Paid Peering: Pricing and Adoption Incentives

Costas Courcoubetis, Kostas Sdrolias, and Richard Weber

Abstract: **Large access providers (ISPs) are seeking for new types of business agreements and pricing models to manage network costs and monetize better the provision of last-mile services. A typical paradigm of such new pricing norms is the proliferation of paid peering deals between ISPs and content providers (CPs), while on top of this, some ISPs are already experimenting with usage-based tariffs, usually through data-plans, instead of the typical fixed-based charging. In this work we define as common platform, the infrastructure in which a single ISP transacts with several CPs through peering agreements. In this context, we examine whether, and under which market conditions, the profitability of the involved stakeholders improves when the establishment of this platform is accompanied by a monetary compensation from the CPs to the ISP (paid peering), v.s. a scenario where their deal is a typical settlement-free one. In both cases, we assume that the ISP implements a usage-based access pricing scheme, implying that end-users will pay more for higher transaction rates with the CPs. Our framework captures some of the most important details of the current market, such as the various business models adopted by the CPs, the end-users' evaluation towards the ISP's and CPs' level of investments and the traffic rates per transaction for the offered services. By analysing the equilibrium derived by a leader-follower game, it turns out (among other practical takeaways) that whether or not the profitability of a CP improves, it highly depends on whether its business model is to sell content, or if it obtains its revenue from advertisements. Finally, we extract that consumer surplus is considerably higher under paid peering, which in turn implies improved levels of social welfare [1].**

Index Terms: **Interconnection economics, network neutrality, tussle analysis.**

I. INTRODUCTION

THE debate around the economics of interconnection is currently on the rise, following an era in which the most common types of interconnection between two networks have been either *transit* — where one party pays the other for access to the Internet — or *settlement-free* peering, implying the direct interconnection between two administratively separate entities for the purpose of directly exchanging traffic between them and hence bypassing several, and probably congested, transit paths [2]. Throughout our paper we use the term *free* peering as an abbreviation for the settlement-free scenario, when this is the case between ISPs or between ISPs and CPs.

Typically, free peering deals have been established between entities who are symmetric in terms of the exchanged traffic rates. When the involved players exchange similar volumes for traffic, there is no need for monetary payments between them, and their free peering agreements mainly aim at reducing transit costs.

A. Motivation

Recently some ISPs, including Comcast, a dominant access provider in the U.S., have offered *paid* peering, which is a variation of free peering, but with a monetary compensation from one party to the other. ISPs argue that free peering deals are justified when the involved parties exchange comparable traffic volumes, but this is not the case when a CP delivers several orders of magnitudes more traffic (as with streaming services) than it receives (typically from simple HTTP requests). Additionally, ISPs claim that a possible direct interconnection agreement between an ISP and a CP not only requires the acquisition of some additional interconnection ports, but also other upgrades, especially on the last-mile, to support the increased traffic generated by the end-users due to the improved quality of experience (QoE) [3].

Disliking the cost implications of the above argument, several CPs, including Netflix, which accounts for the 34.21% of the total downstream traffic in the U.S. [4], have requested an aggressive regulatory intervention by the FCC to prevent ISPs from adopting an extensive paid peering policy [3]. Instead of *paid* peering deals, Netflix promotes "Netflix open connect initiative" [5], which aims to localize substantial amounts of traffic inside the access network, on the condition that the ISP agrees to a free (unpaid) peering deal. Netflix recently announced that it delivers its traffic to the 99% of ISPs without transfer payments, and accused Comcast of taking advantage of its monopoly power over their large customer-base to impose peering payments [6]. Nevertheless, and after the observed degradation of the QoE of Netflix content in Comcast's access network [7], the dispute between the two parties ended-up with a paid peering deal, in which Netflix has agreed to pay Comcast, hoping it will provide increased QoE to the end-users [8].

Overall one can argue that the direct interconnection between the ISP and the CPs will improve the end-users' QoE, by having bypassed the various transit paths that were traversed before the agreement. What is referred to 'fast lanes' by many researchers, can be the result of the improved access infrastructure by the ISP after CPs choose a direct interconnection that leads to increased ISP revenues. In this paper we use the term 'HQ (high quality) infrastructure' to refer to the upgraded ISP infrastructure which results when paid peering is the dominant strategy in the CP – access ISP market, in contrast to the infrastructure when it is not.

Recently, in the matter of Protecting and Promoting the Open

Manuscript received April 15, 2016.

C. Courcoubetis is with the Department of Engineering Systems and Design, Singapore University of Technology and Design, Singapore, email: costas@sutd.edu.sg.

K. Sdrolias (correpsonding author) is with the Department of Computer Science, Athens University of Economics and Business, Greece, email: sdrolias@aueb.gr.

R. Weber is with the Statistical Laboratory, Faculty of Mathematics, University of Cambridge, UK, email: rw1@cam.ac.uk.

Digital objective identifier: 10.1109/JCN.2016.000131

[1] A short version of the present work has been presented in IEEE ICC'16 [1]

Internet, the FCC has issued a "notice of proposed rulemaking" [9], which seems to prohibit ISPs from charging CPs, for premium access to their end-users. This ambiguous rule, forced many network experts[2] to raise a warning flag that eventually paid peering lies into the soft underbelly of the long-standing network neutrality debate. In fact, the network neutrality principle imposes that ISPs should treat Internet data the same, without discriminating or charging differentially by user, content or mode of communication. Thereafter, they argue that subsidized direct interconnection deals will lead to the de facto establishment of a multi-tier access network, where those who afford to pay the ISPs, will obtain an enhanced access to the eyeballs; and this fact will actually improve further their competitive advantage over smaller players or start-ups.

B. Our Approach

Motivated by the aforementioned 'tussle' between ISPs and CPs, we design a framework to analyse and compare the effects of various parameters on the profits of ISPs and CPs under two extreme situations: when all parties use paid peering, and when no such revenue sharing takes place[3].

Our model assumes that there is a neutral platform, where the ISP makes money only from access usage-based fees from its access customers, or by deploying a non-neutral platform with higher quality delivery of services, in which CPs pay. Thereafter, the two examined modes of pricing, i.e., free and paid, could be alternatively denoted as neutral and non-neutral, respectively.

We consider a single type of access ISP, and CPs of two types: Ad-powered CPs (aCPs), whose profits are generated solely by advertisements, and subscription-based CPs (eCPs), which charge end-users for content transactions. The eCP class includes CPs that provide content to customers by charging fixed monthly fees, by assuming some equivalent average price per content transaction and a price-elastic demand[4].

In Fig. 1, we illustrate the entities of our market and the direction of the money flows. The dotted lines indicate payments corresponding only to the paid peering regime.

It is the demand for user transactions that creates traffic for the ISP, and for both eCPs and aCPs this demand increases with the popularity of the content and sizes of infrastructure investment made by the CPs and the ISP[5]. Crucially, in the case of eCPs (but not aCPs), it also decreases in p, the price charged to end customers for content.

In the paid peering case, the ISP sets charges to the CPs to maximize its profit, assuming full information on the CP's business model, and that all stakeholders choose their corresponding optimal infrastructure investment levels. By intuition alone, it is not clear what the net effects of paid peering might be. On one hand, it may create incentives for the ISP to invest more in infrastructure, overall benefiting the CPs. On the other hand, it may just push up the profits of the ISP, and even worse, discourage CPs to invest. Nevertheless, sufficiently large investments by ISPs might also lead to an increase of the CPs' revenues as end-users become more engaged with their services, due to the increased QoE.

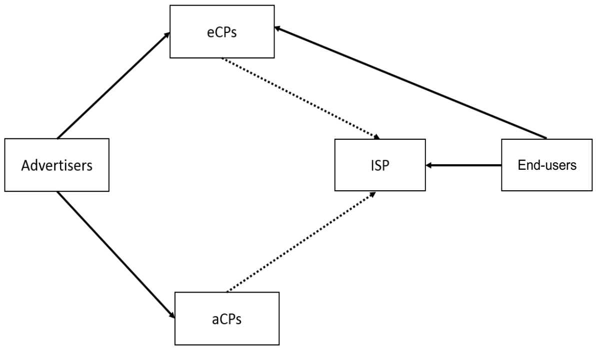

Fig. 1. The figure presents the actors of our model and the direction of the money flows. The solid lines imply payments taking place in both the free and paid peering regime. More precisely, the ISP receives access usage-based fees by the end-users, which also pay a fee per transaction to the eCPs. On top of that, we further assume that eCPs are also receiving ad-based revenues, which is the only source of income for aCPs. The dotted lines from the CPs to the ISP indicate the peering payments, taking place only under the paid peering regime. Under free peering, the ISP's revenues are derived solely from its access customers.

Hence, the overall net benefit for the various stakeholders is unclear and cannot easily be deduced by simple arguments. Our aim is to investigate precisely these aspects of paid peering and, by taking account of its influence on investments and pricing strategies of eCPs, deduce the impact on stakeholder profitability and social welfare.

Of course our model is simple, as most tractable economic models must be. It is not meant to capture all details of the actual market. But we believe that it shows and justifies some important trends that will be seen if paid peering becomes the dominant in the CP–ISP market. Our goal is not to advocate paid peering but to explain its potential benefits and shortcomings to players in the Internet ecosystem, and to determine the key parameters which affect economic performance.

C. Main Research Questions

i. *Which parameters are crucial to determine optimal peering prices?* Given the algebraic complexity of the model, it is rather surprising that we obtain simple intuitive analytic formulas for the optimal paid peering volume prices, which involve for each type of CP some measure of its profitability, the traffic volume per transaction, the impact of its infrastructure investments, and for the ISP the price he can charge end users for access services. We also discover that eCPs will add a markup on their original content price to compensate the effects of being charged by the ISP. This sends a signal to the ISP that he should not overcharge an eCP since this will directly reduce the revenue of the ISP due to the drop of demand for CP's content and hence of chargeable edge ISP traffic. More specifically, we conclude that the peering payments from the eCPs to the ISP do not depend on their (if any) advertisement-based profits, and are

---

[2] `http://www.interstream.com/blog/drpeering-3`

[3] This 'greenfield' model has been traditionally used by regulators when developing bottom-up models where the network is designed from scratch.

[4] Based on the business model of Hulu Plus (a VoD provider in the U.S.), we assume that eCPs may also earn additional ad-based revenues.

[5] Examples of infrastructure investments for CPs include HD encoding, HD streaming, or using CDNs and building caches closer to the end-users, while for ISPs these are deep caching, or FTTx, etc. On top of these expenditures, we can always add the direct peering costs at the Internet exchange points (IXPs).

computed based on the subscription revenues earned from the end-users. This observation is significant since some of the current most powerful eCPs do not earn additional advertisement profits. It is a direct consequence from the significant effects peering payments have on eCP prices to their customers, and hence on the demand for user transactions that translates to ISP traffic. This is not the case for aCPs where peering prices influence very indirectly traffic demand through their effect on increasing/reducing aCP infrastructure investments (also the case for eCPs).

ii. *How is the volume per transaction influencing the paid peering charges?* We find that the greater the volume per content transaction for a CP, the less are the per-unit payments to the ISP. This implies that providers of light volume applications such as search, are in a way, forced to pay greater fees per unit of traffic and hence to contribute more in proportion to their traffic in the financing of the common ISP infrastructure, compared to video streaming applications. This finding provides some elements of fairness, as a possible charging scheme that would not take into account the inherent characteristics of each type of service could be proved unaffordable for volume-intensive CPs, such as VoD providers. This also confirms the findings presented in [10]. In this work, based on a quantitative model of fair profit sharing between CPs and access ISPs fed with actual market data, it is shown that the predicted fair payments for heavy bandwidth applications such as video were orders of magnitude less than for light ones such as search.

iii. *What is the impact of paid peering on the profits of the CPs of different types?*
We find that when the common platform includes only one type of CPs, then eCPs are consistently better off under paid peering compared with aCPs. An explanation is that the ISP has to take into account that eCPs are able to transfer the peering payments to end-users as a markup due to the better infrastructure quality (see (7)), and both ISP and eCP gain if the resulting demand for transactions increases. Hence, the incentives of the eCP and the ISP are more aligned, resulting in eCPs gaining higher profits under paid peering compared to the aCPs which may observe a significant reduction of their profits. If the ISP overcharges the eCPs, this extra charge will lead to a reduction of the end-users' usage levels and consequently diminish the ISP revenues.
On the contrary, as mentioned above, aCPs can be charged by the ISP with no direct reduction of user traffic (and hence of ISP profits), besides the indirect effect that such charges have on decreasing infrastructure investment incentives of aCPs (which influences negatively demand).
Nevertheless, we show that the introduction of eCPs in the HQ platform, eventually mitigates the aCPs' economic loss, due to the imposition of peering fees.

iv. *Is the social welfare improved under paid peering?* Throughout our simulations, we have identified that the ISP's level of investments will be always higher under the paid peering scenario. This directly affects the end-users' rate of transaction with the various CPs, and hence the consumer surplus will be improved in the HQ infrastructure, especially as the access price per transaction increases. Moreover, we deduce that in general, the social welfare is improved under paid peering.

The paper is structured as follows. In Section II we discuss related work and in Section III we define our economic model and its key parameters. This is followed, in Section IV, by an indicative study of how the profits of the players are derived in the free and the paid peering regimes. In Section V we present the parameterization of our model and our key results, while in Section VI we discuss some assumptions of our framework and possible extensions, concluding in Section VII.

II. RELATED WORK

In their seminal work Clark *et al.* [11] presented some design techniques taking into account contemporary tussles on the Internet. Houidi and Pouyllau [12] extended this work by focusing on tussles specifically between ISPs and CPs.

Revenues sharing issues between Service Providers have been examined by He and Walrand [13], while Shrimali *et al.* [14] looked at a related scenario by providing a game-theoretic analysis of paid peering agreements. Courcoubetis and Weber [15] have proposed some improved sharing policies for common infrastructures. Li and Huang [16] have studied the effects when the ISP price-discriminates among the various CPs, in the context of complete and incomplete information.

Norton [2] and the discussion in NANOG[6] summarize technical and practical views around interconnection economics. Clark *et al.* [17] have investigated the complexities of the commercialized Internet and provided insights on interconnection bargaining processes, while Lodhi *et al.* [18] have designed a framework to analyze network formation decisions in the Internet. Ma *et al.* [19] have provided a framework based on the Shapley value to compute a fair split of profits and costs between ISPs, Transit Providers, and CPs.

Platform competition issues in the context of two-sided markets are have been presented in [20], while Musacchio *et al.* [21] and Njoroge *et al.* [22] have studied a related investment incentives problem between ISPs and CPs using a two-sided market model. Wu *et al.* [23] have further examined bargaining power issues between an ISP and a CP, using a Stackelberg model, again in the context of two-sided markets. Altman *et al.* [24] have investigated the implications of usage-based pricing in non-neutral networks assuming both cooperative and non-cooperative case-studies.

Our economic model is one-sided and borrows ideas from the framework, presented in [21], regarding the form of the demand function for content in (1) and (2). Our model extends existing ones and assumes:

i. Different business models for the CPs where content transactions may be charged to end-users or be paid by advertisers.
ii. Different types of content with different application-related parameters, such as the traffic volume per transaction.
iii. Prices for content that are determined endogenously in the market assuming CPs provide independent services.
iv. Fixed broadband access prices, determined exogenously by competition or by the history of the market).

---

[6] `www.nanog.org`

III. THE ECONOMIC MODEL

A. *Traffic Rates*

In our model a CP is a provider of content goods, and the provision of a single unit of such a good corresponds to a single content transaction between an end-user and the CP. We assume customers are homogeneous in respect of their demand for content. We use subscripts *a* and *e* when referring to parameters of aCPs and eCPs, respectively.

The key economic quantity upon which we focus is the demand for transactions. We start with the case of an aCP with no end-user payments, where demand in terms of the expected rate of end-users' transactions with the aCP is

$$D_a = d_a c_a^u t^w e^{-sb_a/\theta} . \tag{1}$$

This is a Cobb-Douglas demand function, similar to [21]. Functions of this type are widely used to relate the amounts of multiple types of input to the amount of output they produce. In (1) we express the link between the level of investments t and c_a, of the ISP and the aCP respectively, and the output transaction rate, D_a. In particular:

- d_a denotes the popularity or attractiveness of aCP's content.
- s is the access price charged to end-users per unit of traffic, and θ models the price elasticity of the user demand for access traffic, when charged with s per unit. We assume s is exogenously defined, since competition effectively makes this price fixed, as the current situation in the Internet suggests. Although today most access pricing uses flat rates, we expect users to pay more due to tiered pricing;this extra effect is modelled by s. Taking s near zero could be a proxy for flat rate. We note that using usage based pricing is being actually discussed in practice, see [25], while Comcast has already expressed its ambition to move towards this type of billing, and currently trials a number of such pricing policies strategies [26]. In Section VI, we provide a more thorough discussion on the emergence of usage-based pricing.
- b_a denotes the traffic volume per transaction. We capture the fact that demand for transactions will decrease as the price sb_a per transaction increases. The decrease is less for greater values of θ;
- $w, u \geq 0$ and $w + u < 1$, so that demand is a concave increasing function of both c and t with decreasing returns to scale[7]. Concavity models saturation effects with increasing infrastructure investments and $w + u < 1$ ensures concavity over the combined infrastructure levels (e.g., even if $c = t$).

Turning to eCPs we capture the effect of a price p per transaction being charged by the eCP to the end-users, by modifying (1) to give

$$D_e = d_e c_e^u t^w e^{-sb_e/\theta} e^{-p/\theta_e} . \tag{2}$$

Now θ_e represents the end-user' price sensitivity towards price p, b_e is the eCP's traffic volume per transaction, while c_e denotes each eCP's level of investments. In (2) we see that price p acts to reduce the end-users' rate of transactions with the eCPs, while greater θ_e implies more price-inelastic users, with respect to p.

Table 1. Summary of parameters.

Parameter	Symbol
Significance of CPs' investments	u
Significance of ISP's investments	w
ISP's marginal cost	k
CPs' marginal cost	z
Access fee (per unit of transaction)	s
eCPs' transaction price	p
eCPs population	n_e
aCPs population	n_a
Popularity of eCP's content	d_e
Popularity of aCP's content	d_a
eCP advertising rates	γ_e
aCP advertising rates	γ_a
eCP traffic rates	b_e
aCP traffic rates	b_a
ISP price sensitivity	θ
eCP price sensitivity	θ_e

Our overall objective is to investigate how the network quality and prices affect the adoption incentives and the decisions of the various CPs. Hence, we do not introduce competition effects in the demand functions, assuming that each CP has its own customer base. Nevertheless, our model encompasses that CPs with more interesting content (e.g., higher d_e or d_a), receive higher click rates.

Converting the aforementioned transaction rates to equivalent traffic rates ρ_a, ρ_e we obtain:

$$\rho_a = b_a D_a, \quad \rho_e = b_e D_e. \tag{3}$$

Based on the above traffic rates, we next present derived profit functions of the multiple economic entities.

B. *Profit Functions*

For simplicity we assume that all CPs of the same type have similar parameters and that the ISP is able to discriminate among the types of the CPs. Hence all eCPs (aCPs) will have to pay under paid peering a price q_e (q_a), per unit of traffic volume. We allow q_e, q_a and p to be negative to also model subsidies.

The revenue the ISP makes from CPs is due to the traffic generated by the transactions of the CPs plus the peering charges in the case of paid peering. Its costs are due to investment quantity t. Hence the ISP profit is

$$\pi_{isp} = n_e\rho_e(s + q_e) + n_a\rho_a(s + q_a) - kt, \tag{4}$$

where n_e and n_a denote the number of eCPs and aCPs respectively, and k the marginal investment cost of the ISP. In the case of free peering we set $q_e = q_a = 0$.

The profits of each aCP and eCP are given respectively by

$$\pi_a = \rho_a(\gamma_a/b_a - q_a) - zc_a \tag{5}$$

and

$$\pi_e = \rho_e(\gamma_e/b_e + p/b_e - q_e) - zc_e, \tag{6}$$

where

- z denotes the marginal cost of investment for a CP (assumed equal for eCPs and aCPs);
- γ models the revenue per transaction of each CP due to advertisements. For example $\gamma = \mathrm{CTR}\times$'*ad-revenue-per-click*',

---

[7] The exponents u, w model the responsiveness of end-user demand to infrastructure investments by the ISPs and CPs.

where CTR (the 'click-through rate') is a widely adopted metric used to calculate the proportion of visitors who initiated transaction with an advertisement that redirected them to another site where they might purchase an item[8]. Here we assume that aCPs, and possibly eCPs, have ad revenues, with different popularity (hence different values of γs); In any case, we reasonably consider price-inelastic advertisers, and so we assume that γ is fixed and CPs cannot maximize on the ad-based charges.

- $q_a = q_e = 0$, in the free peering scenario.

In Table 1 we summarize the various parameters of our framework.

Our next step is the analysis of the profits of the eCPs and aCPs in the two cases of paid and free peering.

IV. ANALYSIS

Each provider's aim is to maximize its profit, that is, the revenue generated by its transaction with the end-users, minus investment costs. For each pricing scheme we study we find the symmetric equilibria for both free and paid peering, and compare each provider's profits in both regimes. This equilibrium is found within the context of a two-stage leader-follower game.

As the last-mile investments are more expensive and long-term than the content-side ones, the ISP acts as a Stackelberg leader and maximizes its profits with respect to q_e, q_a, and t. The CPs act as followers who, given the parameters chosen by the ISP, maximise their net benefits by choosing p, c_e, and c_a. More precisely, the stages are as follows.

i. Stage 1: ISP chooses q_e, q_a and t ($q_e = q_a = 0$ in the free peering case) by anticipating the optimal decisions by CPs in Stage 2.
ii. Stage 2: Based on the parameters chosen by the ISP, eCPs choose optimally p, c_e and aCPs choose c_a.

Note that the access price s is defined exogenously by the competitive access market, and not as in [21] to be freely chosen by the ISP.

A. Optimal Strategies

We solve the Stackelberg game by considering first the optimal strategy in stage 2 of the game, in which each eCP chooses the optimal end-user charge p and the optimal level of investments c_e, and aCPs choose only c_a.

Proposition 1 *The optimal p per transaction charged by each eCP to end-users is*

$$p = \theta_e + q_e b_e - \gamma_e. \tag{7}$$

Based on the solution for stage 2, the the ISP decides its profit-maximizing q_e, q_a, and t. The peering fees have the simple following form:

Proposition 2 *The optimal q_a is*

$$q_a = \frac{\gamma_a(1-u) - b_a s u}{b_a} \tag{8}$$

and the optimal q_e is

$$q_e = \frac{\theta_e(1-u) - b_e s}{b_e}. \tag{9}$$

It is interesting that a profit maximizing ISP does not always wish to extract positive payments from the CPs, since such payments would impact upon their infrastructure investments. This idea is summarised in the following corollary.

Corollary 1 *If the ISP earns significant access revenues, e.g., when s is large, then it may be that q_a and q_e are negative, implying that the ISP prefers to subsidize the CPs so they invest more in their infrastructure.*

The next corollary, although it follows directly from the structure of the optimal charges, will be justified by the economic interpretation of (8) and (9).

Corollary 2 *The optimal paid peering fees do not depend on the number of eCPs and aCPs in the market.*

Proposition 3 *The optimal c_a and c_e will be given by:*

$$c_a = \left(\frac{d_a t^w u e^{-(b_a s/\theta)}(\gamma_a - b_a q_a)}{z}\right)^{\frac{1}{1-u}}, \tag{10}$$

$$c_e = \left(\frac{d_e t^w u e^{-(b_e s/\theta + p/\theta_e)}(\gamma_e + p - b_e q_e)}{z}\right)^{\frac{1}{1-u}}. \tag{11}$$

Proposition 4 *If the ISP is not allowed to charge CPs for their transaction with its customer base, its optimal level of investments will be*

$$t = \left(\frac{sw(B^e_{fp} + B^a_{fp})}{k(1-u)e^{\frac{\theta + s(b_e + b_a)}{\theta}}}\right)^{\frac{1-u}{1-u-w}}, \tag{12}$$

with

$$B^e_{fp} = n_e \rho_e e^{\frac{s b_a}{\theta} + \frac{\gamma_e}{\theta_e}} \left(\frac{d_e \theta_e u}{z e^{\frac{\theta + s b_e}{\theta} - \frac{\gamma_e}{\theta_e}}}\right)^{\frac{u}{1-u}}$$

while

$$B^a_{fp} = n_a \rho_a e^{\frac{\theta + s b_e}{\theta}} \left(\frac{\gamma_a d_a u}{z e^{\frac{\theta + s b_a}{\theta}}}\right)^{\frac{u}{1-u}}.$$

In the paid peering regime, the optimal level of investments for the ISP, after having fixed the payments from the eCPs and aCPs, will be given by

$$t = \left(\frac{w(B^e_{pp} + B^a_{pp})}{k}\right)^{\frac{1-u}{1-u-w}}, \tag{13}$$

where

$$B^e_{pp} = n_e d_e \theta_e e^{-2+u+\frac{\gamma_e + b_e}{\theta_e} + \frac{b_a}{\theta}} (u/z)^{\frac{u}{1-u}}$$

and

$$B^a_{pp} = n_a d_a e^{-s b_a/\theta} (u^2/z)^{\frac{u}{1-u}}.$$

---

[8] `http://www.webopedia.com/TERM/C/CTR.html`

The solution of the equations leading to (12) and (13) suggests that the relative effect on the investments of the ISP in the cases of free and paid peering depends only on the fraction of the different types of CPs in the market and not on their absolute numbers.

Proposition 5 *If t_f, t_p are the investments of the ISP in the case of free and paid peering respectively, then t_f/t_p depends on n_a, n_e only through n_a/n_e.*

The proofs of the aforementioned propositions and formulas are presented in the Appendix.

B. Economic Interpretation

We first provide an economic interpretation to the key quantities that define the optimal charges. Let us rewrite (8), (9) using $\hat{q}_a = q_a b_a$, $\hat{q}_e = q_e b_e$, $\hat{s}_a = s b_a$, $\hat{s}_e = s b_e$ to reflect charges per transaction by the respective players. We find that the optimal $\hat{q}_a$ is

$$\hat{q}_a = \gamma_a(1-u) - \hat{s}_a u = \gamma_a - u(\gamma_a + \hat{s}_a), \tag{14}$$

the optimal $\hat{q}_e$ is

$$\hat{q}_e = \theta_e(1-u) - \hat{s}_e = \theta_e - \hat{s}_e - u\theta_e\,, \tag{15}$$

and the price of the eCP is

$$p = \theta_e + \hat{q}_e - \gamma_e\,. \tag{16}$$

We are now able to make some interesting observations. We start with the case of the eCP. From (16)

$$p + \gamma_e - \hat{q}_e = \theta_e\,, \tag{17}$$

i.e., the price p is chosen so that the *net profit* per transaction of the eCP is *always* θ_e (both in free and paid peering).

Consider the base case in which $u = 0$, i.e., eCP investments have no effect on customer traffic, and hence these investments will be zero. From (15) it turns out that $\hat{q}_e + \hat{s} = \theta_e$, suggesting that the profits per eCP's transaction of the ISP and eCP are equal, providing some elements of fairness. Hence larger values of $\hat{s}$ imply the need for lesser charges $\hat{q}_e$.

When $u > 0$, then $\hat{q}_e + \hat{s} = (1-u)\theta_e$ and the ISP gets a smaller share of the eCP's profit. The charge drops since the ISP benefits more by charging the eCP less, in order for the later to invest more in infrastructure and increase traffic. This net drop of the ISP revenues per transaction is proportional to u. We summarise the above in the following proposition.

Proposition 6 *In the case of an eCP, the following hold:*

i. In cases of both free and paid peering, the eCP obtains a fixed net profit per transaction equal to θ_e, independent of the rest of the system parameters.

ii. In the case of paid peering the charge $\hat{q}_e$ is chosen so that the total net benefit per eCP transaction of the ISP $\hat{q}_e + \hat{s}$ is equal to $(1-u)$ times the profit of the eCP.

Corollary 3 *For small values of u (as practical sense suggests) the ISP and the eCP obtain similar profit per transaction, equal to θ_e.*

We turn now to the case of aCPs. From (14) we observe that if $u = 0$, then $\hat{q}_a = \gamma_a$. This implies that the ISP *acting as a monopolist* takes back all the profit of the aCP, on top of his access revenue per transaction $\hat{s}$. This high charge has no effect on the traffic of the aCP since it is not being transferred to the aCP's customers generating the traffic (as in the case of an eCP, see (16)). It could result in higher prices for ad spaces in the longer term, if this market is not competitive. But again any increase of γ_a would lead to an increase of $\hat{q}_a$.

When $u > 0$ and hence aCP investments have a positive effect on user traffic, ISP charges are decreased to leave some surplus to the aCP to invest in infrastructure. As expected, (14) suggests that higher values of u imply a faster drop of the charges as a function of $\hat{s}$, to promote higher investments by the aCP.

Proposition 7 *In the case of an aCP, for small value of u, the ISP acts as a monopolist and takes away all the profits from the aCP.*

Let us compare now the profit situation of the different CPs for small values of u[9]. In the case of eCPs, these obtain the same profit per transaction θ_e independently of free or paid peering. Hence their preference for paid peering v.s. free peering will be determined by which regime offers a higher rate of transactions (i.e., more traffic).

Note from (15) that for $u \approx 0$ and $\hat{s}_e = \theta_e$, the charge in the paid peering case is zero, hence the price p charged to the eCP customers is the same as in the case of free peering. The same holds for the investment levels c_e, t, implying that the transaction rates are the same. If $\hat{s}_e$ decreases from θ_e, $\hat{q}_e$ increases leading to higher p and also to higher investment t from the part of the ISP. These have contradicting effects on the user traffic. If w is small, then the increase of t will have a small positive impact on the user traffic, hence the overall effect will be determined by a decrease of the transaction rates due to the higher price p. If w is large, then this effect might be reversed.

If $\hat{s}_e$ is larger than θ_e, the eCP gets a subsidy from the ISP and reduces its price p to its customers (see (16)), increasing their demand. However, the ISP now invests less in infrastructure than in the case of free peering (since it has less total revenues). But it easy to prove that the overall effect is to strictly increase user demand for transactions compared to the case of free peering. Since using $q_e = 0$ is always a choice for the ISP in the case of paid peering, if $\hat{q}_e < 0$ maximises its profits, it must come with increased traffic to justify a lower total revenue for transaction $\hat{s}_e + \hat{q}_e$ (compared to $\hat{s}_e$ in the case of free peering). Hence if $\hat{s}_e > \theta_e$, the eCP is always better off under paid peering.

Paid peering clearly has a devastating effect on the profit of an aCP, unless u and $\hat{s}_a$ are large and then a substantial fraction of profits will be left to aCP. Similarly as for eCPs, large $\hat{s}_a$

---

[9] A reasonable practical approximation is to take $u = 0$ and $w > 0$ since most of the significant infrastructure upgrades that will positively affect user experience are expected to come by upgrading network infrastructure (in the access network, peering points, deep caching, etc). Content infrastructure is not as expensive and is easily over-dimensioned from the start. Hence we expect in practice much smaller values of u compared to w. See for example: `https://freedom-to-tinker.com/blog/felten/last-mile-bottleneck-and-net-neutrality/`

can result in subsidies to the aCP and also increase the total transaction rates.

For small values of s near zero, both eCPs and aCPs profit from paid peering. Without paid peering, the ISP has almost zero revenue from the user traffic, and hence invests $t \approx 0$ which 'kills' all demand for transactions. Hence even a small ISP revenue obtained by paid peering can cause a significant improvement compared to free peering. We summarise our previous observations in the following proposition. Recall that s parameterizes revenue that the ISP is able to capture from user traffic.

Proposition 8 *Comparing the profits of CPs under free and paid peering when u is small:*

i. *aCPs are expected to be much worse off under paid peering unless s is either near zero or very large.*
ii. *eCPs are expected to be better off under paid peering if the investments of the ISP have a significant effect on user traffic, i.e., when w is large. When w is small, we expect the eCPs to be better off under free peering for an intermediate range of values for $\hat{s}_e$, where $\hat{s}_e < \theta_e$.*

Corollary 4 *If the infrastructure investments by the ISP have a small impact on user traffic, the incentives of the eCPs and the aCPs are (mostly) aligned regarding adopting paid or free peering. In the opposite case, eCPs are in favour of paid peering, while aCPs prefer free peering unless the ISP has very low or very high revenue from user traffic.*

Below, we summarize some key practical remarks that can be easily made by reference to the formulas of the optimal p, q_a and q_e.

Remark 1: eCPs will pass to the end-users the ISP charge, while this is not possible for aCPs. As (7) implies, eCPs will charge end-users by taking into account the peering payments. This has major implications in the profitability of the eCPs in the paid peering regime, as we discussed earlier.

Remark 2: Is paid peering always detrimental for aCPs? If one argues that (i) attractiveness of content needs substantial infrastructure investments by aCPs, i.e., u is substantially above zero, and (ii) the income of the ISP from access traffic is also substantial, i.e., $s \gg 0$, then the aCP is left with a large fraction of its revenues from the ads. This fact, combined with increased traffic volumes, can make paid peering not such a bad alternative.

Remark 3: Unit payments for heavy traffic applications are lower. As the traffic volume per transaction increases (b_e and b_a), the per unit of volume peering payments from the CPs (either eCPs or aCPs) to the ISP decrease, eventually becoming negative (see (8) and (9)). Consequently, our model implies that light types of content will contribute under paid peering proportionally more to the ISP revenues, part of which will go to last-mile investments needed to support the deliverance of heavier and QoE demanding applications. It is the incentive of the ISP to offer quantity discounts to volume-intensive CPs for delivering their content over the fast-lane, as the traffic generated by these CPs becomes very important to its profits.

Moreover, by comparing the paid peering fees of the eCPs and aCPs ((8) and (9)), we observe that q_a decreases more slowly compared to q_e as the traffic volume per transaction increases (since $u < 1 - w$). This shows that eCPs obtain even more favourable discounts for greater transaction traffic, providing some strong indications of better profitability under paid peering.

Remark 4: Doubling the volume per transaction b has the same effects on profits of all players as doubling the value of s. This is easy to see since the profits of all players are defined on a per transaction basis. Hence larger values of volume per transaction have the same effects as larger values of ISP profitability per unit volume (s). This also explains why unit payments for heavy traffic applications are less. As b increases, for a constant s, the profit per transaction of the ISP increases. It is then more willing to allow for a larger profit margin (per transaction) of the CP's in order to induce more investments.

Remark 5: External market conditions do not affect the peering payments. As mentioned in Corollary 2, the ISP does not need to take into account the exact number of eCPs and aCPs in choosing the optimal payments. Thus, the negotiations between an ISP and a CP can be done on a bilateral basis, without considering the external conditions of the market. This suggests a simpler adoption of paid peering since optimal charges have a very simple form.

Having highlighted some key qualitative findings of our framework, we next study the effects of the ISP's pricing decisions on the profits of the different actors of our model, under several market scenarios.

V. NUMERICAL DEMONSTRATION OF PRICING EFFECTS

We now consider how paid peering payments affect the profits of each player in the ecosystem, and to what extent. We do this by comparing the symmetric Nash equilibria of the free v.s. the paid peering regimes (being discussed in Section IV), in terms of the ratios r_{isp}, r_e, and r_a, where for $x \in \{isp, e, a\}$,

$$r_x = \frac{\pi_x(\textit{free peering})}{\pi_x(\textit{paid peering})}.$$

A. Parameterization

Throughout our experiments, we consider the access price s as an independent variable, and plot results against it. We do not consider it to be a control parameter since it is exogenously defined, but rather focus on its impact on the results.

We assume that the end-users' demand elasticity towards p in the free peering regime is equal to 1. Based on this assumption we find that $\theta_e = 0.36$. As we have already shown that q_e does not depend on the ad-based profits of an eCP, and the most significant active eCPs are subscription-based exclusively, we assume that $\gamma_e = 0$. This further implies that $p > 0$ (see 7), and consequently, the eCPs will always get paid by the end-users, if we are in the free peering regime.

Our main purpose is to investigate how the business model of a CP affects its profits, if it agrees on a paid peering deal with the ISP. Therefore, we use its free peering net benefit as a benchmark, and initially assume that $\pi_e(\textit{free peering}) = \pi_a(\textit{free peering})$, implying $\gamma_a = 0.13$.

Table 2. Model parameterization.

u	w	k	z	n_e	n_a	d_e
0.1	0.4	3	1.5	10	10	5

d_a	γ_e	γ_a	b_e	b_a	θ	θ_e
5	0	0.13	0.7	0.7	3	0.36

As we consider a symmetric scenario, we assume that both the eCPs and the aCPs have same popularity of content, i.e, $d_a = d_e$, and that the number of eCPs is equal to the number of aCPs, i.e., $n_e = n_a$. Our symmetric scenario further implies equal traffic volumes per transaction, i.e., $b_e = b_a = 0.7$.

Moreover, we assume that the ISP's marginal investment cost is greater than the CPs (i.e., $k > z$), taking into account that the last-mile infrastructure is a far more capital intensive investment, in comparison with CP-side investments, such as the establishment of CDNs or cloud services, closer to the end-users. We also take $u = 0.1$ and $w = 0.4$, assuming that investments in the last-mile will be more valuable to the end-users compared to an additional cache deployed by the CPs.

Mention that what really determines our results, is not the exact value of a parameter, unless we refer to u or w, but its numerical relation with the rest variables of our ecosystem. We summarize the initial input of our model in Table 2.

B. Evaluation

In the first part of our evaluation, we investigate some general trends that define the incentives of the various players to establish the HQ infrastructure. We do so by examining an ecosystem in which only aCPs or eCPs participate, e.g., $n_e = 0$ or $n_a = 0$, respectively. Subsequently, we relax some of those assumptions, to look at several market scenarios and identify which general conclusions can be extracted from our framework, especially as regards paid peering profits of stakeholders.

Finally we will examine a more realistic scenario where both types of providers co-exist in the HQ infrastructure. In this case the level of profitability of each type of provider defines which part pays for a larger share of the HQ platform. Hence, we mainly like to investigate the sensitivity of the various CPs on free-riding issues.

Ratio of profits for each type of actor. We apply the numerical values, shown in Table II, but as already mentioned we initially assume that either aCPs or eCPs participate in the HQ platform. We demonstrate our results starting from the ISP's ratio of profits and gradually explain how the parameters of our framework affect the ratio of profits of the eCPs and aCPs. Based on that input we plot the ratio of profits of the multiple actors against values of the access price s.

As Fig. 2 shows, the ISP in the paid peering regime is generally more profitable than in the equivalent free peering case, for any end-user access price s. Notice that that the only case it is indifferent of the pricing regime is when the market conditions forced it to set q_e or $q_a = 0$. In any other case, it prefers paid peering, although it may actually pay the CPs (when s is high). Fig. 2 corresponds to the case where $n_a = 0$, but similar conclusions can be reached for a scenario where only aCPs exist in the ecosystem. Furthermore, we have found that the ISP's level of investments is always higher in the HQ infrastructure.

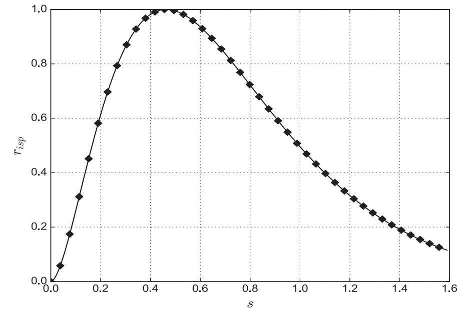

Fig. 2. The figure shows the ratio of profits of the ISP (r_{isp}), for $n_a = 10$ and $n_e = 0$. We observe that we always have that (r_{isp})≤ 1, implying there is no case in which the ISP will be more profitable in the free peering. This finding holds for any combination of n_a and n_e.

We next proceed with the investigation of the ratio of profits of the eCPs and aCPs (r_e and r_a), under several alternative market scenarios.

The general story is the following.

i. For small value of s, the ISP makes very little revenue in the free peering regime and hence invests very little, reducing the demand of the CPs. This is corrected in the paid peering regime where positive qs generate additional revenue to the ISP. This profit sharing mode incentivises more investments and everybody is better off. As s increases, more investments take place by the ISP making everybody better off in absolute terms (profit ratios below 1), but the marginal benefit of these investments to the CPs decreases (profit ratios increase towards 1, making paid peering marginally less beneficial).

ii. Intermediate values of s. Here the effect of the paid peering revenue to the ISP investment incentives level off and the "greed" of the ISP takes more effect. The ISP was already making enough revenue in the free peering regime and the additional peering revenues will not greatly influence his investment decisions. Hence some CPs *may* be worse off compared to free peering (profit ratios above 1).

iii. Large values of s. The ISP makes so much revenue from access traffic that it finds it more profitable to subsidize CPs ($q < 0$) for increasing the total infrastructure investments and hence increase further its profits. This eventually benefits all players in the ecosystem (profit ratios below 1).

These trends are observed throughout our simulation results, but we notice some striking differences between the ratio of profits of aCPs and eCPs, which support our discussion in Proposition 8. As Fig. 3(a) shows the range of the intermediate values of s where $r_a \geq 1$ is significantly larger than the range of s in which $r_e \geq 1$. Moreover, we constantly have that $r_a > r_e$, implying that the profits of the aCPs in the paid peering will always be lower than the eCPs' profits. Mention that we have initially assumed that both types of providers will have equal free peering profits. Eventually, for large s, q_e, and q_a become a subsidy and everybody again prefers paid peering.

It is difficult to deduce the actual value of s in the current Internet market to determine where we lie in practice in the above figures. To avoid speculations we propose the "positive peering charge assumption", which is a simple heuristic to find a plau-

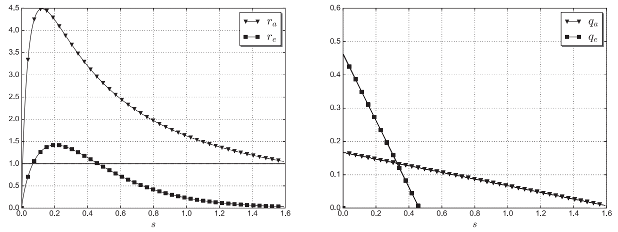

(a) (b)

Fig. 3. The figure presents the ratios of an aCP (r_a) and an eCP(r_e) (blue and purple line respectively). These figures correspond to a market in which only CPs of the same type are active. It turns out that for small values of u, we always have that $r_a > r_e$, while aCPs prefer free peering for a significantly larger range of s. Fig. 3(b) depicts that q_e will fall faster than q_a as the access price s increases, implying that the ISP decides to offer some further incentives to the eCPs to join the HQ platform, for higher values of s. Mention that for values of s higher than the ones shown in Fig. 3, then both q_e and q_a will be negative.

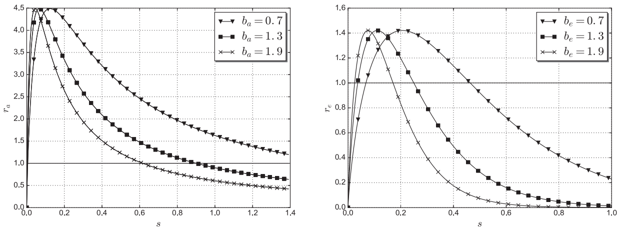

(a) (b)

Fig. 4. The figure suggests that the more bandwidth-intensive is an aCP, the more profitable will be in the paid peering regime, while it also reduces the range of s where $r_a > 1$. Similar conclusions we are able to make for the eCPs, in Fig. 4(b). The above figures correspond to a scenario where only one type of CP directly interconnects with the ISP.

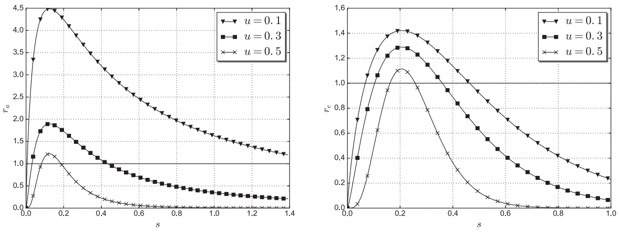

(a) (b)

Fig. 5. The figure shows how different values of u (sensitivity of demand with regard to each CP's investments), affect the preference of a CP for free peering vs paid peering, when a single type of CP peers with the ISP. We observe that if investments have strong effects on demand (large values of u), paid peering becomes almost overwhelmingly the best strategy for the CPs. Figs. 5(a) and 5(b) suggest that this parameter has a much stronger influence for aCPs than for eCPs. For small to moderate values of u, eCPs seem to prefer paid peering for a substantial range of access prices s.

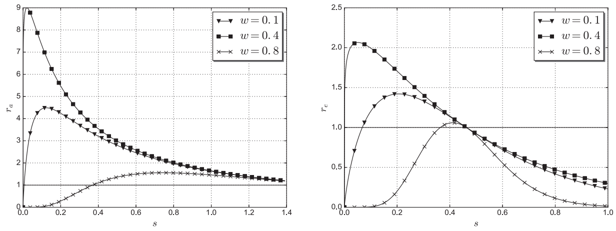

(a) (b)

Fig. 6. The figures (a) and (b) depict r_a and r_e, respectively, for different values of w (sensitivity of demand with regard to the ISP's investments). In Fig. 6(a) we assume that $n_e = 0$, while $n_a = 0$ in Fig. 6(b). Unless the ISP obtains substantial profits, due to the significant increase of the end-users transaction rates (when w is large), both aCPs and eCPs are worse off under paid peering, with the aCPs suffering the greatest economic loss.

sible range of values for s that has practical meaning. The idea is that we expect the current market parameters to justify positive peering charges and no subsidisation from the CP side to the ISP. Hence we propose as our practical range for s the values for s for which $q_a, q_e > 0$.

Fig. 3(b) uses the aforementioned practical range of s in which q_e and $q_a > 0$. It first verifies that as s increases q_e falls faster than q_a, as ((8) and (9)) imply. We notice that q_e becomes negative for values of s higher than 0.48, while the ISP pays aCPs when $s > 1.74$. As the eCPs will pass to the end-users the additional payments, the ISP decides to offer them some extra discounts for high access price, in order to (indirectly) decrease p.

In general when the "positive peering charge assumption" is applied, we notice that only if s is near 0, and hence ISP infrastructure investments are very low, a peering charge can be beneficial for the CPs. The shape of the graphs is of course influenced by our choice of the demand function.

We next proceed with the investigation of the impact of several market parameters on r_e and r_a. As highlighted in Section IV we identified that the following parameters have the most significant effect on the paid peering profits of the CPs, when we examine them in isolation.

i. The volumes of traffic per transaction, i.e., b_e and b_a.
ii. The end-user traffic sensitivity to the ISP's and CPs' investments, i.e., the values of w and u, respectively.

We initially investigate the impact of different values of b_e and b_a, and in what way varying values of w and u affect the CPs' paid peering profits.

Effect of the volume of traffic b_e and b_a. As discussed in Section IV, the more volume-intensive is a CP (either eCP or aCP), the less it has to pay to the ISP per volume of traffic. Moreover, as Figs. 4(a) and 4(b) show, the more content-heavy is a CP (large b per transaction), the more it is benefited by the establishment of the paid peering regime, as its ratio of profits (free v.s. paid peering profits) along with the range of s where r_a or $r_e > 1$ are both decreasing.

We offered earlier an explanation of this: higher values of b imply higher profits per transaction for the ISP. This has a direct effect in reducing the charge per transaction for the eCP by the amount of the extra profit of the ISP. Similarly, in the case of aCPs, when $u > 0$, higher ISP profits imply higher aCP profit margins to induce investments.

Effect of u and w. By comparing the plots of r_a and r_e in Figs. 5(a) and 5(b) respectively, we notice that although both types of CPs will observe an increase of their paid peering profits if u is high, this effect is significantly more profound on r_a. This is explained by Propositions (6) and (7). Starting from a value of u near zero, when this increases, aCP profits increase from zero to some positive value. For eCPs, this change of left-

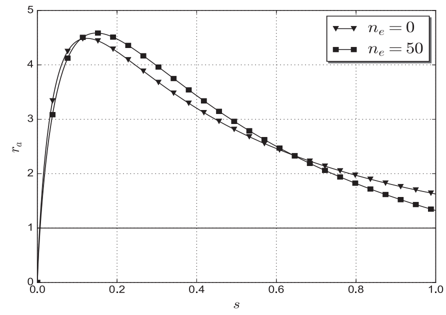

Fig. 7. The figure demonstrates r_a for different values of n_e, while $n_a = 10$. We deduce that aCPs prefer paid peering for a larger range of s, if the ecosystem also includes eCPs. In this case, and for high values of s, end-users will actually be subsidized by eCPs, which eventually increase the ISP's profitability and its level of investments, leading to higher click rates towards the various aCPs. For the rest variables of our model we have relied on Table 2.

over profits is not so acute.

The trends regarding the effect of w are as mentioned in Corollary 4. We observe that eCPs may be always more profitable in the paid peering regime for high values of w, since this implies rates of transactions that are high enough to compensate the reduction of user demand due to higher end-user prices. But aCPs seem are always worse off under paid peering, even for large value of w. This is because the ISP takes away most of their revenues, and any reasonable increase in the transaction rates due to increased investments is not able to compensate for the substantial reduction of profit per transaction they incur under paid peering.

Examining aCPs and eCPs in coexistence. By considering a scenario where n_a and $n_e > 0$, we study which type of CP benefits more by co-existing with the other. Let us try to answer this question by using our economic analysis results. In Fig. 7(a), we illustrate r_a for $n_e = 0$ (blue line) and $n_e = 50$ (purple line), while in both cases we assume that $n_a = 10$. As shown, although for small values of s the maximum value of r_a is marginally higher for the latter scenario, we deduce that as s increases, aCPs will be favored by the participation of more eCPs in the common HQ infrastructure, since this will increase the ISP's level of investments. On the contrary, we do not detect such substantial variations of r_e for different populations of aCPs, and hence plots for this scenario are omitted.

Another question is how the relative profitability of the CPs affects their ratio of profits. To answer this, we still rely on the input of Table 2, but with varying values of γ_a in order to investigate two extreme cases, that may occur in the free peering case:

i. π_e*(free peering)* $= 5\pi_a$*(free peering)*, i.e., $\gamma_a = 0.032$.
ii. π_e*(free peering)* $= \pi_a$*(free peering)*/5, i.e., $\gamma_a = 0.58$.

Fig. 8(a) corresponds to Case i, while Fig. 8(b) depicts r_e and r_a for Case ii.

In Case i, since aCPs obtain small revenues from their customers, most of the ISP's income comes from the eCPs, leaving some profits to the aCPs. Since most of the ISP infrastructure investment comes from the eCPs, the aCPs enjoy a higher quality ISP infrastructure by contributing only a small share. This 'free-ridding' effect makes them much better off compared to the case they would alone in the system. For example, if $s = 0.5$, according to Fig. 8(a), we obtain that $r_a \approx 1$, while if they were alone in the system they would be 3.6 times worse off, as the blue line in Fig. 3(a) shows.

In Case ii, the reverse occurs and most of the ISP revenue comes from the aCPs, offering a free-ride to the eCPs. Notice that for $s = 0.25$, Fig. 8(b) suggests that now eCPs are not any more worse off under paid peering, for any value of s.

We further notice that when $\gamma_a \ll \theta_e$, aCPs' free peering profits will be up to 3 times higher than the equivalent paid peering ones. On the contrary, they will observe a substantial decline of their economic performance (up to 5 times), when $\gamma_a \gg \theta_e$. In the latter case, we also observe a significant increase of the range of s where π_a*(free peering)* $>$ π_a*(paid peering)*.

To sum up, it appears that aCPs will be significantly harmed by the establishment of paid peering if they earn more, under free peering, than the eCPs. In an attempt to provide a practical set of values for θ_e and γ_a, we refer the reader to our discussion in Section IV-B. Our interpretations of the various quantities suggest that in practice θ_e, corresponding to the net profit per transaction of the eCP by selling digital goods, should be larger than γ_a, i.e., $\theta_e \gg \gamma_e$. In practice, we expect that θ_e is of the order of a \$, but γ_a to be much less. Therefore, our initial input as shown in Table 2, or the modified one considered in Case i, seem to be more close to the current market practices than Case ii. Hence eCPs will contribute for a more significant part of the ISP revenues, and the free-ridding effect on the aCPs makes paid peering less devastating to their profits.

Finally, in order to provide a more spherical view of our model and as the whole debate may touch upon regulatory issues, we next investigate the implications of paid peering on Social Welfare.

Consumer surplus and social welfare issues: Social welfare (SW) is defined by the sum of the consumer surplus (CS) with the profits of the several self-interest entities. We compute the CS as the difference between the integral of the end-users' total demand for content and their total expenses to the ISP and the eCPs:

$$CS = \int_0^s (n_aD_a + n_eD_e)ds - s(n_a\rho_a + n_e\rho_e) - (pn_eD_e). \quad (18)$$

We remind the reader that D_a and D_e denote the transaction rates of end-users with and aCP and eCP, respectively. Hence, their utility will be given by the integral of the their total demand when n_a and n_e aCPs and eCPs, respectively are active in the market. Moreover, end-users will be charged by the ISP with a price s per unit of traffic, generated by their transactions with the CPs of the ecosystem, while on top of that they will also pay a price p each time they transact with an eCP's content.

Hence, we compute the SW as

$$SW = CS + \pi_{isp} + n_a\pi_a + n_e\pi_e. \quad (19)$$

As the imposition of peering fees affects the level of investments of the multiple players, the rate of transactions and consequently the CS will be different in the free and paid regimes.

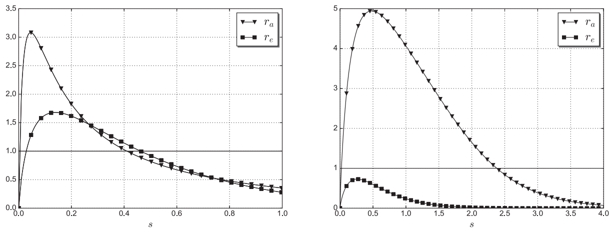

(a) (b)

Fig. 8. The case where both aCPs and eCPs agree on a paid peering deal. Fig. 8(a) corresponds to the case where eCPs are more profitable (per transaction) than the aCPs (i.e., $\gamma_a = 0.032$) while the reverse holds in Fig. 8(b), where $\gamma_a = 0.58$. It turns out, that less profitable types of CPs will free-ride on the HQ infrastructure, which implies that their economic performance under paid peering, will be relatively improved.

Throughout our simulations we have identified that CS is significantly improved under paid peering, and this finding becomes more profound as s increases. The fact that the ISP's level of investments is constantly higher for the HQ infrastructure (as already discussed), has a direct impact on the end-users rate of transactions, due to the improved levels of QoE. Thereafter, a practical conclusion of our analysis is that if usage-based pricing prevails, end-users will be benefited by the imposition of peering fess to the various CPs, especially if they pay high usage-based fees, compared to their price elasticity for access traffic (denoted as θ).

In order to compare the SW under both regimes, we use the following ratio:

$$r_{SW} = \frac{SW(\textit{free peering})}{SW(\textit{paid peering})}.$$

We use the input presented in Table 2, and investigated r_{SW} for ranging values of access price s. Throughout our experiments, we deduced that for the intermediate values of s we may have that r_{SW} will be *marginally* higher than 1, but for small or high values of s, SW seems to be improved in paid peering.

Of course these conclusions are for the symmetric case presented in Table 2. We have additionally considered a more realistic scenario which assumes that $b_e > b_a$, and found out that the more heavy are the eCPs, the greater is the SW under paid peering.

This is explained by the fact that the ISP will obtain higher revenues per transaction, while q_e will be lower for high values of b_e, as (9) implies. Consequently, p will be lower in comparison with the initial input, leading to less decrease of the user demand. Contrariwise, r_{SW} will not be significantly affected by higher values of b_a, since this will not have a direct impact on end-users demand, having assumed small u.

VI. DISCUSSION

In this section we revisit and discuss some assumptions of our framework.

Access infrastructure: Throughout this work, we consider a network that is designed bottom up, i.e., from scratch, specifically for the given comparison without assuming some pre-existing situation. In fact, regulators use this bottom-up model to compute costs of services, since the existing network might be inefficient and one needs some independent benchmark to get numbers. So we use this approach to justify our model used in both pricing cases. However, we have crosschecked our results by investigating a scenario in which the ISP builds upon the existing (before paid peering) access infrastructure, and verified the robustness of our main results, discussed in Section V, no matter the actual level of the existing investments, e.g., low or high.

Usage-based access pricing: Although flat-rate models, in which end-users pay a fixed price (independent of usage), have been the dominant norm of pricing in the last-mile, recently a number of ISPs are considering moving towards usage-based tariffs. Some examples of such ISPs include Comcast (as already noted), Verizon and Deutche Telekom. Their main argument is that usage-based billing establishes a direct mechanism to manage traffic as it ties pricing to end-users' demand.

An alternative approach could be that even if the ISP uses traditional fixed-based billing, there can be a price $s > 0$ since by increasing customer demand because of new advanced content services (like HD video) the overall traffic generated by a typical customer will increase and it will incur extra cost by needing to buy more expensive contracts with larger total volume allowances. It is very difficult to model accurately such effects of traffic increases, and to keep our model simple we use s to denote this ex-post measured extra cost per user per click. On longer time scales, the ISP could affect s by changing its access contract portfolio, but we are not concerned with this control aspect in the present work. In our model the implied cost s per click has two effects. It reduces uniformly the demand for content, and provides a revenue for the ISP.

In any case, our model of the end-users' demand for content is general enough to be also applied in wireless networks, where usage-based pricing is commonplace. Indeed, since nowadays there is an ongoing discussion, mostly in the U.S and the E.U. on the economic implications of sponsored mobile data (see for example [27]), (i.e., whether CPs should be permitted to cover a part of users' cost of downloading mobile data), our framework could be extended to examine such issues in wireless networks.

Competition issues-What if eCPs are price takers? Our analysis so far, was based on one extreme form of competition, the case of eCPs selling independent content and hence acting like monopolists choosing freely their price p. In an extension of our work we have examined the other extreme scenario in which eCPs are in perfect competition with each other, and hence p is defined exogenously. In this case, it turns out that for a significant range of s, the eCPs' profitability is deteriorated under paid peering, and in many cases their ratio of profits is even worse than the equivalent metric for the aCPs. Thereafter, it turns out that in the case of a perfectly competitive market for the eCPs, their business model trivializes and becomes similar to the aCP market (i.e., the only strategic action would be the infrastructure investment level), implying that under paid peering most of their surplus of the CPs will flow to the ISP.

VII. CONCLUSION

In this work we provide some further insights on the complexities around the paid peering debate. The whole discussion is highly controversial; in general ISPs favor the imposition of peering payments, while CPs do not.

We have established a framework which provides some practical insights on whether the CPs' opposition to paid peering is rational or not, since one may argue that their transferred payments can incentivize the ISP to make greater investment in the common infrastructure. We have studied an ecosystem consisting of a single access ISP that directly interconnects with multiple CPs, doing so in two extreme situations: When paid peering is used for all the involved CPs, and when all parties agree on a settlement-free deal. In the case of paid peering we suppose that the ISP charges the CPs so as to maximize its profit.

We have analyzed the Stackelberg equilibrium, finding the total amount of investments and profits of the various stakeholders, when their revenues are affected by the total investment in infrastructure and the charges that result from the different agreements. We deduce that, although in general the overall social welfare seems to be improved by paid peering, it turns out whether or not there is benefit to a CP highly depends on its business model. The ISP, having implemented usage-based pricing, will not overcharge the eCPs, due to their ability to pass the peering costs to the end-users. Conversely, the fact that aCPs can not directly affect the rate of clicks of the end-users, allows the ISP to act as a monopolist and extract all their ad-based profits, especially if their level of investments is not significant for the end-users. Interestingly enough, although we have assumed that eCPs may also earn additional ad-based revenues, this does not affect the peering payments to the ISP.

Peering payments also depend on the volume (in traffic units) per customer transaction, since we find that CPs with high volume per transaction will be charged less per traffic sent, and so seem to be benefited more by the introduction of paid peering fees. We have also detected a crucial role for the end-users' evaluation of the level of investments of the various market players. Another relevant factor is the access price that an ISP charges its customers. However, we have modeled this as exogenously defined by competition in the access network, with the ISP being a price taker.

As a final remark, we point out at a strong relation of our analysis to the network neutrality debate. Even if one could argue that paid peering deals are "business as usual" and should not be correlated with network neutrality issues, our free peering scenario could be associated with the current neutral network, while the paid one corresponds to a non-neutral ecosystem in which the ISP imposes traffic tolls to each CP transacting with its customer-base, in order to upgrade its access infrastructure and consequently improve its levels of profitability.

REFERENCES

[1] C. Courcoubetis, K. Sdrolias, and R. Weber, "Pricing the fast-lanes: A qualitative study on the implications of paid peering agreements," in *Proc. IEEE ICC*, 2016.

[2] W. B. Norton, *The Internet peering playbook: Connecting to the core of the Internet*, DrPeering Press, 2012.

[3] C. Dovrolis, "The evolution and economics of Internet interconnections", *submitted to the Federal Communications Commission (FCC)*, 2015.

[4] Sandvine, Global Internet phenomena report, 1H 2014.

[5] Netflix, Netflix Open Connect, [Online]. Available: https:// openconnect.itp.netflix.com/

[6] Arstechnica, Netflix says 99 percent of its links with ISPs are unpaid, [Online]. Available: http://arstechnica:com/information-technology/2014 /08/netflix-sends-99-percent-of-its-traffic-over-free-connections-to-isps

[7] Arstechnica, Netflix performance on Verizon and Comcast has been dropping for months, [Online]. Available: http://arstechnica:com/information-technology/2014/02/netflix-performance-on-verizon-and-comcast-has-beendropping-for-months

[8] Gigaom, Confirmed: Comcast and Netflix have signed a paid peering agreement, [Online]. Available: https://gigaom:com/2014/02/23/ confirmed-comcast-andnetflix-have-signed-a-peering-agreement

[9] T. Wheeler, M. Clyburn, J. Rosenworcel, A. Pai, and M. O'Rielly, "Notice of proposed rulemaking, in the matter of protecting and promoting the open Internet," *Federal Communications Commission (FCC)*, 2014.

[10] C. Courcoubetis, L. Gyarmati, N. Laoutaris, P. Rodriguez, and K. Sdrolias, "Negotiating premium peering prices: A quantitative model with applications," *ACM Trans. Internet Tech.*, vol. 16, no. 2, Apr. 2016.

[11] D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden, "Tussle in cyberspace: Defining tomorrow's Internet," *IEEE/ACM Trans. Networking*, vol. 13, no. 3, pp. 462–475, June 2005.

[12] Z. B. Houidi and H. Pouyllau, "The Price of tussles: Bankrupt in cyberspace?," in *Proc. workshop on pricing and incentives in networks*, June 2012.

[13] L. He and J. Walrand, "Pricing and revenue sharing strategies for Internet service providers," *IEEE/ACM Trans. Networking*, vol. 24, no. 5, pp. 942–951, May 2006.

[14] G. Shrimali and S. Kumar, "Paid peering among Internet service providers," in *Proc. ACM GameNets*, 2006.

[15] C. Courcoubetis and R. Weber, "Economic issues in shared infrastructures," *IEEE/ACM Trans. Networking*, vol. 20, no. 2, pp. 594–6081, Aug. 2011.

[16] S. Li and J. Huang, "Price differentiation for communication networks," *IEEE/ACM Trans. Networking*, vol. 22, no. 3, pp. 703–716, June 2014.

[17] P. Faratin *et al*, "Complexity of Internet interconnections: Technology, incentives and implications for policy," in *Proc. TPRC*, 2007.

[18] A. Lodhi, A. Dhamdhere, and C. Dovrolis, "GENESIS: An agent-based model of interdomain network formation, traffic flow and economics," in *Proc. IEEE INFOCOM*, 2012.

[19] R. T. B. Ma, D. M. Chiu, J. C. S. Lui, V. Misra, and D. Rubenstein, "Internet economics: The use of shapley value for ISP settlement," *IEEE/ACM Trans. Networking*, vol. 18, no. 3, pp. 775–787, 2010.

[20] J.C. Rochet and J. Tirole, "Platform competition in two-sided markets," *J. the European Economic Association*, vol. 1, no. 4, pp. 990–1019, 2003.

[21] J. Musacchio, G. Schwartz, and J. Walrand, "A two-sided market analysis of provider investment incentives with an application to the net-neutrality issue", *Review of Network Economics*, vol. 8, no .1, pp. 22–39, 2009.

[22] P. Njoroge, A. Ozdaglar, N. Stier-Moses, and G. Weintraub, "Investment in two-sided markets and the net neutrality debate," *Review of Network Economics*, vol. 12, no. 4, pp. 355–402, 2013.

[23] Y. Wu, H. Kim, P. H. Hande, M. Chiang and D. H. K. Tsang. Revenue Sharing among ISPs in Two-Sided Markets. in *Proc. IEEE INFOCOM*, 2011.

[24] E. Altman *et al., A study of non-neutral networks with usage-based prices, Incentives, Overlays, and Economic Traffic Control*, Springer Berlin Heidelberg, pp. 76–84, 2010.

[25] Cisco, "Moving toward usage-based pricing: A connected life market watch perspective," *Cisco Internet Business Solutions Group*, Mar. 2012.

[26] FierceCable, Comcast says it publicly outlined its usage-based pricing philosophy back in 2012, [Online]. Available: http://www.fiercecable.com/cable/comcast-says-it-publicly-outlined-its-usage-based-pricing-philosophy-back-2012

[27] C. Joe-Wong, S. Ha, and M. Chiang, "Sponsoring mobile data: An economic analysis of the impact on users and content providers," in *Proc. IEEE INFOCOM*, 2015.

APPENDIX

Proof: **[eCPs' optimal strategy.]**

Each eCP is faced with the problem of maximizing its profit of

$$\rho_e(\gamma_e/b_e + p/b_e - q_e) - zc_e.$$

This is the same problem as maximizing

$$d_e c_e^u t^w e^{\frac{\gamma_e}{\theta_e}} e^{-\frac{q_e b_e}{\theta_e}} e^{-\frac{s b_e}{\theta}} e^{-\frac{\gamma_e + p - q_e b_e}{\theta_e}} (\gamma_e/b_e + p/b_e - q_e) \\ - z c_e.$$

The optimum p is therefore such that $\gamma_e + p - q_e b_e = \theta_e$ which implies

$$p^* = \theta_e + q_e b_e - \gamma_e,$$

which makes the maximal value of the profit to be

$$b_e D_e c_e^u t^w e^{\frac{\gamma_e}{\theta_e}} e^{-\frac{q_e b_e}{\theta_e}} e^{-\frac{s}{\theta}} \theta_e e^{-1} - z c_e.$$

Furthermore, we might write this as $(c_e)^u A_e - z c_e$ for $A_e = b \rho_e t^w e^{\frac{\gamma_e}{\theta_e}} e^{-\frac{q_e b_e}{\theta_e}} e^{-\frac{s}{\theta}} \theta_e e^{-1}$, and this is maximized by

$$c_e^* = (A_e/z)^{\frac{1}{1-u}},$$

taking maximized value of

$$\frac{(A_e)^{\frac{1}{1-u}} (1-u) u^{\frac{u}{1-u}}}{z}.$$

Having found the formulas for the optimal p and c_e, we now show that there is a relative maximum for π_e at these points. To do so, we define as $D(c_e, p)$ the determinant of the Hessian matrix of π_e. After some simple computations we obtain that:

1. If we are in the free peering regime, i.e., $q_e = 0$:

   $$D(c_e, p) = c_e^{2(u-1)} d_e^2 e^{2(\frac{\gamma_e}{\theta_e} - \frac{\theta + b_e s}{\theta})} t^{2w} (1-u) u,$$

   which by assumption is always > 0, since $u + w < 1$. Furthermore, we get that

   $$\frac{\partial^2 \pi_e}{\partial c_e^2} = c_e^{u-2} d_e e^{\frac{\gamma_e}{\theta_e} - \frac{\theta + b_e s}{\theta}} t^w u \theta_e (u-1) u \theta_e,$$

2. If $q_e \neq 0$:

   $$D(c_e, p) = c_e^{2(u-1)} d_e^2 e^{2(\frac{\gamma_e - b_e q_e + theta_e}{\theta_e} - \frac{\theta + b_e s}{\theta})} t^{2w} (1-u) u,$$

   which, similarly to the free peering case, is always > 0. Additionally, it is derived that:

   $$\frac{\partial^2 \pi_e}{\partial c_e^2} = c_e^{u-2} d_e e^{\frac{\gamma_e - b_e q_e}{\theta_e} - \frac{\theta + b_e s}{\theta}} t^{2w} u \theta_e (u-1),$$

   which again is always below 0.

Thereafter, we conclude that the obtained stationary point (c_e^*,p^*), consists a relative maximum for π_e, either under the free or the paid peering situation. □

Proof: [**aCPs' optimal strategy.**] Now consider what happens under the constraint $p = 0$, and let $Z_a = d_a e^{-s b_a/\theta} (\gamma_a - b_a q_a) t^w u$.

Then each aCP's optimal level of investments will be given by:

$$c_a^* = (Z_a/z)^{\frac{1}{1-u}}.$$

To determine whether the aforementioned critical point is relative maximum, we compute the second order derivative, which is equal to:

$$b_a (Z_a/z)^{\frac{1}{1-u}(u-2)} d_a e^{-\frac{s b_a}{\theta}} (\gamma_a/b_a - q_a) t^w u (u-1).$$

Clearly, when the ISP interconnects with the aCPs via free peering contracts, i.e., $q_a = 0$, the obtained c_a^* maximizes r_a, since $u - 1 < 0$. Subsequently, we obtain the optimal q_a determined by the ISP, and based on its formula, we derive that the aforementioned critical point maximizes the profits of an aCP in the paid peering regime, as well. □

Proof: [**ISP's optimal strategy**] After all eCPs have determined their optimal startegies, the ISP's problem is to maximize its profits given by:

$$-kt + n_e \rho_e (s + q_e).$$

So q_e will be chosen to maximize

$$\rho_e (c_e)^u e^{-p/\theta_e} (s + q_e) = (A_e u)^{\frac{u}{1-u}} e^{-p/\theta_e} (s + q_e).$$

By substituting for p and ignoring multiplicative terms not depending on q_a we have:

$$\begin{aligned} &(c_e)^u e^{-\frac{p}{\theta_e}} (s + q_e) \\ &\quad \propto \left(e^{-\frac{b_e s}{\theta_e}} e^{-\frac{b_e (s + q_e)}{\theta_e}} \right)^{\frac{u}{1-u}} e^{-\frac{b_e s}{\theta_e}} e^{-\frac{a + b_e q_e}{\theta_e}} (s + q_e) \\ &\quad \propto e^{-\frac{b_e (s + q_e)}{\theta_e (1-u)}} (s + q_e). \end{aligned}$$

Thus the maximizing q_e should be chosen so that

$$b_e (s + q_e) = \theta_e (1 - u),$$

and hence the optimal q_e will be given by

$$q_e^* = \frac{\theta_e (1-u) - b_e s}{b_e}.$$

After each aCP has maximized with respect to its investment c_a, the ISP will be faced with the problem of choosing paid peering prices q_a to maximize:

$$-kt + n_a (b_a d_a (c_a)^u t^w e^{-s b_a/\theta}) (s + q_a).$$

This is stationary where

$$\frac{1}{u} \frac{1}{b_a (s + q_a)} - \frac{1}{1-u} \frac{1}{\gamma_a - b_a q_a} = 0$$

and so

$$q_a^* = \frac{\gamma_a (1-u) - b_a s u}{b_a}.$$

Based on the formula of the optimal q_a, it can be easily shown that the obtained formula for the stationary point c_a, maximizes each aCP's profits under paid peering.

If the ISP is not allowed to charge CPs for their transaction with its customer base, its optimal level of investments will be

$$t^* = \left(\frac{s w (B_{fp}^e + B_{fp}^a)}{k (1-u) e^{\frac{\theta + s(b_e + b_a)}{\theta}}} \right)^{\frac{1-u}{1-u-w}}, \tag{20}$$

with

$$B_{fp}^e = n_e \rho_e e^{\frac{s b_a}{\theta} + \frac{\gamma_e}{\theta_e}} \left(\frac{d_e \theta_e u}{z e^{\frac{\theta + s b_e}{\theta} - \frac{\gamma_e}{\theta_e}}} \right)^{\frac{u}{1-u}},$$

while

$$B^a_{fp} = n_a \rho_a e^{\frac{\theta + s b_e}{\theta}} \left(\frac{\gamma_a d_a u}{z e^{\frac{\theta + s b_a}{\theta}}} \right)^{\frac{u}{1-u}}.$$

We also find that:

$$\frac{\partial^2 \pi_{isp}}{\partial t^2} = e^{\frac{s(b_e - b_a)+\theta}{\theta}} s t^{w-2} w(u + w - 1)(\lambda_a + \lambda_e),$$

with:

$$\lambda_a = \frac{n_a b_a d_a e^{1+\frac{s b_e}{\theta}} \left(\frac{z e^{\frac{s b_a}{\theta}}}{t^w \gamma_a d_a} \right)^{\frac{u}{u-1}}}{(u-1)^2}$$

and

$$\lambda_e = \frac{n_e b_e d_e e^{\frac{\gamma_e}{\theta_e} + \frac{s b_e}{\theta}} \left(\frac{z e^{1+\frac{s b_e}{\theta} - \frac{\gamma_e}{\theta_e}}}{t^w u \theta_e d_e} \right)^{\frac{u}{u-1}}}{(u-1)^2}.$$

It turns out that $\partial^2 \pi_{isp}/\partial t^2$ is always below 0, since $u + w < 1$. Thus, the obtained stationary t, maximizes π_{isp}, under free peering.

In the paid peering regime, the optimal level of investments for the ISP, after having fixed the payments from the eCPs and aCPs, will be given by:

$$t^* = \Big(\frac{w(B^e_{pp} + B^a_{pp})}{k} \Big)^{\frac{1-u}{1-u-w}}, \tag{21}$$

where

$$B^e_{pp} = n_e d_e \theta_e e^{-2+u+\frac{\gamma_e + b_e}{\theta_e} + \frac{b_a}{\theta}} (u/z)^{\frac{u}{1-u}}$$

and

$$B^a_{pp} = n_a d_a e^{-s b_a/\theta} (u^2/z)^{\frac{u}{1-u}}.$$

By examining the eigenvalues of the 3×3 Hessian Matrix corresponding to r_{isp}, for the obtained values of t, q_e, and q_a, we find that it is always negative definite, and hence r_{isp} attains local maximum at (t^*, q_e^*, q_a^*).

Moreover, it can be easily shown (by dividing (21) with (20)), that the ratio of investments in the paid v.s. free peering does not depend on the populations of eCPs and aCPs. □

Costas A Courcoubetis was born in Athens, Greece and received his Diploma (1977) from the National Technical University of Athens, Greece, in Electrical and Mechanical Engineering, his M.S. (1980) and Ph.D. (1982) from the University of California, Berkeley, in Electrical Engineering and Computer Science. He was MTS at the Mathematics Research Center, Bell Laboratories, Professor in the Computer Science Department at the University of Crete, Professor in the Department of Informatics at the Athens University of Economics and Business, and since 2013 Professor in ESD Pillar, SUTD. His current research interests are economics and performance analysis of networks and internet technologies with applications in the development of pricing schemes that reduce congestion and enhance stability and robustness, regulation policy, smart grids and energy systems, resource sharing and auctions. Besides leading a large number of research projects in these areas he has also published over 100 papers in scientific journals such as Operations Research, Mathematics of Operations Research, Journal on Applied Probability, ToN, IEEE Transactions in Communications, IEEE JSAC, SIAM Journal on Computing, etc. and in conferences such as FOCS, STOC, LICS, INFOCOM. GLOBCOM, ITC, ACM SIGMETRICS. He is co-author with Richard Weber of "Pricing Communication Networks: Economics, Technology and Modeling" (Wiley, 2003).

Kostas Sdrolias is a student member of IEEE Communications Society and a Ph.D. candidate in the Department of Computer Science at Athens University of Economics and Business. His research interests include networks economics, game theory and mechanism design. His scientific research has been published in major conferences and journals such as IEEE ICC and ACM Transactions on Internet Technology.

Richard Weber is Churchill Professor of Mathematics for Operational Research in the Department of Pure Mathematics and Mathematical Statistics at the University of Cambridge. His research interests range over communications and operations management, control of queues, stochastic networks, on-line bin packing, queueing theory, ergodicity and stability of Markov processes, optimal search, stochastic scheduling, Gittins index, dynamic resource allocation, search games, algorithmic mechanism design, game theory and microeconomics. He has written numerous academic papers and two books: Pricing Communication Networks: Economics, Technology and Modelling, 2003, with Costas Courcoubetis, and Multi-armed Bandit Allocation Indices, 2nd edition, 2011, with John Gittins and Kevin Glazebrook.