

Capacity Expansion of Neutral ISPs via Content Peering Charges: The Bargaining Edge

Anand Kalvit, Gaurav Kasbekar, D. Manjunath and Jayakrishnan Nair
IIT Bombay

ABSTRACT

Many internet service providers (ISPs) operate under network neutrality regulations which forbid smart data pricing schemes such as those that provide differential QoS or differential pricing, leading to lower profitability. Increasing bandwidth-hungry content is making the consumers demand improved ISP infrastructure. With the risk of poor consumer experience squarely on the ISP, the ISPs are forced to invest in their infrastructure with little scope for monetisation via innovative user pricing. And they are asking the content providers (CPs) to pick up some of the tab for ISP capacity expansion. In this paper we explore the possibility of network neutral capacity expansion sponsored by voluntary peering charges from CPs.

We consider the scenario where CPs peer with an ISP and take the lead in paying peering charges with the caveat that this has to be used for capacity expansion. Since ISP capacity expansion can benefit all the CPs, and possibly even the ISP, selfish CPs will determine their charges strategically. We consider three models for the CPs to interact in determining the charge—a cooperative model, a non-cooperative model, and a bargaining model. Our analysis reveals a rather surprising result. We show that the bargaining model leads to a higher investment in the ISP infrastructure than even the cooperative model. This leads us to recommend policies that promote transparency in the interconnection agreements between CPs and ISPs.

CCS CONCEPTS

• **Social and professional topics** → **Net neutrality**; • **Networks** → *Network performance analysis*; *Public Internet*;

KEYWORDS

network neutrality, paid peering, internet economics, interconnection markets

ACM Reference format:

Anand Kalvit, Gaurav Kasbekar, D. Manjunath and Jayakrishnan Nair. 2017. Capacity Expansion of Neutral ISPs via Content Peering Charges: The Bargaining Edge. In *Proceedings of 11th EAI International Conference on Performance Evaluation Methodologies and Tools, Venice, Italy, December 5–7, 2017 (VALUETOOLS 2017)*, 8 pages.
<https://doi.org/10.1145/3150928.3150931>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

VALUETOOLS 2017, December 5–7, 2017, Venice, Italy

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-6346-4/17/12...\$15.00

<https://doi.org/10.1145/3150928.3150931>

1 INTRODUCTION

The public discourse on net neutrality—arguments for and against differential pricing and differential QoS—has been vigorous for quite some time now and many countries have adopted strong regulations favouring network neutrality. Interconnection agreements, especially between the access ISPs and content provider (CP) networks like those of YouTube, Netflix, and Facebook, or content delivery networks (CDNs) like Akamai and Lightstream can also produce effects similar to neutrality violation. However, not much attention has been paid to the interconnection markets and they have largely remained unregulated.

That interconnection agreements would determine the structure of the internet was recognized in [1] in the prescient observation that “interconnection agreements do not just route traffic in the internet, they also route money.” It was further recommended in [1] that any policy on interconnections should focus on transparency into the workings of the interconnection markets.

In the present paper, we present economic models for the interconnection market between an access ISP and content provider (CP) networks operating in a net neutrality regime. The objective is to provide insights into the interconnection market structure and guidelines for policies that are socially beneficial in the long term. We show that transparency can actually lead to significantly improved infrastructure investments as compared to an unregulated regime.

1.1 Background

The internet has traditionally been a hierarchical network with Tier 1 ISPs at the top and access ISPs and their subscribers (users) at the bottom of the hierarchy. This hierarchical topology started to break down with several networks peering around Tier 1 ISPs and resulting in ‘donut peering’ structures [2]; many of these are paid peering arrangements (as opposed to settlement-free peering). With a small number of content providers dominating internet traffic,¹ such donut peering structures are also being pursued by CPs. For example, the Netflix Open Connect² is a platform for Netflix to directly connect with an access ISP. Google maintains multiple points of presence (POP) in most markets and access ISPs may connect directly to them.

Direct interconnections benefit both the CPs and the end users. The CPs save on transit costs and the users see improved quality of experience (QoE). The improved QoE in turn benefits the CPs with increased customer stickiness. The direct interconnection though

¹According to the Sandvine report (see www.sandvine.com) NetFlix (55.5%), YouTube (17.5%) and seven other streaming applications together contribute about 80% of download traffic on wireline networks. For mobile networks, the top eight CPs contribute more than 80% of the downloads.

²<https://media.netflix.com/en/company-blog/how-netflix-works-with-isps-around-the-globe-to-deliver-a-great-viewing-experience>

imposes a cost to the ISP because the improved QoE can drive up the demand. Hence, as has been pointed out in [3], since the risk of poor user experience is with the ISP, the ISP needs to improve its infrastructure through additional investment.

To recoup such an investment, the ISP can extract the system surplus from the user-side and/or from the CP side. User-side surplus may be extracted through pricing innovations; see [4] for several examples. However, they face several obstacles. Firstly, market expectations have been shaped by flat fee regimes and a simple pricing structure is the *de rigueur* in many markets. Secondly, in many markets, the pricing regimes can be close to marginal cost because of inter-ISP competition. Finally, regulatory requirements like net neutrality stipulations severely limit the ability of the ISP to extract user-side surplus. There appears to be more manoeuvring room for the ISP to extract CP-side surplus. Also, as has been described in the analysis in [5], with increasing asymmetry in the ISP and CP revenues, the ISPs demanding the CPs to contribute to the costs of ISP infrastructure is natural and peering by CPs and/or groups of CPs (via a hosting service provider) is natural. Further, since different CPs benefit differently from improved user QoE, it is also natural that the ISPs charge the CPs in an asymmetric manner. This seems to be happening to some extent with the paid-peering arrangement between Netflix and Comcast being the celebrated case. There is also emerging research interest in characterising what can happen here.

The net neutrality regulations also affect CP-side surplus extraction, but services such as caching and interconnection, or paid-peering, are as yet unregulated. To be sure, there have been some attempts at imposing harsh regulations, e.g., Cogent and Netflix requested regulatory intervention to prohibit Comcast from collecting peering charges [3]. To the best of our knowledge, the public discourse on interconnection regulation has not yet happened, but several governments are exploring this matter. This paper develops economic models to inform such regulations.

1.2 Previous Work

There is substantial literature on the economics of non-neutral networks. The effect of discriminatory QoS on various performance parameters like social surplus, surplus of users, CPs, and ISPs, and the incentive of the ISP to invest in its infrastructure has been studied in, among others, [6–8]. Analysis of the effect of discriminatory pricing is being increasingly studied, e.g., [9–11]. Coalitional game theory was used in [5, 12] to analyse the fair sharing of surplus among the access ISPs and CPs; these papers predict that prevalent settlement arrangements are not stable. Early work on paid peering, e.g., [13–15], implicitly assumed that all the networks were ‘similar’. More recently, [16] suggested a Nash-peering model based on the value peering.

Our interest in this paper is more along the lines of the work of [17–19]. In [17], a Nash bargaining model is used to determine peering prices when there is a churn in the system. In [18], using a transactional model for demand, a Stackelberg game with the ISP as the leader and the CPs as the followers is formulated to study peering prices. In [19], a choice model is used to determine the value of direct peering and the peering dynamics are analyzed.

1.3 Preview

Traditionally, paid peering has implied that the ISP takes the lead in setting the peering charges. This can lead to a potentially non-neutral arrangement between the ISP and the CP. In this paper we propose a neutral network with CPs taking the lead in contributing to the ISP infrastructure via voluntary peering payments. Indeed, they do so in their own self interest. This is not unreasonable because, as we saw earlier, only a small number of CPs dominate internet traffic and they are the ones that would be willing to pay the peering charges. A further motivation for this model is the belief that many CPs wield significant power over the access ISPs, especially in markets where the latter do not enjoy a monopoly.

Our objective here is to compare the ISP investment outcome from the different types of interactions between the CPs. In the next section, we describe the model in detail. In Section 3, we consider three mechanisms that the CPs use to determine the peering charge. First, we consider a cooperative model in which the CPs form a single coalition that seeks to maximise the total CP surplus. Next we consider the non-cooperative setup where the CPs strategically decide the peering charge. Finally, we consider a Nash bargaining problem between the CPs to determine the ISP investment. Our key finding in this section is that the bargaining based peering charges leads to a higher investment than even the cooperative model where the CPs maximize their net profit. A second important finding is that in the non-cooperative setup, in most cases, only one of the two CPs will contribute to capacity expansion and the other will free-ride.

In Section 4, we illustrate the results with numerical examples and characterize the differences in the outcome from the three models. We also provide some policy guidelines for the interconnection market. We conclude with a discussion on extensions and future work in Section 5.

2 MODEL AND PRELIMINARIES

We consider a system with a single access ISP and two CPs, labeled 1 and 2, serving a user population. The CPs may peer with the ISP by paying a voluntary q -charge that is used for capacity expansion by the ISP, thus increasing the quality of service (QoS) that the users see and hence their consumption. We will initially assume that the q -charge to be paid by the CPs to the ISP for improving the QoS is a fixed charge Q_i for CP i . We will also briefly discuss the case when there is a volume-based charge at rate q_i for CP i in Section 5.

The consumption of the users from the two CPs depends on the inherent interest in the corresponding content and also on the quality of service that the network provides. This QoS in turn depends on the investment made by the ISP towards its infrastructure. Let μ_0 be the baseline investment without the CP peering and let μ be the additional investment enabled by the asymmetric q -charge on the CP side. The increased QoS is seen by all the users and, since the ISP is neutral, it does not control the effect of this increase on consumption of any specific CP.

The effect of interest and QoS on the consumption by a user of content of CP i is given by $x_{i0} + x_i(\mu)$ where x_{i0} is the ‘baseline’ consumption, i.e., the consumption without additional investment afforded by q -charges on the CP, and $x_i(\mu)$ is the extra consumption

that is enabled by the additional investment by the ISP. We make the reasonable assumption that $x_i(\mu)$ is increasing and concave in μ with $x_i(0) = 0$ and $\lim_{\mu \rightarrow \infty} x'(\mu) = 0$. We assume that the profitability of a CP is linear in the consumption, i.e., CP i has a revenue of v_i per unit of traffic.

The q -charges that are paid by the CPs are used by the ISP towards enhancing its infrastructure; thus μ is a function of Q_1 and Q_2 . This in turn increases consumption which can be profitable to both the CPs and also to the ISP. Thus the surplus of each CP will be a function of both Q_1 and Q_2 ; denote this by $f_i(Q_1, Q_2)$. For the flat charge case, we can now write

$$f_i(Q_1, Q_2) = v_i(x_{i0} + x_i(\mu(Q_1, Q_2))) - Q_i \quad (1)$$

We will make the reasonable assumptions that $\mu(0, 0) = 0$. Also note that $f_i(0, 0) = v_i x_{i0}$. Without loss of generality, we assume $f_i(0, 0) = 0$ since our interest is in the analysis of the incremental benefits of paid peering by the CPs with a neutral ISP.

We consider two models for the user side charges. There is either a fixed charge p per user with unlimited usage, or there is a volume based charge of p per unit consumption.

We will assume a complete information game, i.e., the v_i , $x_i(\mu)$ and $\mu(Q_1, Q_2)$ are known to both CPs. We will also assume that the user charge of p is exogenously determined and is not part of the strategy space of either the ISP or the CPs. Our interest is to analyze the incentives for CPs to pay for ISP capacity expansion. As we see above, this is determined by the Q_i and, when applicable, by the extra revenue earned from the increased usage.

We will consider the following kinds of interactions between the CPs and the ISP.

- We first consider a cooperative game between the two CPs with the objective of maximizing the total revenue, determining μ cooperatively, and distributing the costs using a Shapley value mechanism. This is an ‘ideal’ objective and is useful for purposes of comparison with more ‘realistic’ schemes.
- Next we consider the case when the CPs lead by non-cooperatively determining (Q_1, Q_2) .
- The third interaction model that we consider is for the CPs to determine the (Q_1, Q_2) using a bargaining framework.

In all of these cases, we will assume that $\mu \propto Q_1 + Q_2$.

Finally, we describe the user consumption model. As we said earlier, our objective here is to model the effect of increased ISP investment on the user consumption and hence the benefits to the CPs when the ISP is neutral. To keep the model parsimonious, and tractable, we assume the ISP investment to be a proxy for user QoS and assume the following model for $x_i(\mu)$.

$$x_i(\mu) = a_i \phi(\mu) \quad (2)$$

where $\phi(\cdot)$ is an increasing concave function with $\phi(0) = 0$ and $\lim_{t \rightarrow \infty} \phi'(t) = 0$. The parameter a_i indicates the user preference towards CP i 's content.

Our interest is to analyze the following quantities under the different interaction models described above: the ISP investment μ and the CP surplus f_i . We will use the superscript C to indicate these quantities under the cooperative framework, B for the bargaining framework, and N for the non-cooperative framework. For example

μ^C would be the ISP investment with cooperation, and f_i^B the surplus of CP i when they bargain on the q -charge.

3 FLAT PEERING AND USER CHARGES: ANALYTICAL RESULTS

In this section, we consider the case where the CPs pay the ISP a flat peering charge towards capacity expansion, and users pay a flat fee (not dependent on data usage) to the ISP for internet access. We study an idealised cooperative peering model between CPs, as well as more practical models based on the Nash equilibrium and the Nash bargaining framework. The cooperative model corresponds to the scenario where the CPs form a single coalition seeking to maximize its aggregate profit. At the other extreme, the Nash equilibrium based model captures a non-cooperative setting where each CP seeks to maximize its own revenue, in the absence of any communication between the CPs. Finally, the bargaining framework corresponds to a scenario where the CPs can communicate and ‘bargain’ to arrive at an agreement on the peering charges paid by each CP. The above models enable us to analyse the impact of non-cooperation as well as bargaining between the CPs.

Our analysis reveals that from the standpoint of capacity expansion, bargaining based peering arrangements are more efficient than the cooperative setting, which in turn is more efficient than the non-cooperative setting. In other words, the strategic interaction that enables a bargaining solution between the CPs is actually beneficial to the user base, since it leads to the maximum capacity expansion. Interestingly, except when the CPs are perfectly symmetric, any Nash equilibrium has only one CP contributing towards capacity expansion. This is undesirable not just because it leads to lower capacity expansion, but also because it can result in a push towards (non-neutral) preferential treatment to traffic of the ‘sponsoring’ CP.

From a regulatory standpoint, our analysis suggests that it is socially beneficial to have a transparent platform for CPs to commit on their contributions towards network neutral capacity expansion. Indeed, this might be preferable to the present practice where CPs enter into bilateral (and seemingly non-neutral [20]) peering arrangements with ISPs, the terms of which are kept private.

Throughout this section, we assume that the capacity expansion μ is proportional to the contributions of the CPs, i.e., $\mu = \gamma(Q_1 + Q_2)$. Without loss of generality, we set $\gamma = 1$. In this case, the CP surplus functions are given by

$$f_i = v_i x_i(Q_1 + Q_2) - Q_i = \theta_i \phi(Q_1 + Q_2) - Q_i \quad (i \in \{1, 2\}).^3$$

where $\theta_i := a_i v_i$.

We begin our analyses by first considering the cooperative regime.
Cooperative peering

We now consider the (idealised) setting where the CPs act as a single coalition seeking to maximize its net surplus, given by

$$v_1 x_1(\mu) + v_2 x_2(\mu) - \mu.$$

³Under the flat pricing model, it is not meaningful to consider the ISP as being strategic, i.e., to not couple μ to $Q_1 + Q_2$ but instead treat μ as a strategic decision of the ISP in response to (Q_1, Q_2) . Indeed, under such a model, since the ISP surplus is not tied to the capacity expansion (q -charges as well as user payments being flat and not volume based), the ISP has no incentive to increase capacity. If the q -charges are volume-based (we discuss this case briefly in Section 5), it is indeed meaningful to consider models with a strategic ISP.

The first term above is the revenue from the increase in user data consumption due to capacity expansion μ , and the second is the q -charge paid by the coalition. Note that this case is an ‘idealisation’ in the sense that we disregard the strategic interactions between the CPs. Thus, it is natural to benchmark the capacity expansion under this model to that under more practical settings where the CPs selfishly seek to maximize their own surplus.

The concavity assumptions on the x_i lead to the following elementary characterization of the cooperative capacity expansion, denoted by μ^C .

LEMMA 3.1.

$$\mu^C = \begin{cases} 0 & \text{if } \theta_1 + \theta_2 \leq \frac{1}{\phi'(0)} \\ (\phi')^{-1}\left(\frac{1}{\theta_1 + \theta_2}\right) & \text{otherwise} \end{cases}.$$

Note that if $\theta_1 + \theta_2 \leq \frac{1}{\phi'(0)}$, then the CP coalition does not have an incentive to invest in capacity expansion, since the resultant usage increase does not generate sufficient revenue.

While the above cooperative model does prescribe the capacity expansion uniquely, it does not prescribe the q -charges paid by the respective CPs. Cooperative game theory provides several solution concepts to capture this settlement between the CPs [21]. However, since our focus is primarily on the capacity expansion corresponding to the cooperative regime (which forms the benchmark against which we compare the capacity expansion under the non-cooperative and bargaining models), we do not address the issue of cooperative settlement in this paper.

Non-cooperative peering

Next, we consider the non-cooperative setting, where CPs decide on their contributions towards capacity expansion selfishly and without coordination. In this case, it is then natural to consider Nash equilibria between the CPs, with Q_i being the ‘action’ of CP i . Note that $Q^N = (Q_1^N, Q_2^N)$ is a (pure) Nash equilibrium if

$$\begin{aligned} Q_1^N &\in \arg \max_{Q_1} f_1(Q_1, Q_2^N), \\ Q_2^N &\in \arg \max_{Q_2} f_2(Q_1^N, Q_2). \end{aligned}$$

A key issue with the non-cooperative setting is the so-called ‘tragedy of commons’, where the CPs under-invest in capacity expansion (often to their own disadvantage) by acting selfishly. Indeed, as we will see, except when the CPs are perfectly symmetric, a Nash equilibrium involves at most one of the CPs making a positive contribution towards capacity expansion. To formalise this, we define the following types of Nash equilibria. A Nash equilibrium (Q_1^N, Q_2^N) is said to be

- Type 0 if none of the CPs peer, i.e., $Q_1^N = Q_2^N = 0$,
- Type 1 if only one of the CPs peers, i.e., $Q_1^N = 0$ and $Q_2^N > 0$, or $Q_2^N = 0$ and $Q_1^N > 0$,
- Type 2 if both CPs peer, i.e., $Q_1^N, Q_2^N > 0$.

The following theorem characterizes the conditions for existence of different types of Nash equilibria.

THEOREM 3.2. *A (pure) Nash equilibrium always exists.*

- (1) If $\max(\theta_1, \theta_2) \leq \frac{1}{\phi'(0)}$, then the only Nash equilibrium is Type 0, i.e., $(0, 0)$.

- (2) If $\max(\theta_1, \theta_2) > \frac{1}{\phi'(0)}$ and $\theta_1 > \theta_2$, then the only Nash equilibrium is Type 1 of the form $(Q_1^N, 0)$, where Q_1^N is the unique solution of $\phi'(Q_1) = \frac{1}{\theta_1}$.

- (3) If $\max(\theta_1, \theta_2) > \frac{1}{\phi'(0)}$ and $\theta_2 > \theta_1$, then the only Nash equilibrium is Type 1 of the form $(0, Q_2^N)$, where Q_2^N is the unique solution of $\phi'(Q_2) = \frac{1}{\theta_2}$.

- (4) If $\max(\theta_1, \theta_2) > \frac{1}{\phi'(0)}$ and $\theta_2 = \theta_1$, then there is a continuum of Type 1 / 2 Nash equilibria (Q_1^N, Q_2^N) , satisfying

$$Q_1^N + Q_2^N = \mu^N, \quad Q_1^N, Q_2^N \geq 0,$$

where μ^N is the unique solution of $\phi'(\mu) = \frac{1}{\theta_1} = \frac{1}{\theta_2}$.

Theorem 3.2 partitions the $\theta_1 \times \theta_2$ space into four regions, and provides a precise characterization of the Nash equilibria in each region. Figure 1 provides a pictorial depiction of these regions. Interpreting θ_i to be the size of CP i , we see that when the sizes of both CPs are small (precisely, less than or equal to $\frac{1}{\phi'(0)}$), the only Nash equilibrium is $(0, 0)$. Else, except when the sizes are exactly matched, the unique Nash equilibrium is Type 1, with the ‘larger’ CP being the only contributor towards capacity expansion. This means that in general, non-cooperative peering results in highly asymmetric contributions by the CPs.

A direct corollary of Theorem 3.2 is the following.

COROLLARY 3.3. *The capacity expansion μ^N under any Nash equilibrium between the CPs is unique.*

PROOF OF THEOREM 3.2. If $\max(\theta_1, \theta_2) \leq 1/\phi'(0)$ (Condition (1)), then for $i = 1, 2$, $\phi'(0) \leq 1/\theta_i$. Thus, the dominant response of CP i to any action of the other CP is to set $Q_i = 0$. It therefore follows that Type 0 is the unique equilibrium for this case.

Next, if $\max(\theta_1, \theta_2) > 1/\phi'(0)$ and $\theta_1 > \theta_2$, (Condition (2)), it is easy to check that $(0, 0)$ is not a Nash equilibrium, and that $(Q_1^N, 0)$ is a Nash equilibrium if and only if $\phi'(Q_1) = \frac{1}{\theta_1}$. Moreover, if we assume that (\hat{Q}_1, \hat{Q}_2) is an equilibrium with $\hat{Q}_2 > 0$, we get the conditions

$$\phi'(\hat{Q}_1 + \hat{Q}_2) = \frac{1}{\theta_2}, \quad \phi'(\hat{Q}_1 + \hat{Q}_2) \leq \frac{1}{\theta_1},$$

which are clearly inconsistent.

The conclusions under Condition (3) are proved using an identical argument.

Finally, if $\max(\theta_1, \theta_2) > 1/\phi'(0)$ and $\theta_2 = \theta_1$ (Condition(4)), then $\phi'(0) > 1/\theta_i$. It is easy then to check that $(0, 0)$ is not a Nash equilibrium. Moreover $(\hat{Q}_1, \hat{Q}_2) \neq (0, 0)$ is a Nash equilibrium if and only if

$$\phi'(\hat{Q}_1 + \hat{Q}_2) = \frac{1}{\theta_1} = \frac{1}{\theta_2}.$$

Thus, we have a continuum of Type 1 / 2 Nash equilibria as claimed. \square

Our next result states that non-cooperative peering always leads to a lower capacity expansion compared to the cooperative model.

THEOREM 3.4. *If $\theta_1 + \theta_2 \leq \frac{1}{\phi'(0)}$, then $\mu^N = \mu^C = 0$. If $\theta_1 + \theta_2 > \frac{1}{\phi'(0)}$, then $\mu^N < \mu^C$.*

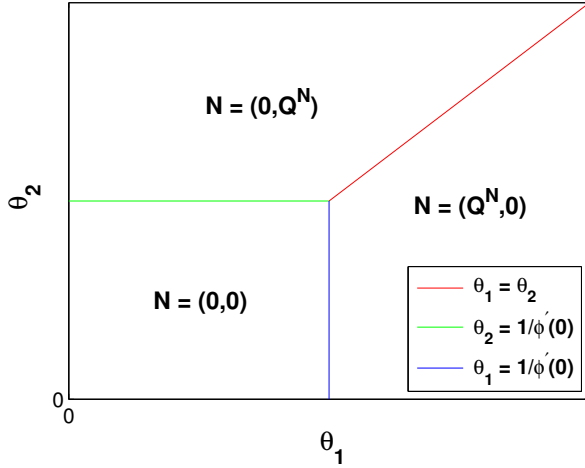


Figure 1: Conditions for existence of different types of Nash equilibria. Note the partition of the $\theta_1 \times \theta_2$ space into 3 regions. There exists a unique Nash equilibrium at each point in the space except for those lying on the infinite ray $\theta_1 = \theta_2 > 1/\phi'(0)$. At each point along this infinite ray, there exists a continuum of Nash equilibria, all with the same μ^N .

PROOF. Recall from Theorem (3.2) that the conditions for existence of different types of Nash equilibria partition the $\theta_1 \times \theta_2$ space into four non-overlapping subsets (see Figure 1). It is therefore sufficient to prove the statement of the theorem for each subset.

For the subset corresponding to Condition (1), this is elementary. For the subsets corresponding to Conditions 2, 3, and 4, note that μ^N satisfies

$$\phi'(\mu^N) = \min\left(\frac{1}{\theta_1}, \frac{1}{\theta_2}\right).$$

On the other hand, μ^C satisfies

$$\phi'(\mu^C) = \frac{1}{\theta_1 + \theta_2} < \min\left(\frac{1}{\theta_1}, \frac{1}{\theta_2}\right).$$

It now follows that $\mu^N < \mu^C$ given the concavity assumptions on $\phi(\cdot)$. \square

Finally, we consider the price of anarchy (POA), defined as

$$\text{POA} = \frac{\mu^C}{\mu^N}.$$

The POA captures the inefficiency due to the non-cooperation between the CPs with respect to capacity expansion. Clearly, the POA is well defined when $\theta_1 + \theta_2 > \frac{1}{\phi'(0)}$ (otherwise, we have $\mu^C = \mu^N = 0$). Moreover, if $\theta_1 + \theta_2 > \frac{1}{\phi'(0)}$, Theorem 3.4 implies that $\text{POA} > 1$. For an explicit characterization of the value of the POA, we consider specific examples of usage functions. For the logarithmic usage function $x_i = \theta_i \log(1 + b\mu)$ (where $b > 0$), it can be shown that (assuming $\theta_1 + \theta_2 > \frac{1}{\phi'(0)} = \frac{1}{b}$)

$$\text{POA} = \begin{cases} \infty & \max(\theta_1, \theta_2) \leq \frac{1}{b} \\ \frac{b(\theta_1 + \theta_2) - 1}{b \max(\theta_1, \theta_2) - 1} & \max(\theta_1, \theta_2) > \frac{1}{b} \end{cases}.$$

For the bounded exponential usage function $x_i = \theta_i(1 - e^{-b\mu})$ (where $b > 0$), it can be similarly shown that (assuming again that $\theta_1 + \theta_2 > \frac{1}{\phi'(0)} = \frac{1}{b}$)

$$\text{POA} = \begin{cases} \infty & \max(\theta_1, \theta_2) \leq \frac{1}{b} \\ \frac{\ln(b(\theta_1 + \theta_2))}{\ln(b \max(\theta_1, \theta_2))} & \max(\theta_1, \theta_2) > \frac{1}{b} \end{cases}.$$

In both the above cases, note that even within the region $\max(\theta_1, \theta_2) > \frac{1}{b}$, the POA is unbounded. Interestingly, the POA approaches 1 when $\max(\theta_1, \theta_2) \gg \min(\theta_1, \theta_2)$, i.e., when the CP sizes are highly asymmetric.

Bargaining based peering

Finally, we consider the setting where the CPs ‘bargain’ to arrive at an agreement on their peering payments. Note that this requires that the CPs are able to communicate with one another. We invoke the classical Nash bargaining solution from the bargaining literature to capture the agreement between the CPs. Our main result is that the bargaining solution is even more efficient than the cooperative regime with respect to capacity expansion. As discussed before, this has significant implications from a regulatory standpoint.

To define the Nash bargaining solution (NBS), we first define the set of feasible, non-negative surplus pairs:

$$\mathcal{F} := \{(f_1(Q_1), f_2(Q_2)) \mid Q_1, Q_2 \geq 0\} \cap \mathbb{R}_+^2.$$

A Nash bargaining solution (NBS) $f^B = (f_1^B, f_2^B)$ is defined to be a solution of the following maximization.

$$\begin{aligned} \max \quad & \hat{f}_1 \hat{f}_2 \\ \text{such that} \quad & (\hat{f}_1, \hat{f}_2) \in \mathcal{F} \end{aligned} \quad (3)$$

Note that we are maximizing the product of the CP surpluses, subject to the constraint that each surplus is non-negative. It is important to note that the axiomatic development of the Nash bargaining formulation assumes that the payoff space, which is the set of all payoff pairs of both players, is convex. This in turn implies that the NBS is unique. In the present setting, we are unable to prove the convexity of the set \mathcal{F} (although numerical experiments suggest that the set is indeed convex). However, we prove via direct arguments that the optimization (3) has a unique maximizer (see Lemma 3.6).

The Nash bargaining framework involves a *disagreement outcome*, which is the vector of payoff pairs if the two parties fail to arrive at an agreement. In the present setting, it is natural to take the *disagreement outcome* to be $(0, 0)$, which corresponds to the CPs not peering. Note that under the Nash bargaining framework, the CPs arrive at an agreement if and only if the optimization (3) has a optimal value that is strictly positive.

Our first result characterizes the condition for a bargaining agreement between the CPs.

LEMMA 3.5. *The CPs arrive at a bargaining agreement if and only if $\theta_1 + \theta_2 > \frac{1}{\phi'(0)}$.*

PROOF. Suppose that $\theta_1 + \theta_2 > \frac{1}{\phi'(0)}$. Then there exists $\bar{\mu} > 0$ such that

$$v_1 x_1(\bar{\mu}) + v_2 x_2(\bar{\mu}) - \bar{\mu} > 0.$$

It is easy to see that one can find $\bar{Q}_1, \bar{Q}_2 \geq 0$, such that

$$\begin{aligned}\bar{Q}_1 + \bar{Q}_2 &= \bar{\mu}, \\ \bar{f}_1 &:= v_1 x_1(\bar{\mu}) - \bar{Q}_1 > 0, \\ \bar{f}_2 &:= v_2 x_2(\bar{\mu}) - \bar{Q}_2 > 0.\end{aligned}$$

Since we have demonstrated a point $(\bar{f}_1, \bar{f}_2) \in \mathcal{F}$ that yields a strictly positive objective value for the optimization (3), it follows that we have a bargaining agreement.

Next, suppose that there exists a bargaining agreement. This implies that there exists $\hat{Q}_1, \hat{Q}_2 \geq 0$ such that $\hat{f}_i := f_i(\hat{Q}_1, \hat{Q}_2) > 0$ for $i = 1, 2$. This in turn implies that

$$v_1 x_1(\hat{Q}_1 + \hat{Q}_2) + v_2 x_2(\hat{Q}_1 + \hat{Q}_2) - (\hat{Q}_1 + \hat{Q}_2) > 0,$$

which proves that the capacity optimization under the cooperative regime has a positive solution. It now follows that $\theta_1 + \theta_2 > \frac{1}{\phi'(0)}$ (see Lemma 3.1). \square

Note that existence of a non-trivial NBS is equivalent to the capacity expansion under the cooperative regime being strictly positive (see Lemma 3.1). Thus, the set of system parameters over which we have a non-trivial NBS is a strict superset of the set of system parameters over which a non-trivial Nash equilibrium exists.

Our next result establishes the uniqueness of the NBS.

LEMMA 3.6. *If $\theta_1 + \theta_2 > \frac{1}{\phi'(0)}$, then the optimizer $f^B = (f_1^B, f_2^B)$ of (3) is unique.*

PROOF. Suppose, for the purpose of obtaining a contradiction, that there exist two optimizers $\hat{f} = (\hat{f}_1, \hat{f}_2)$ and $\bar{f} = (\bar{f}_1, \bar{f}_2)$ of (3). Clearly, there exists (\hat{Q}_1, \hat{Q}_2) and (\bar{Q}_1, \bar{Q}_2) such that

$$f_i(\hat{Q}_1, \hat{Q}_2) = \hat{f}_i, \quad f_i(\bar{Q}_1, \bar{Q}_2) = \bar{f}_i, \quad i = 1, 2.$$

It is easy to see that

$$\left(\frac{\hat{f}_1 + \bar{f}_1}{2} \right) \left(\frac{\hat{f}_2 + \bar{f}_2}{2} \right) > \hat{f}_1 \bar{f}_2 = \bar{f}_1 \hat{f}_2 > 0. \quad (4)$$

Consider now, for $i = 1, 2$,

$$\begin{aligned}\tilde{f}_i &:= f_i \left(\frac{\hat{Q}_1 + \bar{Q}_1}{2}, \frac{\hat{Q}_2 + \bar{Q}_2}{2} \right) \\ &= v_i x_i \left(\frac{\hat{Q}_1 + \bar{Q}_1}{2}, \frac{\hat{Q}_2 + \bar{Q}_2}{2} \right) - \frac{\hat{Q}_i + \bar{Q}_i}{2} \\ &< \frac{1}{2} v_i x_i(\hat{Q}_1 + \hat{Q}_2) + \frac{1}{2} v_i x_i(\bar{Q}_1 + \bar{Q}_2) - \frac{\hat{Q}_i + \bar{Q}_i}{2} \\ &= \frac{\hat{f}_i + \bar{f}_i}{2}.\end{aligned}$$

The bounding above invokes Jensen's inequality. It now follows, from (4), that

$$\tilde{f}_1 \tilde{f}_2 > \hat{f}_1 \bar{f}_2 = \bar{f}_1 \hat{f}_2.$$

Since $(\tilde{f}_1, \tilde{f}_2) \in \mathcal{F}$, we have a contradiction. \square

Having established the uniqueness of the NBS, the next step is to characterize the capacity expansion under the NBS. However, it turns out that the capacity expansion under the NBS $f^B = (f_1^B, f_2^B)$ is not necessarily unique. Indeed, we show that the NBS is associated with at least one and at most two values of capacity expansion.

When there is a unique capacity expansion associated with the NBS, we show that it is equal to the cooperative capacity expansion μ^C . When there are two possible values, the smaller value is less than μ^C , whereas the greater value exceeds μ^C . If we thus follow the convention that for the same surplus vector, the CPs choose the greater capacity expansion (resulting in a greater benefit to the users), we conclude that the capacity expansion under NBS actually exceeds that under the cooperative regime. This is formalised in the following theorem.

THEOREM 3.7. *If $\theta_1 + \theta_2 > \frac{1}{\phi'(0)}$, then there exists a unique $\mu^B \geq \mu^C$, $Q_1^B, Q_2^B \geq 0$, such that*

$$\begin{aligned}\mu^B &= Q_1^B + Q_2^B, \\ f_i(Q_1^B, Q_2^B) &= f_i^B, \quad i = 1, 2.\end{aligned}$$

PROOF. Let $g(\mu) := v_1 x_1(\mu) + v_2 x_2(\mu) - \mu$ denote total surplus of CPs 1 and 2. Clearly, $f_1^B + f_2^B \leq g(\mu^C)$.

Now, suppose that $f_1^B + f_2^B < g(\mu^C)$. Given convexity properties of g , it follows that there exist unique values $\underline{\mu}, \bar{\mu}$ such that $0 < \underline{\mu} < \mu^C < \bar{\mu}$, and $g(\underline{\mu}) = g(\bar{\mu}) = f_1^B + f_2^B$. Note that $\underline{\mu}$ and $\bar{\mu}$ are the only possible capacity expansions under the NBS f^B .

Now, let $(\underline{Q}_1, \underline{Q}_2)$ and (\bar{Q}_1, \bar{Q}_2) be defined as follows. For $i = 1, 2$,

$$f_i^B = v_i x_i(\underline{\mu}) - \underline{Q}_i, \quad f_i^B = v_i x_i(\bar{\mu}) - \bar{Q}_i.$$

Note that $(\underline{Q}_1, \underline{Q}_2)$ and (\bar{Q}_1, \bar{Q}_2) are the q -charges corresponding to $\underline{\mu}$ and $\bar{\mu}$ respectively, if feasible. Clearly, $(\bar{Q}_1, \bar{Q}_2) > (\underline{Q}_1, \underline{Q}_2)$. Since at least one of $\underline{\mu}$ and $\bar{\mu}$ is feasible, it follows that $(\bar{Q}_1, \bar{Q}_2) \geq (0, 0)$. This completes the proof, setting $\mu^B = \bar{\mu}$, and $(Q_1^B, Q_2^B) = (\bar{Q}_1, \bar{Q}_2)$.

The case of $f_1^B + f_2^B = g(\mu^C)$ is trivial; in this case, we have $\mu^B = \mu^C$. This completes the proof. \square

Note that while the cooperative regime does generate a higher aggregate surplus for the CPs, the bargaining framework results in a higher capacity expansion. This means that even though strategic interaction between the CPs can lower CP surplus, it is beneficial to the user base.

We conclude by defining the benefit of bargaining (BOB), which captures the relative benefit of the bargaining solution over the cooperative model with respect to capacity expansion:

$$\text{BOB} = \frac{\mu^B}{\mu^C}.$$

Note that the BOB is well defined for $\theta_1 + \theta_2 > \frac{1}{\phi'(0)}$, and is lower bounded by 1. However, a closed form characterization of the BOB is infeasible for even the simplest usage models. In Section 4 that follows, we thus resort to numerical experiments to gain additional insights on the efficiency of the bargaining based solution.

4 FLAT PEERING AND USER CHARGES: NUMERICAL RESULTS AND POLICY IMPLICATIONS

Given the analytical results in Section 3, the goal of this section is to gain additional insights on the impact of non-cooperation and

bargaining between the CPs via numerical experiments, and to summarize the policy implications of our results.

We first consider the price of anarchy (POA). Figure 2 illustrates how POA varies across the $\theta_1 \times \theta_2$ space for the logarithmic usage function. We note that the POA becomes unbounded and grows to ∞ as one approaches the unit square $[0, 1] \times [0, 1]$ from outside. Also, observe that the POA is maximum when $\theta_1 \approx \theta_2$, i.e., when the CPs are of comparable size. This is because in this case, only one of the CPs ends up contributing towards capacity expansion under any Nash equilibrium, whereas both CPs make comparable contributions under the cooperative regime. It is also worth noting that the POA decreases as the CP sizes get more asymmetric. This is to be expected, since even the cooperative regime would involve the larger CP making the dominant contribution in this case.

Next, we consider the benefit of bargaining (BOB), which captures the relative improvement in capacity expansion under bargaining-based peering compared with the cooperative setting. Figure 3 shows how the BOB varies across the $\theta_1 \times \theta_2$ space, again for the logarithmic usage functions. Note that the tessellation in the image is a result of the numerical resolution at which μ^B was computed. Moreover, we are unable to sweep all the way upto the $\theta_1 + \theta_2 = 1$ line because of the computational constraints created due to μ^B being infinitesimally small close to this line. We observe that the BOB is modest when the CPs are of comparable size, i.e., $\theta_1 \approx \theta_2$. On the other hand, the benefit of bargaining grows as the CP sizes are highly asymmetric. Interestingly, this implies that the capacity expansion under bargaining-based peering is always substantial relative to that under non-cooperative peering: When $\theta_1 \approx \theta_2$, we have a large POA but $\text{BOB} \approx 1$, whereas when $\max(\theta_1, \theta_2) \gg \min(\theta_1, \theta_2)$, we have $\text{POA} \approx 1$ but a large BOB.

Finally, we compare the aggregate CP surplus under the bargaining and non-cooperative models. Figure 4 shows the difference between the total revenues generated at NBS and NE. We observe an interesting dichotomy here. Surprisingly, when the CP sizes are highly asymmetric, it turns out that the aggregate CP surplus is actually greater under non-cooperative peering than under bargaining-based peering. In other words, compared to the non-cooperative setting, the configuration resulting from the bargaining framework leaves the CPs worse off, and the user base better off. On the other hand, when the CP sizes are comparable, the bargaining framework generates a higher surplus for the CPs than the non-cooperative setting.

Policy Implications

In this paper, we propose network neutral capacity expansion funded by voluntary peering charges paid by CPs. The results of Sections 3 and the present section show that this is indeed feasible. Moreover, the observation that bargaining based peering results in the highest capacity expansion suggests that policy makers should set up a transparent platform for CPs to make commitments for internet infrastructure expansion. In contrast, the present practice of confidential and bilateral peering arrangements between CPs and ISPs leads not only to potentially lower infrastructure investments, but also to non-neutral internet access for users.

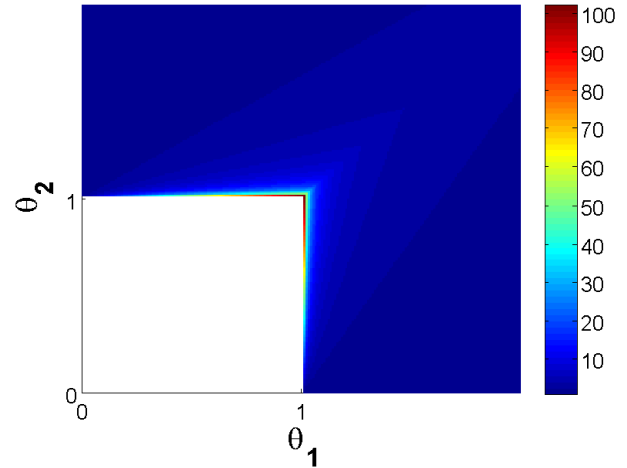


Figure 2: Price of anarchy (POA) over the $\theta_1 \times \theta_2$ space under the model $x_i = a_i \log(1 + \mu)$

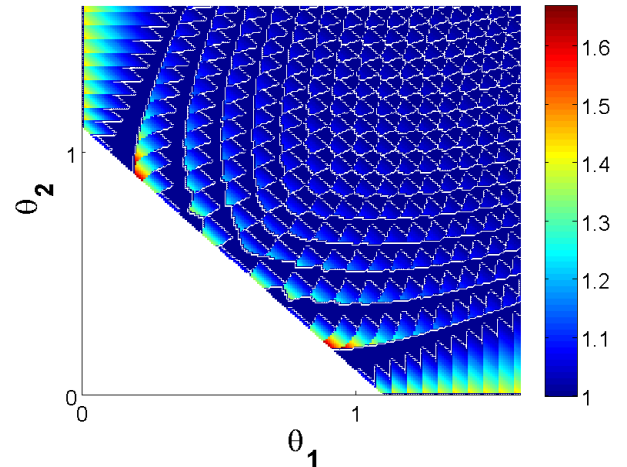


Figure 3: Benefit of bargaining (BOB) over the $\theta_1 \times \theta_2$ space under the model $x_i = a_i \log(1 + \mu)$

5 DISCUSSION AND CONCLUSION

An immediate extension of the model is to consider the case when the users pay an internet access fee that is proportional to usage. From a peering standpoint, it turns out that this case is identical to the case of flat user pricing in Section 3, and all our previous conclusions apply to this case. To see that this is identical to the case when the users pay flat fee, we observe that the CP surplus functions remain unchanged from the flat access price case. Thus, if $\mu = Q_1 + Q_2$, then the analysis and conclusions of Section 3 apply.

Another variation is to let the ISP be strategic, i.e., $\mu(Q_1, Q_2) \neq Q_1 + Q_2$ but is determined to maximise its profit. It is easy to see that when the q -charges are flat, considering the ISP as being strategic is not a well posed model. To see this note that the ISP surplus is

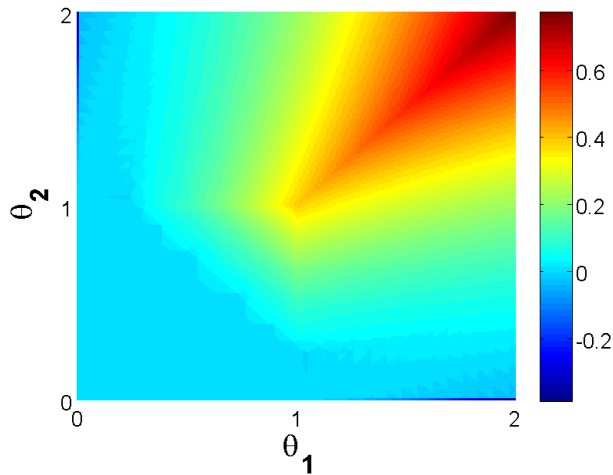


Figure 4: Difference between net CP surplus under bargaining and non-cooperative models, for $x_i = a_i \log(1 + \mu)$

given by $p(x_1(\mu) + x_2(\mu)) + Q_1 + Q_2 - \mu$. Thus, given any (Q_1, Q_2) , the optimal response of the ISP is independent of (Q_1, Q_2) .

A third extension that is possible, is to assume that the CPs pay a volume based q -charge, i.e., they pay per byte. This makes the problem mathematically messier because of loss of convexity and we will not pursue it here. However, with significant simplifications including $x(\mu)$ being linear, and some additional tweaks to the model (e.g., cost of capacity expansion is not linear but some power of the investment), closed form expressions for many of the quantities of interest can be obtained. The qualitative conclusion though is similar to what we have obtained in the previous two sections.

We conclude by reiterating that modeling the interconnection market is critical to developing a comprehensive understanding of the economics of the emerging internet. With the increasing concentration of traffic among a few CPs such models may be an even more crucial input for policy experts.

REFERENCES

- [1] D. Clark, W. Lehr, and S. Bauer, "Interconnection in the internet: The policy challenge," in *39th Research Conference on Communication, Information and Internet Policy*, September 23–25 2011.
- [2] G. Goth, "New internet economics might not make it to the edge," *IEEE Internet Computing*, pp. 7–9, January/February 2010.
- [3] C. Dovrolis, "The evolution and economics of internet interconnections," Submitted to Federal Communications Commission, 2015. [Online]. Available: www.cc.gatech.edu/~dovrolis/Papers/FCC-declaration.pdf
- [4] S. Sen, C. Joe-Wong, S. Ha, and M. Chiang, "A survey of smart data pricing: Past proposals, current plans, and future trends," *ACM Computing Surveys*, vol. 46, no. 2, November 2013.
- [5] R. T. B. Ma, D. M. Chiu, J. C. S. Lui, and V. Misra, "On cooperative settlement between content, transit and eyeball internet service providers," *IEEE/ACM Transactions on Networking*, vol. 19, no. 3, pp. 802–815, June 2011.
- [6] 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, March 2009.
- [7] N. Economides and J. Tag, "Network neutrality on the internet: A two-sided market analysis," *Information Economics and Policy*, vol. 24, no. 2, pp. 91–104, June 2012.
- [8] R. ElDeGawy and R. J. La, "Interaction between a content provider and a service provider and its efficiency," in *Proc. of IEEE ICC*, 2015.
- [9] R. T. B. Ma, "Subsidization competition: Vitalizing the neutral internet," in *Proc. of ACM CoNEXT*, 2014.
- [10] C. Joe-Wong, S. Ha, and M. Chiang, "Sponsoring mobile data: An economic analysis of the impact on users and content providers," in *Proc. of IEEE INFOCOM*, 2015.
- [11] L. Zhang, W. Wu, and D. Wang, "Sponsored data plan: A two-class service model in wireless data networks," in *Proceedings of ACM SIGMETRICS*, 2015.
- [12] 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 Transactions on Networking*, vol. 18, no. 3, pp. 775–787, June 2010.
- [13] J. J. Laffont, S. Marcus, P. Rey, and J. Tirole, "Internet peering," *American Economic Review*, vol. 91, no. 2, pp. 287–291, May 2001.
- [14] G. Shrimali and S. Kumar, "Paid peering among internet service providers," in *Proceedings of GameNets*, October 2006.
- [15] —, "Bill-and-keep peering," *Telecommunications Policy*, vol. 32, no. 1, pp. 19–32, February.
- [16] D. Zarchy, A. Dhamdhere, C. Dovrolis, and M. Schapira, arXiv preprint, October 2016. [Online]. Available: <https://arxiv.org/abs/1610.01314>
- [17] C. Courcoubetis, L. Gyarmati, N. Laoutaris, P. Rodriguez, and K. Sdrolis, "Negotiating premium peering prices: A quantitative model with applications," *ACM Transactions on Internet Technology*, vol. 16, no. 2, pp. 1–22, April 2016.
- [18] C. Courcoubetis, K. Sdrolis, and R. Weber, "Paid peering: pricing and adoption incentives," *Journal of Communications and Networks*, vol. 18, no. 6, December 2016.
- [19] R. T. B. Ma, "Pay or perish: The economics of premium peering," *IEEE Journal on Selected Areas in Communications*, vol. 35, no. 2, pp. 353–366, February 2017.
- [20] R. Hastings, "Internet tolls and the case for strong net neutrality," <https://media.netflix.com/en/company-blog/internet-tolls-and-the-case-for-strong-net-neutrality>, 2014, [Online; accessed 31-July-2017].
- [21] R. B. Myerson, *Game theory*. Harvard university press, 2013.