From advertising profits to bandwidth prices

A quantitative methodology for negotiating premium peering

Laszlo Gyarmati, Nikolaos Laoutaris
Telefonica Research
laszlo@tid.es, nikos@tid.es

Kostas Sdrolias
Athens University of Economics and Business
sdrolias@aueb.gr

Pablo Rodriguez
Telefonica Research
pablorr@tid.es

Costas Courcoubetis
Singapore University of Technology and Design
costas@sutd.edu.sg

1. INTRODUCTION

The economics of peering [18], is among the thorniest debates affecting the Internet, 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 the Internet. It was usually set up between local access ISPs of similar size for the purpose of avoiding charges and longer paths through upstream "transit" providers.

This has changed recently with the establishment of peering between dissimilar networks, namely Access ISPs (A-ISPs) and Content and Service Providers (CSPs). In contrast with classic unpaid peering, the peering between A-ISPs and CSPs is primarily driven by the need to offer premium quality to emerging services such as video, search, online social networks, and gaming. Cablevision and Netflix recently signed such a *premium peering* agreement [13].

Classic unpaid peering [18] was justified on the basis of traffic symmetry, which no longer exists since CSPs inject into A-ISPs several orders of magnitude more traffic than they receive from them [6]. This has opened the door to fierce conflicts between A-ISPs and CSPs about who should pay whom and at what rate. 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 [18], without handling the question of "*who can monetize better the superior traffic delivery quality?*". Many of the proposed models miss 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, and consequently can not be used in peering negotiations.

Motivated by the above, we propose a novel framework capable to analyze premium peering agreements quantitatively. Specific contributions can be broken down as follows:

• **Modeling:** We model both costs and new profits due to improved QoE that translates into increased engagement time of existing customers, and incoming churn of new customers taken from competitors. We allocate the total surplus by solving a Nash bargaining problem which outputs fair side-payments.

• **Data driven approach:** 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.

• **Practical takeways:** Based on our analysis, it turns out that CSPs have more ways to monetize improved QoE and thus in most cases they are the ones that pay for premium peering. In addition to this, large ISPs typically receive payments from CSPs for premium peering, whereas smaller ones may offer it for free or even pay for it. Furthermore, the fair volume of payments (extracted from the Nash Bargaining solution) of a particular A-ISP-CSP pair, is highly sensitive to what other pairs have done before and that early movers may have a clear advantage. Finally by breaking down the aggregate payments into per service payments, and comparing them with actual transit and paid-peering prices we conclude that balancing the interconnection economics will require per service peering. We notice that trying to price fairly application-independent 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.

2. CHURN MODEL

In this section, we establish 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.

Let ISP be the set of A-ISPs on the market, while CSP denotes the set of CSPs. We arbitrarily order the K different services of the service set S as $(1, 2, \ldots, K)$ (*i.e.,* service 1 is video, etc.). We define a customer type as $(i,(s,T))$, where $i \in ISP$, $T \in CSP$ and $s \in \{1, \ldots, K\}$ denotes the customer's most valuable service.

Customers' 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 another CSP. Hence, customers' churn is motivated by *changes of quality only regarding their preferred service*, which in any case remains the same.

Phase 1: churn across A-ISPs. If A-ISP j and CSP x establish a premium peering, 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. 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 to switch ISPs because the quality of s is improved. 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, customers churn from type $(i,(s,T))$ to type $(i,(s,T'))$, where $T' = x$ and $T \neq x$ 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: 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 to switch CSPs because the quality of s improves. 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

Our economic framework for computing fair peering prices, captures the revenues and the expenditures of an ISP and a CSP both before and after the premium peering agreement, taking also into account customer churn.

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 ξ to CSP x.

CSPs may obtain both subscription fees and advertisement-based revenue. If $R(s)$ = profit per customer subscribing to service s, 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 of the customers, $\rho(s)$, the rate of the clicks of the advertisements for 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 per Mbps cost of transit for A-ISP i and CSP x, respectively. Then, 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 $V_x = \sum_{\xi\in S}\left(R(\xi) + a(\xi)\rho(\xi)\tau(\xi)\right)n(i)_{\xi=x} - c_t(x)$. These expressions contain solely the monetary aspect of the relation between the two actors i, x.

3.2 Profits after peering

If the A-ISP and the CSP agree to engage in a premium peering, their profits change due to the following reasons. First both types of providers do not incur anymore transit cost but need to pay for the new costs of peering. Second, customers are becoming more engaged with the services of the CSP if they are accessible in premium quality. Third, the improved quality acts as a driving force for the customers' A-ISP and CSP churn. We use the '^' notation to distinguish the various quantities after peering from the respective values before peering. 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 $\hat{V}_x = \sum_{\xi\in S}\left(\hat{R}(\xi) + \hat{a}(\xi)\hat{\rho}(\xi)\hat{\tau}(\xi)\right)\hat{n}(i)_{\xi=x} - c_p(x)$ respectively.

3.3 Nash bargaining solution

We quantify the total monetary benefit of the peering agreement as

$$U = \hat{V}_i - V_i + \hat{V}_x - V_x\,. \tag{1}$$

We use the Nash Bargaining solution [7] as a concept to define fair profit allocations when the possible actions of the players are to establish the peering connection or not peer at all. Formally, we like to find the fair new *total* profits z_i and z_x for the two actors by solving

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

Once we obtain the net fair profits, we compute the respective payment from CSP to ISP as $\hat{V}_x - z_x$. Finally we calculate the fair payment w_x as

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

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 realise before the establishment of the premium peering.

4. PARAMETERIZATION OF THE MODEL

We consider four types of services: video, online social network (osn), search, and gaming, which cover most of the time end-users spend online, and capture a significant portion of advertisement revenues.

Justified by the highly competitive nature of the Internet access market, all A-ISPs charge the same access fees. Furthermore, customers of the A-ISPs do not pay more for premium connectivity and all A-ISPs have initially the same type distribution—but with different size of customer base. In any case, all A-ISPs have identical stickiness. We compute the annual profits of the A-ISPs based on [4] and [11]. A-ISPs' stickiness is derived using historical data [9].

We assume that $h(s)$, the probability of the A-ISP churn, increases linearly in the following order of the services: $h(search) + 3\mu = h(video) + 2\mu = h(osn) + \mu = h(gaming) = h$. We use $h = 0.4$ and $\mu = 0.1$ as parameters.

We focus on the following CSPs: AOL, Google, Microsoft, Yahoo, and Facebook and use their revenues and profits presented on their Q3 2010 financial reports [1, 8, 12, 10]. In the case of the probability of the CSP churn ($g(s)$), the order of the services reverses. Again there is a linear relation between the services: $g(gaming) + 3\nu = g(osn) + 2\nu = g(video) + \nu = g(search) = g$, where $g = 0.4$ and $\nu = 0.1$.

Using real-world datasets we estimate the per service engagement time of the customers based on [15, 14]. We estimate the traffic rate of the services based on the reports of Sandvine [19] and Cisco [3]. We assume that the quality of video doubles as a result of a premium peering agreement. All the other traffic rates remain the same because for these services, the premium quality means dominantly latency.

For computing customers' distribution for video, osn and search we use the reports of Nielsen and comScore [16, 17, 5]. Owed to the lack of specific reports, we assume a uniform distribution of customers in case of games. Based on [2], we estimate the profit per engagement time for each service. Finally transit prices are taken by a Telegeography report[20].

5. QUANTITATIVE EVALUATION

Feeding data from Sect. 4 into our model, we investigate qualitatively the volume and the direction of payments in the U.S. Internet market. We concentrate on the Comcast-Google pair (largest A-ISP and CSP, respectively), assuming that they are the first to introduce premium peering into the U.S. market. Since there is little sense in discussing payments if the premium peering has a small impact on end user engagement time, we assume that the engagement time doubles as a result of a premium peering deal.

5.1 The effect of engagement time increase & customers churn

We will attempt to derive payments under loyalty levels extracted from the real US market. For each CSP, we are able to 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 $\underline{\theta} = 0.36$, AOL's ratio with respect to search, as the lower bound of the CSPs' loyalty, while $\overline{\theta} = 0.80$ as upper bound (the ratio of Microsoft in case of video). Similarly we obtain bounds on the A-ISPs' loyalty. We consider $\underline{\beta} = 0.77$, Cablevision's ratio, as the lower bound of the A-ISPs' loyalty, while $\overline{\beta} = 0.95$ as upper bound (the ratio of Time Warner).

We compute the monthly payments between Comcast and Google for the above four combinations of the extracted loyalty bounds. The results of Table 1 indicate monthly payments from Google towards Comcast in the order of \$15M–\$18M under all combinations of loyalty.

	$\underline{\theta}$	$\overline{\theta}$
$\overline{\beta}$	17.6 \$	16.6 \$
$\underline{\beta}$	16.4 \$	15.4 \$

Table 1: Million \$ paid by Google to Comcast (per month)

Based on our analysis in Sect. 2 , we are able to extract some key conclusions. By establishing a premium peering agreement, Google has two sources of extra benefit. First, as the user engagement time increases, the CSP is able to increase its advertising revenues. Second a premium delivery of Google's content, triggers customer transition to Google from its competitors, which also leads to additional profits. Comcast on its turn, is not able to monetize the premium peering agreement by its existing customers, due to flat monthly payments. The only source of extra income is the incoming population, obtained by its competitors, after the establishment of the premium peering agreement. Thus, the Nash bargaining solution induces Google to transfer some of the benefit over to Comcast.

The volume of the side payments depends on the interplay of both the ISP and CSP churn. By observing Table 1, it turns out that as customers' loyalty on their CSP increases, the premium peering fees, are falling. If end-users' stickiness to the provider of their favorite type of content is high, Google attracts a smaller portion of its competitors customer base. Hence, its profits are less significant in comparison with a scenario in which end-users are more eager to change the provider of their favorite type of content (small CSP loyalty), and consequently the paid peering payments are lower. Similarly, high A-ISP's loyalty implies low incoming churn to Comcast after its premium peering deal, and hence low post-peering additional profits. If this is the case, Google has to transfer a larger amount of payments to Comcast, in order to boost the A-ISP's incentives to accept the premium peering interconnection.

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 relationships with CSPs. Therefore, we compare the premium peering payments between 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, 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 payments is significantly smaller than the corresponding one for Comcast (see Figure 1).

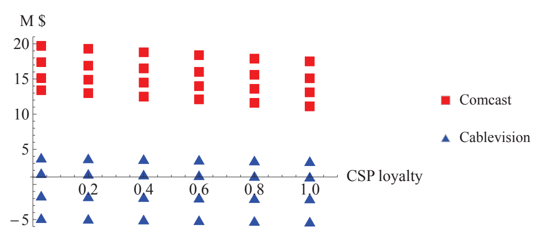

Figure 1: Volume of monthly fees for Comcast-Google and Cablevision-Google pairs, against multiple values of CSP's loyalty. The direction of payments strongly depends on the A-ISPs' market share.

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.

5.3 Timing matters

In order to investigate in what way the timing of the premium peering agreement affects the derived payments, we compare a scenario in which Comcast is the first to introduce such agreements in the U.S market, with a case where it waits for its competitors to first act. 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. Thus, the generated profits and the derived fair side-payments, will be significantly lower in comparison with the aggressive peering case. 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 (see Fig. 2).

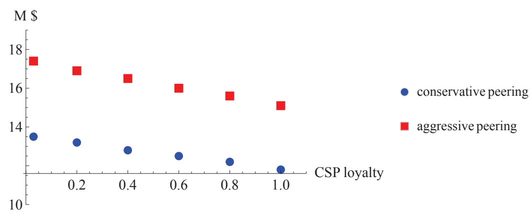

Figure 2: Monthly charges from Comcast to Google for an aggressive vs a conservative A-ISP's peering policy. The volume of payments decreases as Comcast becomes more reluctant to establish a premium peering deal.

5.4 Comparison with real bandwidth prices

Next we will compute an equivalent price per Gbps per month for each one of the considered services, in order to see how our payments correlate with actual bandwidth prices in the market. Assuming that the pair Comcast-Google is the first to establish a premium peering deal, we demonstrate in Table 2 the equivalent price per Gbps per month for each service.

CSP's loyalty	video	search	osn	gaming
0.0	0.65	880	155	387
0.2	0.60	828	143	341
0.4	0.55	776	131	293
0.6	0.51	724	119	243
0.8	0.46	671	107	192
1.0	0.40	618	95	138

Table 2: K $ per Gbps/month transferred between Comcast and Google for the upper bound of 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 that the half the price of the corresponding transit prices [1], and based on current market reports, which claim that transit prices in 2013 have dropped below 1 K $ per Gbps/month [2], we observe that our results are in the same order of magnitude with both the current transit and paid peering prices.

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 parametrization 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 2, 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 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. Finally, we have developed *PeeringCalc*, a web-site that permits users to evaluate our model, on their own case-study. (URL: http://195.251.252.144/PeeringCalc)

[1]See for example: http://drpeering.net

[2]http://www.fiercewireless.com/europe/story/ip-transit-prices-fall-sharply

6. REFERENCES

[1] AOL. Third Quarter 2010 Earnings Summary.

[2] I. A. Bureau. Iab internet advertising revenue report, 2012.

[3] Cisco. Visual Networking Index, 2011-2016, 2012.

[4] CMT - Comision del Mercado de las Telecomunicaciones. Local loop unbundling prices, June 2010.

[5] comScore. comScore Releases November 2012 U.S. Search Engine Rankings. 2012.

[6] A. Dhamdhere and C. Dovrolis. The internet is flat: Modeling the transition from a hierarchy to a peering mesh. In *Proceedings of CoNEXT'10*.

[7] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.

[8] Google. Q3 2010 Quarterly Earnings Summary.

[9] L. R. Group. About 580,000 Add Broadband in the Third Quarter of 2012.

[10] Y. Inc. Q3'10 Financial Highlights.

[11] ITU. Measuring the Information Society. 2011.

[12] Microsoft. First Quarter Fiscal Year 2011 Results.

[13] Multichannel News. CES: Cablevision Cuts CDN Deal With Netflix, 2013.

[14] Nielsen. What Americans Do Online: Social Media And Games Dominate Activity, 2010.

[15] Nielsen. August 2012: Top U.S. Web Brands and Technology News Websites, 2012.

[16] Nielsen. September 2012: Top U.S. Online Video Sites and Sports Brands, 2012.

[17] Nielsen. Social Media Report 2012: Social Media Comes of Age, 2012.

[18] W. B. Norton. *The Internet Peering Playbook: Connecting to the Core of the Internet*. DrPeering Press, 2012.

[19] Sandvine. Global Internet Phenomena Report, 2H 2012.

[20] TeleGeography. IP Transit Prices Continue to Decline, Geographic Differences Remain, 2011.