

P2P-ISP Cooperation: Risks and Mitigation in Multiple-ISP Networks

Aliye Özge Kaya
WINLAB, Rutgers University

Mung Chiang
Princeton University

Wade Trappe
WINLAB, Rutgers University

Abstract—Several proposals on P2P-ISP cooperation have recently been developed using information sharing for locality-based peering. Their benefits in terms of P2P efficiency, ISP cost, and traffic localization have been demonstrated in the single ISP case. However, potential risks associated with such cooperation have not been well examined. This paper develops a taxonomy and a mathematical model to explore the unintended and counter-intuitive behaviors emerging out of these cooperations in the multiple ISP case. Through both numerical examples and analytical results, we illustrate how such behaviors may become damaging to both P2P providers and ISPs, and how they might be mitigated so that the full benefit of cooperation can be maintained.

I. INTRODUCTION

Currently, Internet traffic management is carried out by ISPs whereas CDN and P2P determine the source-destination pairs. P2P providers have limited capability to infer network topology; therefore their peering decisions conflict with the interests of the underlying network. If ISPs disclose the topology to the P2P providers and the P2P providers disclose peer locations to the ISPs, the P2P traffic can be distributed more efficiently.

Recently, new cooperation approaches were proposed to enable such information sharing between ISPs and P2P providers [1], [2], [3]. They are based on providing the P2P providers with partial information on topology, to enable more preferable peering decisions. In Aggarwal et al. [1], the authors propose an oracle supported by the ISP that ranks the peers according to the ISP's preferences. In Choffnes et al. [2], the peers close to the same content distribution network (CDN) server are assigned to each other. In Xie et al. [3], a trusted party offers peering suggestions to P2P after collecting information from both P2P and ISPs.

It is intuitive that such cooperation improves performance for both ISPs and P2P providers, by assigning peers in close proximity to each other, avoiding network bottlenecks, and preventing downloads from distant locations thus decreasing the backbone traffic. However, through information exchange both P2P providers and ISPs get additional freedom to further manipulate the destination and routing of the traffic, which can lead to unintended or even counter-intuitive outcomes of cooperation. This risk factor is particularly prominent in the Internet reality of multiple ISPs coexisting with their own peering relationships.

Take a brief illustrative example: An ISP can offer to another ISP connection to the rest of the Internet. In a such

relationship the ISP who accepts this service gets charged by the provider ISP both for incoming and outgoing traffic. In a peering relationship they exchange traffic without charging each other. Therefore, ISPs are more willing to use the links to their customers or peers than the links to their providers. BGP allows ISPs to show preferences by the path selection. However, ISPs are not always able to avoid certain paths just by tweaking BGP parameters. At this point, P2P providers can further help the ISP by avoiding peering decisions over those paths. Through P2P-ISP cooperation, an ISP might get the advantage to reroute its P2P traffic over the customer links instead of the provider links. This may increase costs for the customer ISPs since they have to send extra content to the provider ISPs.

The above case is one of the several unintended and counter-intuitive behaviors in P2P-ISP collaborations. This paper develops a taxonomy and formal framework to quantify and mitigate these risks. Consequent results can be used to help ISP answer the question: Should I participate in these P2P-ISP cooperations given a particular set of peering relationships with other ISPs? We will illustrate through both numerical examples and analytical results the potential issues when P2P providers and ISPs collaborate through partial information sharing. However, we do not intend to discourage such collaborations. They are generally helpful to all parties involved, and can have significantly positive impact after their risks are discovered and bounded.

The paper is organized as follows. In Section II, we point to the main properties and risks of the recent cooperation proposals. Section III generalizes possible outcomes of the P2P-ISP cooperation. Section IV formulates an analytical framework that models the interaction between the ISPs and P2Ps. Section V and VI summarize the numerical and analytical results. Section VII attempts to prevent the undesired results of cooperation by enforcing fairness. Section VIII gives guidelines to prevent risks in P2P-ISP cooperation.

II. CURRENT P2P-ISP COOPERATION APPROACHES AND THEIR RISKS

We start with a brief and qualitative discussion on some of the risks of the ISP-P2P cooperation approaches in [1], [2], [3].

A. Approach A: ISP supported Oracle in Fig. 1

Aggarwal et al. [1] propose an oracle service supported by the ISP. Here, the P2P provider suggests a list of peers to the

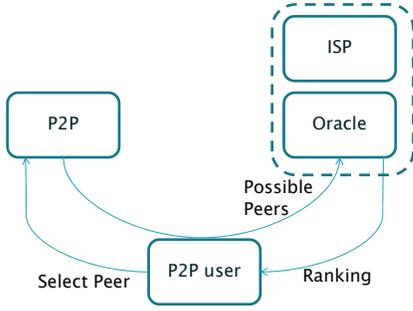


Fig. 1. Approach A: ISP supported Oracle in [1]

end users as shown in Fig. 1. The oracle ranks these peers for the end user. The end user chooses the peers accordingly.

A risk here is that P2P providers may infer the customer-provider and undisclosed peering relationships by querying the oracles of each ISP. ISP oracles can be unintentionally used as traffic monitors and P2P providers can promote their interests. Consequently, the P2P providers can offer additional benefits to favored ISPs by promoting or hiding their peers, leading to preferential treatments to some ISPs.

Alternatively, an ISP oracle has the freedom to rank the peers according to its preferences, which might not always be in the interest of the peers. Assume the ISP prefers a low-cost congested link to a peer-ISP to a costly but uncongested, high capacity link to the provider-ISP. Therefore, it ranks the peers in the peer-ISP higher than the peers in the provider-ISP. Although they are ranked lower the peers in the provider ISP would improve the download experience of the end-users much more since they are reachable through better links. Also, the ISP oracle may provide different information to each P2P provider. This will cause performance gaps between various P2P providers.

B. Approach B: CDN supported Oracle in Fig. 2

CDNs collect location information to redirect users to their low latency servers, which are located in every ISP. Choffnes et al. [2] propose a CDN based oracle called Ono to reduce costly cross-ISP traffic. Here each peer performs periodic DNS look-ups for popular CDN names as shown in Fig. 2. If two peers are redirected to the same servers most of the time, they are assumed to be close to each other. The P2P providers assign these peers to each other. This approach localizes traffic using the hints from CDNs. P2P providers make peering decisions without any direct information from ISPs. Further, this approach enables P2P providers to offer preferential treatments to ISPs. CDN server names can be used as a unique identifier of an ISP, and an ISP can reveal its local CDN server names and ask for preferential treatments.

C. Approach C: P4P in Fig. 3

In the P4P approach proposed by Xie et al. [3], ISPs reveal the network topology to an oracle called iTracker, which lets the P2P providers choose the appropriate peers. iTracker is run by each ISP or by a trusted third party. It assigns each peer to a

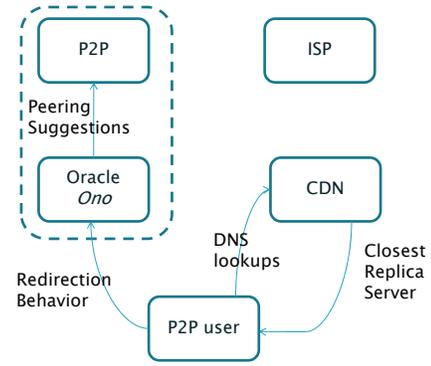


Fig. 2. Approach B: CDN supported Oracle in [2]

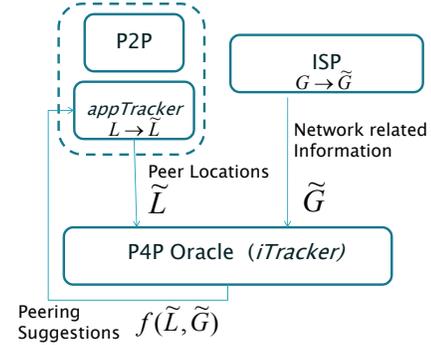


Fig. 3. Approach C: P4P in [3]. L represents the peer locations and G represents the network information. \tilde{L} and \tilde{G} are generalized versions.

cluster with an identifier called PID, and computes a distance metric called p4p distance between these PID locations, based on the hints from ISPs. Another entity called appTracker, owned by the P2P provider, assigns peers close in terms of p4p distance to each other. Without revealing the network topology to the P2P providers, this strategy helps to localize the traffic, promising higher rates to the customers while avoiding transit fees and congestion for the ISPs.

Assigning multiple PIDs within the same ISP's network, the iTracker scrambles the revealed information such that P2P providers cannot directly observe the exact network topology. However, it is still possible for the appTracker to reverse engineer the network topology to a certain accuracy and granularity, e.g., by clustering PID's with similar preferences. Some ISPs also might purposefully reveal their PID's to the appTracker hoping for preferential treatment by the peer assignment. ISPs may hide the true topology from iTrackers and trick the appTracker by peering decisions. Therefore, similar risks as in Approach A remains.

III. CONSEQUENCES OF P2P-ISP COOPERATION

Table I summarizes the main risks as discussed in the last section. Table II further classifies the risks into (more obvious) "drawbacks" and (less obvious) "counter-intuitive" behaviors, and put them right next to the obvious benefits of ISP-P2P collaborations.

TABLE I
COMPARISON OF APPROACHES A, B AND C

	A	B	C
Manipulation of traffic by the ISP alone	ISP has full freedom to manipulate traffic.	not possible	partially (might misrepresent its own network)
Can the P2P provider infer network topology?	partially (use ISP oracle as traffic monitor)	partially (CDNs replica serves as identifiers of ISPs)	partially (cluster PID's with similar preferences, iTracker as traffic monitors)
Preferential Treatment by the P2P providers (hiding or promoting peers)	possible (manipulate the list of the suggested peers)	possible (use CDN Replica servers as identifiers for the preferred ISPs)	possible (misrepresent location of the peers, ignore hints from iTracker)

TABLE II
OUTCOMES OF COOPERATION

Benefit	Drawback	Counter-Intuitive
P2P Efficiency (Higher rates, less delay)	Lose privacy and sensitive information like topology	Unfairness (better benefits for non-cooperating ISP)
Reduce inter-ISP cost	One sided preference setting	Selfish behavior leading degradation of the benefits
Reduce bandwidth consumption	Difficult to verify the exchanged information	

A. Benefits of Cooperation

Having extra information on network topology or peer locations, ISPs and P2P get additional freedom to manipulate the traffic. Efficiently used, this leads to (1) Localization of traffic and reduction of inter-ISP traffic, (2) Reduction of backbone bandwidth usage, (3) Improved rates for the users.

B. Drawbacks of Cooperation

- *No valley peer assignment.* ‘No valley and prefer customer policy’ is a common practice in today’s Internet [4] for choosing a route to a destination: A peer route is preferred over a provider route and customer routes are most preferred. The P2P-ISP cooperation enables the ISPs to use the same policy to show the preference for the destination. Consider an ISP that reveals its customer-provider links to the P2P provider and wants the P2P provider to avoid those links when doing peering assignments. If the P2P provider observes the preference of the ISP, the traffic is only allocated using the in-network links and the links to the customer ISPs. Customer ISPs will face higher costs and congestion since they have to send more content to the provider ISPs.
- *Preferential Treatment.* Consider a P2P provider that observes the preferences of some ISPs more than the others. The favored ISPs and their end users will get more benefits than others. Similarly, an ISP could selectively

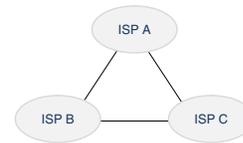


Fig. 4. Network Model

release information in a way to favor some P2P providers over others.

- *Exchanging Misleading Information.* An ISP may represent its network topology inaccurately to the P2P provider by hiding inter-ISP links or falsely reporting congestion to force the P2P provider to do peering assignments within the network. It may manipulate the peering decisions of the P2P provider. In order to reduce the inter-ISP traffic for an ISP, the P2P provider might not reveal all of the peers in that ISP. This induces the peers in a neighboring ISP to download from a distant location.

C. Counter-intuitive Outcomes of Cooperation

Cooperation may lead to unfair scenarios where non-cooperating ISPs (or cooperating but “less preferred ISPs”) get more benefits than cooperating and “preferred ISPs”. They might benefit by not cooperating. Selfish behavior by either the P2P providers or the ISPs may lead to degradation in the benefits of cooperation.

Benefits for non-cooperating ISPs: The cooperating ISPs may not be willing to send content to the peers in non-cooperating ISPs, since the location of peers in those cannot be inferred and might require costly paths to reach. If no other ISP is willing to provide content, the non-cooperating ISPs will be forced to supply their traffic demand within the network. Usually the average rate to download a content within the network is much higher than to download it through an inter-ISP link. If the non-cooperating ISP has enough peers to support its demand, it will achieve higher rates than the cooperating ISP, whose users downloads also from the peers in other ISPs. Knowing that the neighboring ISP is cooperating would be a valuable information to decide on joining cooperation.

As an illustrative example, consider the network in Fig. 4. Each ISP requests a traffic volume V . If an ISP cooperates with the P2P provider, it gets the file within the network with probability $1/2$ or from any of the other two ISPs with probability $1/4$. If it does not collaborate, it will get the file with probability $1/3$ from any ISP including itself. Let the rate for downloading a file from the peers within the network be R and the average rate to download from a peer from another ISP $R/2$. If none of the ISPs cooperates with the P2P provider, the average download rate for the end users is $0.6R$. If all cooperate, the rate is $0.75R$. If ISP C is not cooperating, ISP A and ISP B are reluctant to send content to the ISP C. Observing their preference, the P2P provider induces ISP C to download the whole content within the network. Assuming that ISP C has enough peers, the average download rate for

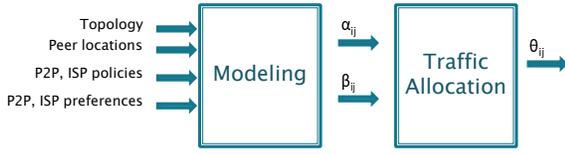


Fig. 5. Overview of the framework

ISP C users will be R , which is higher than the average rate for the cooperating ISPs. In this case, ISP C is better off by not cooperating. If ISP A and ISP B are providers to most other ISPs, they are eager to send content even to non-cooperating ISPs. So, the P2P provider induces the non-cooperating ISP C to download all traffic content from ISP A and ISP B. The average rate for ISP C customers is $0.5R$ less than its average rate when none of the ISPs had cooperated with the P2P provider.

IV. MODELING COOPERATION IN MULTI-ISP NETWORKS

To quantify the risks associated with the three approaches of P2P-ISP cooperation, we will first develop a mathematical language before the numerical and analytical results are presented. The common entity in all the approaches summarized in Section II is an oracle, which collects information and gives hints to the P2P provider. The main difference between these approaches is whether the ISP gets individual and indirect involvement (Approach A), no involvement (Approach B), or collective, direct involvement (Approach C) in oracle. These hints may depend on the network topology, peer locations, ISP and P2P preferences and their incentives. Observing these hints, the ISP downloads content more from some ISPs than the others. We abstract the oracle's hints by a metric β_{ij} and the possible manipulations by the P2P provider by a metric α_{ij} . Fig. 5 shows the overview of the model.

A. Oracle hint metric: β_{ij}

The willingness of an ISP to exchange traffic with another ISP depends on network policies, distance, congestion, reliability of the link, cost relationships, etc. Providers are willing to exchange traffic with their customers. On the contrary, the customers prefer to avoid the links to their providers. Among all 26000 ISPs in the Internet, 20–30 percent are providers and the rest are customer ISPs [4]. The oracle needs to negotiate the conflicting preferences of the ISPs on both sides of a link.

We model the hints the oracle provides to the P2P provider about ISP preferences by the metric β_{ij} . The oracle assigns the weight β_{ij} to the link from ISP i to ISP j , showing the joint preference of ISPs i and j for sending unit traffic volume from ISP i to ISP j . If ISP j requests a traffic volume, the P2P provider considers β_{ij} s for all incoming links to ISP j .

Let β_{ij}^s be the willingness of ISP i to send content to ISP j and β_{ij}^r be the willingness that ISP j to receive content from ISP i as conveyed to the oracle. Specifically, β_{ij} is a compromise between β_{ij}^r and β_{ij}^s . If the ISPs do not convey β_{ij}^r and β_{ij}^s , the oracle determines β_{ij} based on other network metrics it infers.

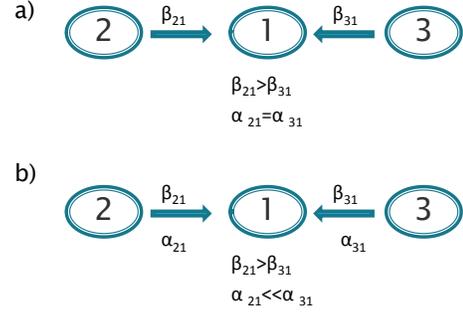


Fig. 6. a) ISP 2 provides ISP 1 more content than ISP 3 if P2P observes their preferences b) By assigning lower α to the link from ISP 2, the P2P provider prevents ISP 2 from providing most of the content

B. Manipulation by the P2P Provider: α_{ij} metric

Since the P2P provider is the party making the peering decisions, it has the freedom to ignore or enforce preferences of some ISPs (captured by metric β_{ij}) more than others, to reroute or localize traffic. We model such P2P incentives to manipulate traffic by α_{ij} . The metric α_{ij} assigned by the P2P provider to manipulate the traffic flow from ISP i to ISP j . This can be independent from what the ISPs themselves prefer through the β_{ij} metrics. As an example, by assigning a large α_{ij} to a link with lower β_{ij} , the P2P provider induces ISP i to send more content to ISP j . Similarly, by assigning a small α_{ij} to a link with higher β_{ij} , it can prevent ISP i from sending as much content to ISP j as originally preferred by ISP i .

By varying α , P2P providers might observe the preferences of some ISPs more than the others to fulfill special agreements with them. Additionally, by assigning higher α values to in-network links than to the other links, they can choose to localize the traffic. In Figure 6, ISP 1 receives more content from ISP 2 than ISP 3 since $\beta_{21} > \beta_{31}$, if P2P provider does not manipulate the traffic. However, by setting $\alpha_{21} \ll \alpha_{31}$, it can prevent ISP 2 from providing more content than ISP 3.

C. Peering Decisions by the P2P provider

The P2P provider is the party doing the peering decisions. If an ISP requests a traffic volume, the P2P provider determines which ISP should send what fraction of that content. By allocating the traffic it takes into account its own interests and the preferences of the ISPs, which are reflected in the hints from the oracles. We have modeled the P2P provider's interests by α_{ij} and the hints from the oracles by β_{ij} .

Let θ_{ij} be the fraction of the traffic volume that ISP j downloads from ISP i . We define the utility of the P2P provider as ISP j downloads θ_{ij} percent of its traffic volume from ISP i by a log function to achieve proportional fairness (with a vertical shift of 1 to avoid negative utility):

$$U_{ij}(\theta_{ij}, \alpha_{ij}, \beta_{ij}) = \alpha_{ij} \log(1 + \beta_{ij} \theta_{ij}). \quad (1)$$

Recall that β_{ij} is an indicator of the joint preference of ISP i and ISP j for sending unit traffic from ISP i to ISP j . Thus $\beta_{ij} \theta_{ij}$ indicates to the P2P provider the satisfaction

of ISP i and ISP j from allocating θ_{ij} fraction of the traffic. We also assume that the P2P provider places more weight in its own preferences than the ISP hints and this is reflected in the choice of the utility function (1) with α_{ij} being outside the logarithmic function. When ISP j requests traffic volume V , the P2P provider allocates fractions θ_{ij} separately for each ISP j by maximizing the sum utility, $\sum_{i=1}^N U_{ij}(\cdot)$ for all the N ISPs from which ISP j downloads traffic:

$$\max_{\theta_{ij} \geq 0} \sum_{i=1}^N \alpha_{ij} \log(1 + \beta_{ij} \theta_{ij}) \quad (2a)$$

$$\text{s.t. } \sum_{i=1}^N \theta_{ij} = 1. \quad (2b)$$

This is mathematically equivalent to the water-filling problem [5] and the optimal solution is

$$\theta_{ij} = \left(\frac{\alpha_{ij}}{\lambda} - \frac{1}{\beta_{ij}} \right)^+. \quad (3)$$

The water levels are modified by the P2P preferences and the ISP preferences act similar to the levels over which water is filled. Parameter λ is calculated from

$$\sum_{i=1}^N \left(\frac{\alpha_{ij}}{\lambda} - \frac{1}{\beta_{ij}} \right)^+ = 1. \quad (4)$$

Let \mathcal{P} be the set of all ISPs i providing content to ISP j , such that $\theta_{ij} > 0$. From (3) and (4) the closed form for θ_{ij} is

$$\theta_{ij} = \frac{\alpha_{ij}}{\sum_{k \in \mathcal{P}} \alpha_{kj}} \left(1 + \sum_{k \in \mathcal{P}} \frac{1}{\beta_{kj}} \right) - \frac{1}{\beta_{ij}}. \quad (5)$$

From (5), we state the following lemma which can be proved easily:

Lemma 1: θ_{ij} is an increasing function of both α_{ij} and β_{ij} . Thus links with high α_{ij} and β_{ij} receive higher traffic volumes.

D. Recovering Approaches A, B, C

The approaches A, B, C differ in how the β_{ij} is determined. In approach A in Section II-A, β_{ij} shows only the preferences of the ISP who employ the oracle. In approach B in Section II-B, ISPs do not have direct involvement on determining β_{ij} . Here β_{ij} is set inversely proportional to the inferred distances between ISPs. In approach C of Section II-C, the β_{ij} is an indicator of the p4p distances which takes into account of the preferences of ISPs on both sides of a link. The ISPs that are close in terms of p4p distance will be assigned higher β_{ij} . Solving, (2), the traffic is fairly allocated between ISPs letting ISP i with higher β_{ij} provide much more content to the the ISP j unless P2P providers interfere by varying α_{ij} . Assigning weights to the links for traffic allocation is common in the existing P2P-ISP cooperating approaches. As an example, in approach C traffic allocation is done by minimizing the p4p distances between clusters of peers which corresponds here maximizing β_{ij} . The undesirable P2P-ISP behavior is captured by ability to choose any value for α_{ij} and β_{ij} .

The risk we just modeled is the capability of P2P providers, through cooperation with ISPs, to induce ISP traffic management strategies that are undesirable. Such risks are also present in the current networks, where they have to infer each other's information. The P2P providers have incentives as modeled by α and but limited information on β_{ij} , which limits their actions. However, if there is cooperation they could obtain it without relying on noisy network measurements and traffic monitoring, thus amplifying these risks.

E. Cost Model

We denote the cost for ISP j to exchange content (send or receive) with ISP i by q_i^j . If $q_i^j < 0$ ISP j exchanges traffic with ISP i through a customer ISP link, it receives $|q_i^j|$ per unit traffic as income. If $q_i^j > 0$ ISP j exchanges traffic with ISP i through a provider ISP link paying $|q_i^j|$ per unit traffic. $q_i^j = 0$ if ISP i is reachable from ISP j through a peering link or if $i = j$. The net cost to download V_j for ISP j is

$$C_{avr,j}^r = \sum_{i \in \mathcal{P}} \theta_{ij} q_i^j V_j \quad (6)$$

The net cost for ISP j to send content to the other ISPs in \mathcal{P} is

$$C_{avr,j}^s = \sum_{i \in \mathcal{P}} \theta_{ji} q_i^j V_i \quad (7)$$

The net cost of ISP j in exchanging traffic with ISPs in \mathcal{P} is

$$C_{avr,j} = C_{avr,j}^r + C_{avr,j}^s; \quad (8)$$

The net income is defined as $I_{avr,j} = -C_{avr,j}$.

F. Efficiency Model

Let R_{ij} be the rate which peers in ISP j experience when they download content from peers in ISP i . Thus, θ_{ij} fraction of V_j is downloaded in $\theta_{ij} V_j / R_{ij}$ time units. The average rate to download V_j is

$$R_{avr,j} = \left(\sum_{i \in \mathcal{P}} \frac{\theta_{ij}}{R_{ij}} \right)^{-1}. \quad (9)$$

We define the rate efficiency η_r and cost efficiency η_c for the total network as

$$\eta_r = \frac{1}{N} \sum_{j=1}^N R_{avr,j}, \quad \eta_c = -\frac{1}{N} \sum_{j=1}^N C_{avr,j}. \quad (10)$$

G. Fairness Model

If ISPs and P2P providers cooperate, the increase in rate and decrease in cost will be compared to the cost and rate if there were no cooperation. Let R_n^j and C_n^j be the average rate and cost when the ISP j does not cooperate with the P2P provider. We define the rate and cost fairness metrics r^j and c^j of ISP j as

$$r^j = \frac{R_{avr,j} - R_n^j}{R_n^j}, \quad c^j = \frac{C_n^j - C_{avr,j}}{C_n^j}. \quad (11)$$

Obviously, it is desired to have $r^j > 0$ and $c^j > 0$ for a cooperating ISP.

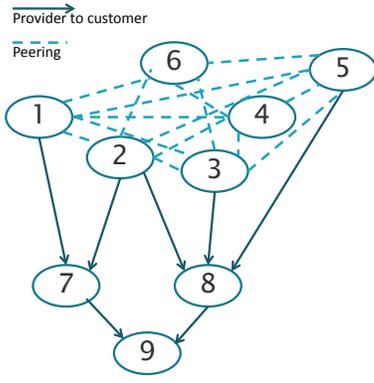


Fig. 7. Multi-ISP topology in numerical examples

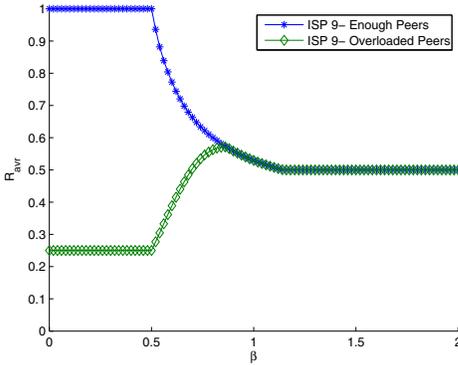


Fig. 8. Average Rate versus β : ISP 9 is not cooperating. The cooperating ISPs do not want to provide content to ISP 9. If ISP 9 has enough peers, it experiences exceptionally high average rates as it is induced to supply its traffic demand within the network ($\beta \rightarrow 0$). If its peers are overloaded it will face disproportionately worse rates as $\beta \rightarrow 0$. As $\beta \rightarrow 2$, ISP 9 is induced to supply to download its whole traffic demand from other ISPs resulting worse rates than the random peer assignment ($\beta = 1$)

V. NUMERICAL EXAMPLES

We consider the topology as depicted in Fig. 7. It is a reannotated graph of a real topology inferred in [6]. ISP 1 to ISP 6 are peers and exchange traffic between each other at no cost. ISP 7 and ISP 8 are their customers. ISP 9 is a customer of both ISP 7 and ISP 8.

A. Disproportional benefits for non-cooperating ISPs

Assume all ISP's in Fig. 7 except the ISP 9 are cooperating with the P2P provider. The P2P provider does not manipulate ISP preferences, and sets $\alpha_{ij} = 1$ for all links. We set $\beta_{i9} = \beta$ for $i = 1 \dots 8$ and $\beta_{99} = 1$. If $\beta > \beta_{99}$, ISP 9 will be induced to download more content from the peers in the other ISPs, if $\beta < \beta_{99}$ it would be forced to supply its traffic demand within the network. $\beta = \beta_{99}$ corresponds to the scenario where the peers are randomly assigned to each other. Fig. 8 shows the variation of the average rate versus β . For $\beta < 1$ we consider two cases: (1) ISP 9 has enough peers to supply its traffic demand. (2) Supplying the whole traffic demand would overload the peers within the network and reduce in-network

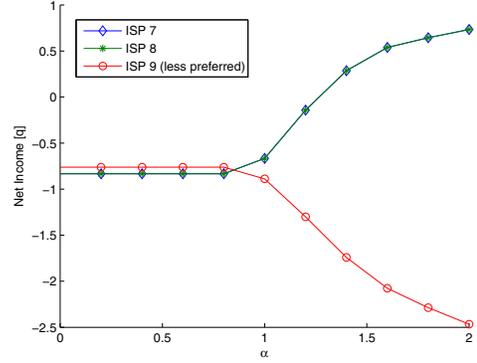


Fig. 9. Net Income versus α : Cost for ISP 9 increase as the P2P provider induces him to exchange traffic with ISP 7 and ISP 8 both ways ($\alpha \rightarrow 2$). Cost decrease for ISP 9 as the P2P provider prevents him to exchange traffic ISP 7 and ISP 8 ($\alpha \rightarrow 0$). As $\alpha \rightarrow 0$ the expenditures of ISP 7 and ISP 8 increase and their income from ISP 9 decrease. They are worse off than ISP 9 in terms of net income.

download rates. In the first case the ISP 9 is better off by not cooperating since it could get maximum possible average rate by supplying its traffic within the network. In the second case its peers are overloaded. It experiences much worse rates than the case where all peers in the network are randomly assigned to each other.

In this example we set the average rate to download within the ISP to R and the average rate to download from other ISPs to $R/2$. This assumption is consistent with the previous work [2] that the download rates from the peers within the same ISP are higher. If ISP 9 does not have enough peers and it is able to supply only $\theta_{max} = 0.25$ of its traffic volume at rate R , then we set the rate to download within the network to $\min(R, \theta_{max}/\theta_{99}R)$ to model the decrease in download rates when the peers are overloaded.

B. Manipulations by the P2P provider

All ISP's in Fig. 7 request a traffic volume of V and exchange traffic between each other. ISP 9 pays q per unit exchanged traffic to ISP 7 and ISP 8. ISP 7 and ISP 8 exchange with each other at cost q per unit traffic since they are connected through their providers. They meet the demand within the ISP with no cost. We set $\beta_{ij} = 1$ for $i \neq j$ and $\beta_{ii} = 2$, showing the preference of ISPs for localizing the traffic.

Assume the P2P provider has incentive to manipulate the traffic for ISP 9. We set $\alpha_{79}, \alpha_{97}, \alpha_{89}, \alpha_{98}$, which are metrics for incoming and outgoing links of ISP 9 equal to α . Other α_{ij} 's are set to 1. As $\alpha \rightarrow 0$ in Fig. 9 the P2P provider prevents ISP 7 and ISP 8 to exchange traffic with ISP 9. ISP 9 is induced to meet most of its traffic demand within the network. ISP 7 and ISP 8 are induced to exchange traffic with each other and their providers, instead of ISP 9. Their net revenue from inter-ISP agreement gets lower than that of ISP 9. As $\alpha \rightarrow 2$ the P2P provider induces ISP 9 to download its traffic demand from ISP 7 and ISP 8. At the same time, ISP

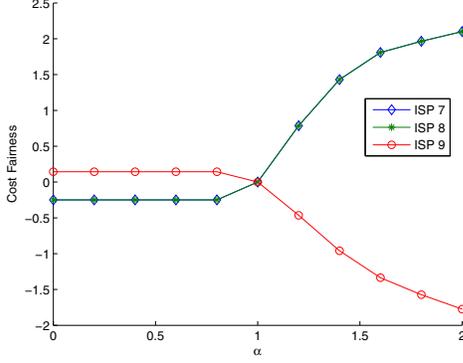


Fig. 10. Cost Fairness versus α : As $\alpha \rightarrow 0$ the cost fairness for ISP 9 is higher than ISP 7 and ISP 8. ISP 7 and ISP 8 get negative cost fairness although they are cooperating. As $\alpha \rightarrow 2$ the cost fairness for ISP 9 decrease as the P2P provider induces him to exchange traffic with ISP 7 and ISP 8 both ways. As $\alpha \rightarrow 2$ the cost fairness increase for ISP 7 and ISP 8 as they exchange much more content with ISP 9.

7 and ISP 8 supply their traffic demand from ISP 9. The cost increases tremendously for ISP 9, whereas ISP 7 and ISP 8 reduce their expenditures and increase their incomes.

Fig. 10 shows the corresponding cost fairness metric versus α . As $\alpha \rightarrow 0$ we see that ISP 7 and ISP 8 get worse cost fairness than ISP 9. Moreover, ISP 7 and ISP 8 get negative cost fairness although they are cooperating since they are induced to exchange content with their providers instead of ISP 9. As $\alpha \rightarrow 2$ the cost fairness for ISP 7 and ISP 8 increase since they exchange much more content with ISP 9. On the contrary, as $\alpha \rightarrow 2$ the cost fairness decreases for ISP 9 as it is induced to exchange more content with its provider ISPs.

VI. CONDITIONS FOR RISK-FREE COOPERATION

In this and the next section, we focus on mitigation aspects of some particular risks of P2P-ISP cooperation discussed so far. Consider an ISP j . Let the average rate to download $\theta_{jj}V$, the traffic within the network, be R , and the average rate to download $(1 - \theta_{jj})V$ from other ISPs be R/K . Since the average rate to download the content within the ISP is greater than from out of network, $K > 1$. The peers in a cooperating ISP j will experience the average rate

$$R_{j,c} = \left(\frac{\theta_{jj}}{R} + \frac{K(1 - \theta_{jj})}{R} \right)^{-1} = \frac{R}{\theta_{jj} + K(1 - \theta_{jj})}. \quad (12)$$

If an ISP j cooperates, θ_{jj} is determined from (5). If ISP j does not cooperate, the peer assignment is random and $\theta_{ij} = 1/N$ for all ISPs i . Therefore, the average rate for the peers in ISP j will be

$$R_{j,n} = \frac{N}{\frac{1}{R} + \frac{K(N-1)}{R}} = \frac{NR}{KN - K + 1}. \quad (13)$$

Lemma 2: The peers in an ISP j will always receive a higher average rates when the ISP cooperates than when it does not if

$$\theta_{jj} > \frac{1}{N} \quad (14)$$

Proof: For $K > 1$, the peers in ISP j will get always better average rates when ISP j cooperating and $R_{j,c} \geq R_{j,n}$. Substituting for $R_{j,c}$ and $R_{j,n}$ from (12) and (13) we obtain the desired result after some algebraic manipulations. ■

Theorem 1: Let $|\mathcal{P}|$ be the cardinality of \mathcal{P} . If $\alpha_{ij} = \alpha, \forall i$ for a given ISP j , the peers in that ISP j get higher average rates, when ISP j cooperates and if the hint metrics satisfy the following condition:

$$\beta_{jj} > \frac{|\mathcal{P}|}{1 + \sum_{k \in \mathcal{P}} \frac{1}{\beta_{kj}} - \frac{|\mathcal{P}|}{N}}. \quad (15)$$

Proof: Setting (5) in (14) we obtain,

$$\frac{\alpha_{jj}}{\sum_{k \in \mathcal{P}, k \neq j} \alpha_{kj} + \alpha_{jj}} \left(1 + \sum_{k \in \mathcal{P}, k \neq j} \frac{1}{\beta_{kj}} + \frac{1}{\beta_{jj}} \right) - \frac{1}{\beta_{jj}} > \frac{1}{N} \quad (16)$$

The result follows from substituting $\alpha_{ij} = \alpha$ for all i , in the condition (16). ■

Remark 1: If $|\mathcal{P}| = N$, then condition (15) simplifies to:

$$\beta_{jj} > \frac{N}{\sum_{k=1}^N \frac{1}{\beta_{kj}}}, \quad (17)$$

the harmonic mean of the hint metrics.

Remark 2: A cooperating ISP might choose to hide its peers from the non-cooperating ISPs to reduce its outgoing traffic. That will reduce number of ISPs from which the non cooperating ISP can get content. Now assume only L of the N ISPs are willing to provide content to the non-cooperating ISP. In this case the rate $R_{i,n}$ (13) will increase since $L < N$. As L decreases it is more likely that condition (14) is violated.

VII. RISK MITIGATION STRATEGIES

We present one of the many possible risk mitigation strategies in this section. Since the P2P provider is the party doing the peering decisions, it has the capability to increase the rate and cost fairness for a specific ISP by reallocating some fraction of the traffic demand V . When an ISP j joins cooperation it lets the P2P provider know R_n^j and C_n^j and its minimum fairness levels r^j and c^j as requirements to continue the cooperation. Similarly, the P2P provider sets thresholds on fairness levels to prevent some ISPs from getting disproportional gain or loss from cooperation.

Let θ be the traffic allocation vector including θ_{ij} s for ISP j determined as a solution of (2). If the traffic allocation θ results in a rate fairness r (11) and average rate is R_{avr} but the desired fairness level is \hat{r} , the P2P provider chooses to allocate only κV based on (2) where $0 < \kappa < 1$. It reallocates the remaining $(1 - \kappa)V$ using another strategy which yields the average rate \hat{R} . Allocating κ fraction of V by solving (2) would again yield the average rate R_{avr} and the traffic allocation vector $\kappa\theta$. Assume the P2P provider has another traffic allocation strategy which could result in an average rate

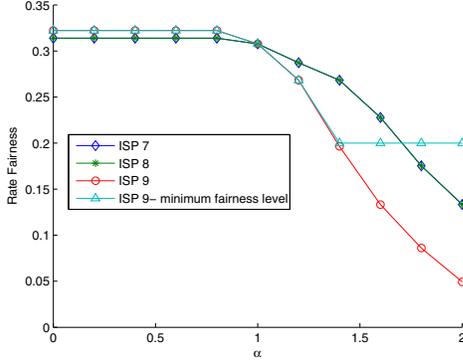


Fig. 11. Rate Fairness: As $\alpha \rightarrow 2$ the rate fairness drops for all ISPs. ISP 9's rate fairness curve decreases steeper than the curves for ISP 7 and 8. The curve denoted with 'ISP-9 minimum fairness level' shows the case where the P2P provider enforces a minimum fairness level $\acute{r} = 0.2$ for ISP 9. The minimum fairness level is reached by reallocating $(1 - \kappa)V$ within the network at higher rate.

\tilde{R} . An example to such strategy could be downloading content within the network at higher rate. Based on this strategy, \tilde{R} is set equal to the average rate to download within the network. To reduce r , the P2P provider could choose the strategy that allocates $(1 - \kappa)V$ evenly to the other ISPs. In this case \tilde{R} will be set to the average rate to download from other ISPs.

If the P2P provider reallocates $(1 - \kappa)V$ using the strategy that results in an average rate of \tilde{R} then fairness level will be raised to \acute{r} . The P2P provider determines κ from:

$$\kappa = \frac{(r + 1)(\tilde{R} - R_n(\acute{r} + 1))}{(\acute{r} + 1)(\tilde{R} - R_n(r + 1))} \quad (18)$$

Similarly, if the traffic allocation θ results in a cost fairness c and average cost is C_{avr} but the desired fairness level is \acute{c} , the P2P provider needs to reallocate $(1 - \kappa)V$ using a traffic allocation strategy that results in the average cost \tilde{C} . κ is determined by:

$$\kappa = \frac{C_n \acute{c} - C_n + \tilde{C}}{C_n c + C_n + \tilde{C}} \quad (19)$$

We plot in Fig.11, for the ISPs in SectionV-B, the fairness metric r versus parameter α . For ISP 9 we consider two cases: (1) If the P2P provider does not control the rate fairness, the rate fairness drops to $r = 0.05$ as $\alpha \rightarrow 2$. (2) If the fairness drops below the threshold $\acute{r} = 0.2$ the P2P provider reallocates some fraction of the traffic in order to reach \acute{r} . The strategy of the P2P provider to download within the network yields an average rate of R . The P2P provider sets $\tilde{R} = R$ and determines κ from (18). Reallocating $(1 - \kappa)V$ at average rate R prevents raises fairness metric back to \acute{r} .

The P2P provider could also detect uncooperative behavior of ISPs by monitoring changes in rate and cost efficiency (10). The disproportionate losses or gains for some ISP would result in noticeable changes in efficiency.

VIII. CONCLUDING REMARKS ON P2P-ISP COOPERATION

We believe that P2P and ISP need to cooperate, and the benefits of sharing information in general can be larger than the downsides. Towards realizing this goal, we focus this paper on the risks of such cooperation, by building a mathematical model for the three recent approaches of cooperation, classifying the potential risks and metrics, and developing numerical examples and analytical results on some of the major unintended or counter-intuitive behaviors, especially when multiple ISPs co-exist and form their own peering relationships.

The risks associated with P2P-ISP cooperation can be mitigated to a great extent with appropriate changes to the existing proposals. Some of the guiding principles are as follows, and their quantification presents an interesting next step in fully understanding the best P2P-ISP collaboration mechanism:

- 1) *Avoid ISP owned oracle for inter-ISP peering decisions.* An ISP supported oracle can give better hints for doing efficient in-network peering decisions. However, ISPs may have conflicting preferences for inter-ISP links. Therefore, such oracles are selfish, when it comes to making inter-ISP peering assignments. A two-tier hierarchy may be used to overcome this problem. Local oracles like the iTracker [3] or the ISP oracle in [1] could be used for peering decisions within the corresponding ISP. For inter-ISP peering decisions, centralized oracles are needed to balance the preferences of multiple ISPs in a fair manner.
- 2) *Auditors.* It is difficult to ensure that P2P providers or ISPs provide correct information to each other. However, the consequences of such malicious behavior can be detected by third parties. To enforce fairness, independent parties can offer tools to monitor P2P traffic. They may deploy in each cooperating ISP traffic monitors for the verification of information or to give statistical information on peering tendencies.
- 3) *Partial observance of the oracle hints.* P2P providers can still choose to make some portions of the peering assignments randomly or based on delay.

REFERENCES

- [1] V. Aggarwal, A. Feldmann, and C. Scheneideler, "Can ISPs and P2P systems co-operate for improved performance?" *ACM SIGCOMM Computer Communications Review*, no. 3, July 2007.
- [2] D. Choffines and F. Bustamante, "Taming the torrent: A practical approach to reducing cross-ISP traffic in P2P systems," *ACM SIGCOMM*, August 2008.
- [3] H. Xie, R. Yang, A. Krishnamurthy, Y. Liu, and A. Silberschatz, "P4P: Provider portal for applications," *In Proceedings of ACM SIGCOMM 2008*, August 2008.
- [4] R. V. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang, "In search of the elusive ground truth: The Internet's AS-level connectivity structure," *SIGMETRICS Perform. Eval. Rev.*, vol. 36, no. 1, pp. 217–228, 2008.
- [5] T. Cover and J. Thomas, *Elements of Information Theory*. John Wiley and Sons, 1991.
- [6] W. Muehlbauer, A. Feldmann, O. Maennel, and S. Uhlig, "Building an AS-Topology model that captures route diversity," *In Proceedings of ACM SIGCOMM 2006*, September 2006.