Is More P2P Always Bad for ISPs?
An Analysis of P2P and ISP Business Models

Qi Liao
Department of Computer Science
Central Michigan University, USA
Email: liao1q@cmich.edu

Zhen Li
Department of Economics and Management
Albion College, USA
Email: zli@albion.edu

Aaron Striegel
Department of Computer
Science and Engineering
University of Notre Dame, USA
Email: striegel@nd.edu

Abstract—**Internet Service Providers (ISPs) face increasing bandwidth pressure from rising access demand by users, especially P2P and VoD applications. Traditionally, P2P has been viewed as tremendously negative from the perspective of the ISP. In this paper, we question this assumption and study the impact of P2P applications on the effective ISP functionality. We perform an economic analysis to show that a higher P2P penetration rate does not necessarily lead to increased ISP bottleneck link bandwidth pressure. Our results show that the local serving rate is critical for the sustainability of the ISP business model as well as for the benefit of P2P users.**

I. INTRODUCTION

The large and rapid growth of data in recent years has given rise to the widely discussed phenomenon of "big data." Though the term has been used to convey various concepts, from a greater scope of information to social media data, big data is centered on huge quantities of data. One of the earliest drivers of data quantities has been peer-to-peer (P2P) networking whereby not only the demand is significant but the number of participants is also massive [1]. Although easily accessible streaming media such as Netflix and YouTube has led to a decrease in P2P file sharing, P2P systems such as BitTorrent still dominate upstream traffic with over 36.35% as of 2013 [2]. Moreover, with continued pressure to scale, P2P continues to receive research attention for massive commercial, legal content distribution [3]–[5]. Several recent findings show the sign for the resurgence of P2P while moving P2P service to the cloud via the notion of "seed box" [6].

Notably, P2P traffic raises concerns among Internet Service Providers (ISPs). The common wisdom is that ISPs are under pressure in the face of P2P, whose potential impacts on ISPs include the increasing load of traffic and the changing traffic patterns [3]. In contrast to other applications such as web browsing and emails that have predictable on/off patterns, P2P applications by nature serve as both client and server and typically have much longer durations. P2P espouses significant increased usage of uplinks as well as downlinks and if unchecked, would consume a disproportionately large share of the transit capacity of ISPs versus traditional applications. The large amount of traffic, especially uplink traffic, imposes significant pressure on the bottleneck link capacity of an ISP. It is this potential imbalance that motivates our study on the impacts of P2P traffic in the era of big data on the business model of a traditional ISP.

Under ISPs' current business model of selling "unlimited" broadband access at flat monthly rates, a handful of P2P users can consume a high volume of network capacity. According to Time Warner Cable, 5% of its broadband customers consume more than 50% of the bandwidth [7]. ISPs have to pay a high cost of carrying traffic while receiving little revenue. Since ISPs practically oversell their bottleneck link to provide access to many subscribers, with an increasing amount of constant traffic, ISPs cannot multiplex the same number of users. The ISPs may then be forced to expand capacity, to control the number of users, or to limit and throttle P2P traffic.

Since it is unlikely that ISPs can truly eliminate P2P, we would argue that P2P users and ISPs need to find an equilibrium point that is beneficial to both sides. We note that more P2P applications do not necessarily hurt ISPs since higher P2P penetration rate will likely increase content local serving rate. With P2P traffic localization such as P4P [8], ISPs gain since data does not have to traverse the transit link and users gain as well with lower latency. Motivated by the desire to seek for a possible mutually beneficial relationship between P2P and ISPs, we model their interactions in a game theoretical framework and solve the resulting optimization problem. We show that there exists a critical value of the local serving rate that would make ISPs indifferent with or without P2P. The fact that in reality ISPs usually are not in favor of P2P implies that the actual local serving rate may be far below the critical level. We also propose the concept of *P2P bots* that can potentially affect the local serving rate, i.e., the key determinant of P2P's net effects on ISPs.

The rest of the paper is organized as follows. Section II discusses the current business model of ISPs and users' choice when shopping for Internet services. A game between competing ISPs is developed and model solutions in presence of P2P are derived. Most importantly, we discuss the insight on how P2P local serving rate can affect the ISP business model. Section III considers the importance and feasibility of the localization of P2P traffic and conducts a simulation study to capture the dynamics of P2P on ISP traffic flows. Section IV briefly discusses some related work while Section V offers several concluding remarks.

II. A MODELING ANALYSIS OF ISPS AND P2P

In this section, we first describe the business model of ISPs and define the profit-maximization optimization problem. We then use a game model to analyze the interactions between competing ISPs. As a base model, we derive insights regarding the optimal choices by ISPs and end users. Then, we discuss the interactions between ISPs and P2P applications and emphasize the challenge P2P imposes on the ISPs' business model of overselling "unlimited" bandwidth access at flat rates. The key P2P local serving rate and how it affects the ISPs' business model are discussed.

A. Traditional ISP business model

ISPs can be roughly grouped into two categories according to their range of services: local ISPs that provide Internet service in small regions to end users, including P2P clients, and transit ISPs that transfer data between local ISPs [9]. Local ISPs[1] serve end users directly: end users pay for Internet services to ISPs, and ISPs incur costs when operating their networks and purchasing upstream connectivity from transit providers.

Bandwidth subscription is usually asymmetric with uplink at a smaller portion. Suppose a typical user i subscribes b_i^d downlink bandwidth and b_i^u uplink bandwidth, and $b_i^u \ll b_i^d$. Combining all n users, the total bandwidth sold by an ISP is $\sum_{i=1}^{n}(b_i^d + b_i^u)$. In principle, if ISPs sell services below the capacity limit, users would be guaranteed the subscribed bandwidth at any time. Nevertheless, few ISPs would adopt this practice with the belief that users are unlikely to saturate services simultaneously. The ISP overselling of bottleneck link capacities can be seen as $\sum_{i=1}^{n}(b_i^d + b_i^u) > B$.

Since AOL started offering unlimited access plans in 1996, the "all-you-can-eat" buffet model has been the major business model for local ISPs. Let $\bar{P}$ be the monthly flat rate paid by end users and $C(B)$ be the cost function of the ISP, where B measures the bottleneck link capacity of the ISP. The profit of the ISP can be written as $Profit = n \times \bar{P} - C(B)$.

The profitability of the ISP depends on many factors. We focus on the portion of the expenses and profits that are directly related to the transit costs of carrying bottleneck traffic. Thus the key determinants of the ISP's profitability include its number of end users, price charged, and bottleneck link capacity.

B. Limited user choice

When subscribing for Internet service, the options of broadband users can choose from are extremely limited, covered usually only by two or three providers (e.g., ADSL, cable, or satellite). We call this an *oligopoly*, an economic term referring to a market dominated by a small number of sellers (oligopolists). ISPs may offer tiered services by charging higher monthly rate for larger bandwidth, basic package vs. premium package for example. The highest monthly rate users are willing to pay depends on the level of utility users expect to receive from the service. We denote the maximum possible rate as $P^* = U(QoS)$ where U is the utility function measuring how much users value the Internet service, which is normally increasing in the quality of service (QoS). Users shop for price difference *ex ante* (prior to purchase) and the ISPs use QoS to retain users *ex post* (after purchase).

C. Optimization problem

In an oligopolistic market with a small number of sellers competing with each other, to attract and retain users by offering better price and better QoS, ISPs face a tradeoff between pricing and investment in transit links. The profit maximization problem for a typical ISP is

$$\begin{aligned} max_{\bar{P},B}\ profit &= n \times \bar{P} - C(B) \\ s.t. \quad & \bar{P} \le P^* \\ & B \ge n\alpha(b^d + b^u) \end{aligned}$$

where α is the bandwidth consumption rate by non-P2P traffic. $b^d = \frac{1}{n}\sum_{i=1}^{n} b_i^d$ and $b^u = \frac{1}{n}\sum_{i=1}^{n} b_i^u$ are the average downlink and uplink bandwidth subscription by end users. The ISP optimal choice of monthly price ($\bar{P}$) and bottleneck link capacity (B) is subject to two constraints. First, the price cannot exceed users' willingness to pay. Second, the bottleneck link capacity has to be sufficient to accommodate bandwidth access demand by users.

Without loss of generality, we consider *duopoly*, a special case of *oligopoly* with only two competing ISPs. Let N be the size of the market that is the total number of users competed for by the ISPs. Assume the two ISPs are similar including cost structure and service tiers. To survive the competition, each ISP has two strategies to choose from: a lower price strategy with an inferior QoS, and a higher price strategy with a superior QoS. If the ISP pursues a lower price policy, it would attract more users *ceteris paribus*[2], but the increased traffic will require more investment in transit links and sacrificed QoS. If the ISP adopts a higher price policy, fewer users will lead to saved cost of transit links without sacrificing QoS. As profitability depends on both revenue and cost, the ISP has to weigh both to choose the optimal strategy.

Table I summarizes the payoff matrix of the two-ISP game with payoffs equaling to expected profit. As shown, the positions of the two ISPs in the game are symmetric, i.e., when two ISPs adopt the same strategy, N user will be equally divided between the two ISPs. If the two ISPs adopt different strategies, the lower-rate ISP would attract all users when users shop for price difference ex ante. To retain users ex post, the lower-rate ISP builds a bottleneck link capacity required to keep QoS. In Table I, P_l is the lower bound for an ISP to cut price without sacrificing profit, and $P_l < P^*$, B_0 is the minimum transit link necessary to accommodate $N/2$ users, and B_h is the minimum transit link necessary to accommodate all N users, i.e.,

$$B_0 = \frac{N}{2}\alpha(b^d + b^u)$$

$$B_h = N\alpha(b^d + b^u)$$

---

[1] Throughout the paper, the term "ISPs" refers to local ISPs, if not specified.

[2] All other things being equal or held constant.

TABLE I
PAYOFF MATRIX OF THE DUOPLISTIC GAME BETWEEN TWO ISPS.

ISP_1 \ ISP_2	lower price (P_l)	higher price (P^*)
lower price (P_l)	$P_l \times \frac{N}{2} - C(B_0)$, $P_l \times \frac{N}{2} - C(B_0)$	$P_l \times N - C(B_h)$, 0
higher price (P^*)	0, $P_l \times N - C(B_h)$	$P^* \times \frac{N}{2} - C(B_0)$, $P^* \times \frac{N}{2} - C(B_0)$

Clearly, $B_h = 2B_0$, and $C(B_h) > C(B_0)$. $C(B_h)$ is not necessarily twice as much as $C(B_0)$, depending on the cost structure of the ISP.

Since the game is symmetric, N users must be equally distributed between the duopolistic ISPs in equilibrium. A Nash equilibrium is a situation in which no player has anything to gain by changing only his or her own strategy unilaterally. The two pairs of strategies of equilibrium candidates are (P_l, P_l) and (P^*, P^*), depending on how profitability changes with pricing.

Having more users by lowering price does not necessarily mean earning more revenue since revenue equals price times quantity. The change in total revenue depends on how price sensitive users are. If demand is price inelastic, the total revenue would actually decrease at a lower price as the increased number of users is insufficient to cover the loss in per-user price, thus lowering price is not the optimal choice when subscription does not respond much to price changes. Raising price instead, not only raises total revenue, but also saves investment in the bottleneck capacity due to the smaller number of users to serve. In this case, the game would have no Nash equilibrium solution when the two ISPs do not cooperate. If ISPs cooperate by pursuing the same strategy of raising price, they can reach the maximum monthly rate of P^*. Clearly in the case of price-insensitive customers, (P^*, P^*) is the optimal solution for both ISPs. The profit level of each ISP will be:

$$\pi^* = P^* \times \frac{N}{2} - C(\frac{N}{2}\alpha(b^d + b^u))$$

which is the highest profit that the ISP may reap given the size of the market.

If users are price sensitive, lowering price would boost revenue by attracting disproportionately more users. If the increased revenue exceeds the cost of building excess bottleneck link capacity to accommodate the increased traffic, (P_l, P_l) will be the Nash equilibrium solution of the game. The profit level of each ISP will be:

$$\pi_l = P_l \times \frac{N}{2} - C(\frac{N}{2}\alpha(b^d + b^u))$$

since $P_l < P^*$, $\pi_l < \pi^*$.

We note that the possible Nash equilibrium solution to the game is not necessarily optimal for ISPs. The optimum would be for both ISPs to play "higher price", which is only achievable when the two ISPs cooperate. It is subject to the tension between cooperation and self-interest common for oligopolistic markets. When the game between the two ISPs is non-cooperative, the optimal outcome for both to charge higher prices cannot be realized, and (P_l, P_l) will be the game solution.

D. ISP business model with P2P challenges

The traditional ISP can be inversely affected by the introduction of P2P. Users of local ISP services are usually *clients* but not *servers*. Large companies that have a huge uplink demand for serving their content are not considered by those last-mile ISPs. If the content distribution load is suddenly transferred from servers to P2P clients, which now act as a server to redistribute content, it creates unwanted upload traffic for local ISPs.

Under the "all-you-can-eat" business model, P2P applications do not provide any specific ISP compensation, i.e., the increasing weight of P2P requests does not bring any extra revenue to ISPs. If the increasing P2P applications occupy more costly bottleneck capacity of ISPs, especially via uplinks, and force ISPs to expand capacity, then the profitability of an ISP would go down significantly. The issue of how a small group of super-users can occupy a much larger share of network resources is a key problem.

To prevent the fall in profitability caused by P2P, the natural thinking would be for the ISP to raise revenue or to lower cost. According to the game theoretical analysis from Section II-C, one possibility is when ISPs cooperate, they can avoid the destructive price war and both be better off by raising price to raise revenue and hence profit. However, there are two problems. First, it is illegal for ISPs to collude on prices. Second, pricing collusion will not solve the problem fundamentally. The question is *whether we can find a feasible solution without changing the current business model of ISPs.*

Whether P2P applications reduce the profitability of ISPs or not depends on whether P2P applications require more bottleneck link capacity of the ISP. We propose there is a critical value of local P2P presence in various circumstances to allow the co-prosperity of both P2P and ISPs. Before our derivation of such critical value (Section II-E), we need to define a few parameters first. There are a few considerations in evaluating the impact of P2P on an ISP that include:

- ρ: P2P penetration rate, measured by the P2P population as a percentage of the user population. Equivalently, it is the likelihood for an average user to use P2P rather than non-P2P applications.
- γ: The percentage of P2P requests originating from outside the ISP. Thus $(1 - \gamma)$ is the local serving rate.

E. Critical P2P penetration rate for ISPs

Let T be the amount of bottleneck traffic generated by a typical user, and its benchmark level in the absence of P2P traffic is $T_0 = \alpha(b^d + b^u)$, which is essentially the average bandwidth occupancy by non-P2P applications.

At the presence of P2P, total bottleneck traffic (combining non-P2P and P2P applications) generated by a typical user is:

$$T_p = (1-\rho)\alpha(b^d + b^u) + \rho\gamma(b^d + b^u)$$

Assuming α denotes non-P2P users using a fraction of their subscribed bandwidth, the largest transit link the ISP can afford, denoted as B_b, satisfies $n \times P^* = C(B_b)$ where n is the number of users of the ISP. Thus, B_b is the maximum affordable transit capacity for the ISP when it breaks even. It is socially optimal to reach this break-even point for ISPs because on one hand, break-even is sufficient to allow the ISP stay in business[3]. On the other hand, it allows for the maximum P2P usage (and other Internet usage) that is feasible. Note the presence of P2P does not change users' willingness to pay since the service package provides unlimited access anyway.

To accommodate all bandwidth requests, the ISP has to keep a bottleneck link capacity of $B' = n \times T_p$. Comparing B' and B_b, if $B' < B_b$, the ISP earns a profit; if $B' = B_b$, the ISP breaks even; if $B' > B_b$, the ISP suffers a loss. If we compare B' to the bottleneck link capacity in the base case, measured by B_0:

$$B' = \{1 + \frac{\rho(\gamma-\alpha)}{\alpha}\} \times B_0$$

In the presence of P2P applications, the change in the ISP's bottleneck link capacity depends on the relative size of γ and α, when the total content demand by users stays the same. If the P2P local serving rate is less than the extra bandwidth occupancy by P2P traffic than non-P2P traffic, i.e., $(1-\gamma) < (1-\alpha)$, P2P applications will lead to increased demand for bottleneck link capacity. In the case of neutral P2P applications, the P2P applications do not discriminate traffic request origins, the P2P local serving rate is close to zero for small local ISPs. As $\gamma \to 1$, more P2P participation will almost for sure impose bottleneck link pressure on the ISP, and hence reduce its profitability.

Equivalently, we can also derive the desired per-user transit link expansion by the ISP to accommodate P2P applications:

$$\frac{B' - B_0}{n} = \rho(\gamma-\alpha)(b^d + b^u)$$

The larger is the baseline profit, the more P2P traffic the ISP can support. Not surprisingly, the necessity for bottleneck link expansion is increasing in γ, and hence decreasing in the local serving rate $(1-\gamma)$. That is, if the local serving rate increases, the ISP would not have to expand bottleneck link capacity as much to accommodate the P2P traffic.

The highest sustainable P2P penetration rate, labeled as ρ', corresponds to $B' = B_b$, that is, it solves

$$n \times P^* = C(n \times ((1-\rho')\alpha + \rho'\gamma)(b^d + b^u))$$

Local ISPs typically get charged by transit providers based on the amount of traffic exchanged, i.e., $cost = C(B)$, as the cost function defined in Section II-A. The cost function is usually a piece-wise linear (non-decreasing) function, and total-volume based charging is in common use [10]. Let c be the per-volume price for the traffic exchanged between local ISPs and transit providers, then the amount that a local ISP owes to the transit provider is $c \times B$. Accordingly, the highest sustainable P2P penetration rate (ρ') is solved as follows:

$$\rho' = \frac{1}{\gamma-\alpha}(\frac{P^*}{c(b^d + b^u)} - \alpha)$$

When P2P is neutral, $\gamma \to 1$. The ISP is affected negatively by P2P applications, like in the current practice. As γ decreases, i.e., the local serving rate increases, the highest sustainable P2P penetration the ISP can support increases.

III. THE IMPORTANCE AND FEASIBILITY OF LOCALIZATION

As shown in the above discussion on how the localization of P2P traffic may affect the relationship between the ISP and P2P, P2P applications do not necessarily hurt the ISP. Rather, it depends on the local serving possibility. A higher local serving rate will allow the ISP to support more P2P penetration at break even. In this section, we explore further the role of local serving in affecting the bottleneck link capacity of the ISP and what the ISP can do to increase the local serving rate.

A. *Increase the likelihood of local serving*

Under the current practice of the ISP, since the local serving is limited, the sustainable P2P penetration ρ' is rather low. There is lack of mechanism for the convergence of actual P2P penetration toward ρ'. Hence, the impact of P2P on the local ISP is negative. Furthermore, ISPs have only limited effective methods to handle the increasing P2P traffic [11]. Therefore, increasing the local serving rate is critical and will help release the bottleneck capacity pressure on ISPs while they serve large volume of traffic flows.

In particular, if $T_0 = T_p$, the ISP would be indifferent between P2P and non-P2P applications. That is

$$\alpha(b^d + b^u) = \{(1-\rho)\alpha + \rho\gamma\}(b^d + b^u)$$

from which we derive $\gamma = \alpha$. In this simplified setup, if the percentage of P2P requests coming from outside the ISP equals to the bandwidth occupancy rate by non-P2P traffic, then P2P and non-P2P applications would be indifferent to the ISP in terms of total bottleneck traffic generated. If $\gamma < \alpha$, the bottleneck traffic would be reduced at the presence of P2P applications, which may give the ISP financial incentives to reward local users from caching content.

Researchers have suggested geographic-aware P2P such as P4P [8], where ISPs provide topology information to guide P2P peering and sharing. While a neutralized P2P model may have data sent across the Internet, P4P assigns priority to nearby peers. Having more P2P participating with the possibility of P4P may increase the chance of finding contents locally, but there is no guarantee that the local serving rate will increase proportionately to the size of local P2P population. Critically in most deployments, there is not a unified P2P construction and as a result, users of the same ISP may belong to disjoint

---

[3] In economics, when a business breaks even, opportunity costs have been covered and capital has received the risk-adjusted, expected return.

P2P groups such that the total number of P2P users may not be the sole determinant.

It is clear that the key determining factor is the local serving rate. One possible method we suggest could be what we refer to as P2P bots. The ISP creates P2P bots, stores and caches popular files, and installs various P2P protocols. The P2P bots controlled by the ISP would only serve P2P requests coming from inside with the ISP potentially blocking most requests from the outside. As a result, the P2P bots of the ISP would build good reputation among their own P2P users and rank low for outside P2P users. Thus, the likelihood for an inside user to download from a P2P bot run by the ISP is high, reducing the downlink traffic for the ISP without generating any additional uplink traffic.

The localization of P2P reduces total P2P bottleneck traffic for the ISP. We define β as the fraction of leftover P2P traffic after localization that still will be exchanged over the bottleneck link. Intuitively, a high local serving rate will result in a small β value. The after-localization P2P traffic per user, denoted by T_L, is then $T_L = \beta T_p$. We use the single (no history) locality model [12] to capture the reduction in P2P traffic after localization. The model considers that only one single copy of the content will need to be downloaded from outside the ISP. Once the initial copy has entered the ISP, content will be exchanged only among local peers. Therefore, β can be calculated as $\beta = \frac{D}{\sum_{k=1}^{D} M(k)}$ where D is the number of distinct P2P groups in which the ISP has clients. $M(k)$ is the number of peers from the ISP downloading content k. Suppose there is no overlap in P2P group members, and P2P population is evenly distributed among the P2P groups, then

$$\beta = \frac{D}{\rho \times n}$$

where ρn is the total P2P population of the ISP.

The fraction of leftover bottleneck traffic after localization is essentially the likelihood for a P2P request to come from outside the ISP after localization, i.e., $\gamma = \frac{D}{\rho n}$ after localization. Compared to $\gamma \to 1$ before localization, by how much can the P2P bottleneck traffic decrease is inversely related to the diversity of the P2P population and positively related to the total P2P population.

By having the ISP set up P2P bots, it becomes beneficial to all interested parties. The ISP gains by being able to partially reduce P2P usage of bottleneck capacity. P2P gains as it increases the quality and stability of P2P because the P2P bots run by ISPs are credible and uninterrupted when serving local users. As more ISPs adopt the practice, the overall quality of P2P will increase. More importantly, this method does not distort user behaviors and users may even gain since being served by P2P bots created by the ISP can be faster and more secure.

This practice is associated with positive externality. Suppose one ISP adopts the practice, its serving of local P2P download requests will reduce other ISPs' uplink traffic as their users are now less likely to receive requests originated from this ISP, thus reducing transit traffic for other ISPs. Similar analogues emerge such as with the NetFlix OpenConnect CDN and recent emerging peer agreements with NetFlix and local ISPs.

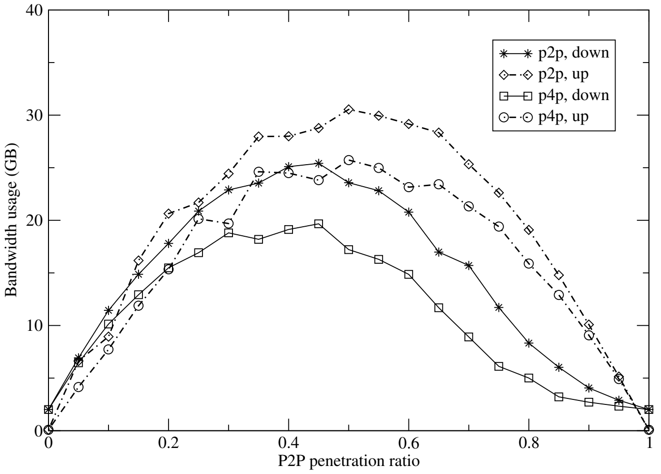

Fig. 1. The impact of P2P penetration rate in the local ISP on the total bandwidth usage on the ISP bottleneck link.

In a competitive environment, as described in the game theoretical analysis in subsection II-C, competition will provide incentives for all ISPs to adopt similar localization policy. When all local ISPs install P2P bots to serve their customers, then only those requests that cannot be fulfilled by ISPs flow out, reducing the overall traffic exchanged with transit ISPs. Therefore, installing P2P bots to serve local users actively is a viable option that is to the advantage of everyone involved.

B. Simulation study

To capture the dynamics and complexity of massive P2P networks, a discrete-event P2P network simulator was developed to complement the mathematical modeling analysis. The topology consists of N_p P2P users as well as N_0 non-P2P users in the local ISP domain, with a further division of local and foreign P2P users depending on whether they belong to local or foreign ISPs, determined by the parameter of local P2P penetration ratio ρ. The TCP packet-level simulation enables capturing the link queueing effect and congestion levels. The oversell index $ind_{oversell}$ determines the percentage of maximum subscribed users when they use their full bandwidth. The varying uplink and downlink bandwidth is assigned to last-mile users based on the metrics that the sum of total users' bandwidth will not exceed certain multiples of the ISP's bottleneck link capacity, i.e. $B = ind_{oversell} \cdot \sum_{i=1}^{N} b_i$. The file tables contain sharable files with varying sizes and follow a Zipf-like popularity distribution. Each file contains a number of chunks that are distributed across the network for concurrent downloading upon request in order to simulate a popular BitTorrent-like system.

Figure 1 shows the simulation result regarding the impact of an increasing number of P2P users on the local ISP's bottleneck link. As the ratio of local P2P users over foreign P2P users increases, the bottleneck link experiences significant increase in both down and up links. The changing user patterns by serving contents they download makes the uplink bandwidth requirement larger than downlink. It is interesting to observe that once the P2P penetration ratio passes approximately half, the bandwidth requirement actually drops as the increasing dominant P2P nodes acts as caching servers for the local ISP, and contents are more likely to be found locally, thus not hitting the pricey bottleneck link. P4P is also simulated by downloading chunks of files from a local peer as much as possible instead of randomly as in P2P cases. While P4P has lower bandwidth requirement in both up and down links, it follows a similar pattern of a bell-like shape.

IV. RELATED WORK

The tension between P2P and ISPs has attracted significant attention among researchers with a growing resurgence for new cloud constructions. For instance, the global impact of the load imposed by the Gnutella P2P system on the autonomous systems(AS)-level underlay was characterized and assessed [13]. Peer Coordination Protocol (PCP) [11] was developed to impose a dynamic rate limit on P2P traffic based on ISPs' constraint. Works on geographical-aware P4P attempt to explore the interest conflicts between ISPs and P2P in a technical way by changing how P2P clients connect with each other [8]. The ALTO (application-layer traffic optimization) approach [14] studies how to provide the right type of network layer topology information to the requesting P2P clients to make P2P data structure efficient. Open research issues have been identified as the need for layer cooperation as a solution to the ALTO problem. In particular, the study in [15] summarized challenges and design issues for the ALTO approach and highlighted the status quo of ALTO standardization. Different from a technical solution, researchers [16] modeled the peering and routing tussle between ISPs and P2P applications and analyzed the economic implications on overlay traffic on ISPs' peering decisions.

Complementary to our economic analysis, there has been work quantifying the impact of P2P localization on ISPs. Cuevas et al. [17] conducted a large scale measurement study of BitTorrent demand demographics, and studied to what extent localization can improve the performance of users, and reduce the amount of transit traffic. Le Blond et al. [18] performed extensive experiments driven by a BitTorrent crawl to evaluate the impact of high locality on inter-ISP links traffic and peers download completion time. Seibert et al. [12] conducted a simulation study to evaluate how localizing P2P traffic within network boundaries will change the profitability of an ISP, and found that the benefits of localization depends on many factors, and the profitability of ISPs only increases when ISPs cooperate with each other.

V. CONCLUSION

As the amount of data transferred over the Internet increases exponentially in the big data era, ISPs face bottleneck bandwidth pressure, in particular from P2P applications. However, the impact of P2P on ISPs has not been well understood. Through modeling analysis and simulation study, we found more P2P users may not necessarily be a bad thing for ISPs. We illustrate the direct impact of the key determining factor, the local serving rate of P2P, on the bottleneck link bandwidth and the business model of ISPs. We also explore the feasibility of potential methods to increase the local serving rate. Our equilibrium solutions suggest it would be mutually beneficial for ISPs to localize P2P traffic, which will allow the co-prosperity of both ISPs and P2P networks.

REFERENCES

[1] M. Parameswaran, A. Susarla, and A. B. Whinston, "P2P networking: an information sharing alternative," *Computer*, vol. 34, no. 7, pp. 31–38, 2001.

[2] Sandvine, "Global Internet phenomena report," vol. 2H, 2013.

[3] C. G. Pablo Rodriguez, See-Mong Tan, "On the feasibility of commercial, legal p2p content distribution," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 1, pp. 75–78, January 2006.

[4] X. Zheng, C. Cho, and Y. Xia, "Optimal peer-to-peer technique for massive content distribution," in *Proceedings of the 27th Conference on Computer Communications*, Phoenix, AZ, April 13-18 2008, pp. 151–155.

[5] G. Neglia, G. Reina, H. Zhang, D. Towsley, A. Venkataramani, and J. Danaher, "Availability in bittorrent systems," in *Proceedings of the 26th IEEE Conference on Computer Communications*, Anchorage, AK, May 6-12 2007, pp. 2216–2224.

[6] C. Labovitz, "The rise and fall of p2p," *$DeepField$ Networks*, April 10 2012.

[7] S. Schwartz and T. McElligott, "Usage-based billing back on agenda," *Billing & OSS World*, Feb. 20 2008.

[8] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. Liu, and A. Silberschatz, "P4P: Provider portal for applications," in *Proceedings of SIGCOMM*, Seattle, WA, August 17-22 2008.

[9] S. Shakkottai and R. Srikant, "Economics of network pricing with multiple ISPs," *IEEE/ACM Transactions on Networking*, vol. 14, no. 6, pp. 1233–1245, 2006.

[10] D. K. Goldenberg, L. Qiu, H. Xie, Y. R. Yang, and Y. Zhang, "Optimizing cost and performance for multihoming," in *Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, Portland, OR, August 30-September 03 2004, pp. 79–92.

[11] H. Xie, Y. R. Yang, and A. SilberSchatz, "Towards an ISP-compliant, peer-friendly design for peer-to-peer networks," in *Proceedings of the 7th International IFIP-TC6 Networking Conference*, Singapore, May 5-9 2008, pp. 375–384.

[12] J. Seibert, R. Torres, M. Mellia, M. M. Munafo, C. Nita-Rotaru, and S. Rao, "The Internet-wide impact of p2p traffic localization on ISP profitability," *IEEE/ACM Transactions on Networking*, vol. 20, no. 6, pp. 1910–1923, December 2012.

[13] A. H. Rasti, R. Rejaie, and W. Willinger, "Characterizing the global impact of p2p overlays on the AS-level underlay," in *Proceedings of the 11th International Conference PAM 2010*, Zurich, Switzerland, April 7-9 2010, pp. 1–10.

[14] V. Gurbani, V. Hilt, I. Rimac, M. Tomsu, and E. Marocco, "A survey of research on the application-layer traffic optimization problem and the need for layer cooperation," *IEEE Communications Magazine*, vol. 47, no. 8, pp. 107–112, August 2009.

[15] J. Seedorf, S. Kiesel, and M. Stiemerling, "Traffic localization for p2p-applications: The ALTO approach," in *Proceedings of IEEE Ninth International Conference on Peer-to-Peer Computing*, 2009, pp. 171–177.

[16] J. H. Wang, D. M. Chiu, and J. C. Lui, "Modeling the peering and routing tussle between ISPs and p2p applications," in *14th IEEE International Workshop on Quality of Service*, 2006, pp. 51–59.

[17] R. Cuevas, N. Laoutaris, X. Yang, G. Siganos, and P. Rodriguez, "Deep diving into BitTorrent locality," in *Proceedings of 2011 IEEE INFOCOM*, Shanghai, China, April 10-15 2011, pp. 963–971.

[18] S. L. Blond, A. Legout, and W. Dabbous, "Pushing BitTorrent locality to the limit," *Computer Networks*, vol. 55, no. 3, pp. 541–557, February 21 2011.