. *Game Theory*. MIT Press.
- Google. 2014. Q2 2014 Quarterly Earnings Summary.
- Leichtman Research Group. 2012. About 580,000 Add Broadband in the Third Quarter of 2012. (2012).
- A. Gupta, L. Vanbever, M. Shahbaz, S.P. Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, R. Clark, and E. Katz-Bassett. 2014. SDX: A Software Defined Internet Exchange. In *Proc. of ACM SIGCOMM'14*.
- hashtags.org. 2013. Infographic Shows How Much Time You Spend On Social Media.
- IAB. 2013. IAB Internet Advertising Revenue Report, 2013 Full Year Results.
- Yahoo! Inc. 2014. Q2 2014 Quarterly Earnings Summary.
- ITU. 2011. Measuring the Information Society.
- R. Kohavi, R. Longbotham, D. Sommerfield, and R. M. Henne. 2009. Controlled experiments on the web: survey and practical guide. *Data Min. Knowl. Discov.* 18, 1 (Feb. 2009), 140–181. DOI: <http://dx.doi.org/10.1007/s10618-008-0114-1>
- S. S. Krishnan and R. K. Sitaraman. 2012. Video stream quality impacts viewer behavior: inferring causality using quasi-experimental designs. In *Proc. of ACM IMC'12*.
- Craig Labovitz, Scott Iekel-Johnson, Danny McPherson, Jon Oberheide, and Farnam Jahanian. 2010. Internet inter-domain traffic. In *Proc. of ACM SIGCOMM'10*.

- Nikolaos Laoutaris and Pablo Rodriguez. 2008. Good Things Come to Those Who (Can) Wait or How to Handle Delay Tolerant Traffic and Make Peace on the Internet. In *Proc. of ACM HotNets-VII*.
- A. Lodhi, A. Dhamdhere, and C. Dovrolis. 2012. GENESIS: An agent-based model of interdomain network formation, traffic flow and economics. In *Proc. of INFOCOM'12*.
- R. T. B. Ma, D. M. Chiu, J. C. S. Lui, V. Misra, and D. Rubenstein. 2010. Internet Economics: The Use of Shapley Value for ISP Settlement. *IEEE/ACM Trans. Netw.* 18 (2010), 775–787. Issue 3.
- Microsoft. 2013. Q2 2013 Quarterly Earnings Summary.
- Nielsen. 2010. What Americans Do Online: Social Media And Games Dominate Activity.
- Nielsen. 2012a. May 2012 – Top U.S. Web Brands.
- Nielsen. 2012b. September 2012: Top U.S. Online Video Sites and Sports Brands.
- Nielsen. 2013. "Binging" is the new viewing for Over-The -Top Streamers .
- William B. Norton. 2012. *The Internet Peering Playbook: Connecting to the Core of the Internet*. DrPeering Press.
- ReelSeo. 2013. And Hulu's Average CPM in 2013 is...
- Sandvine. 2014. Global Internet Phenomena Report.
- socialfresh.com. 2012. Should Facebook be worried about Google Plus?
- TeleInfo. 2012. Top 10 Telecommunication Service Provider Companies in the world in 2012.
- The New York Times. 2013. Ad Blocking Raises Alarm Among Firms Like Google. (2013).
- variety.com. 2014. Netflix Hikes U.S. Streaming Plan by \$1 per Month for New Members.
- VG24/7. 2013. NPD Online Gaming 2013 report .
- Y. Wu, H. Kim, P.H. Hande, M. Chiang, and D.H.K. Tsang. 2011. Revenue Sharing among ISPs in Two-Sided Markets. In *Proc. IEEE INFOCOM'11*.
- ZDNet. 2006. Google's Marissa Mayer: Speed wins. <http://www.zdnet.com/article/googles-marissa-mayer-speed-wins/>