The Internet-Wide Impact of P2P Traffic Localization on ISP Profitability

Original

The Internet-Wide Impact of P2P Traffic Localization on ISP Profitability / Jeff, Seibert; Ruben, Torres; Mellia, Marco; Munafo', MAURIZIO MATTEO; Cristina Nita, Rotaru; Sanjay, Rao. - In: IEEE-ACM TRANSACTIONS ON NETWORKING. - ISSN 1063-6692. - STAMPA. - 20:6(2012), pp. 1910-1923. [10.1109/TNET.2012.2190093]

Availability:

This version is available at: 11583/2502287 since:

Publisher:

IEEE / Institute of Electrical and Electronics Engineers Incorporated:445 Hoes Lane:Piscataway, NJ 08854:

Published

DOI:10.1109/TNET.2012.2190093

Terms of use:

openAccess

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

The Internet-wide Impact of P2P Traffic Localization on ISP Profitability

Jeff Seibert*, Ruben Torres*, Marco Mellia[†], Maurizio M. Munafo[†], Cristina Nita-Rotaru*, Sanjay Rao*

*Purdue University [†]Politecnico di Torino

{jcseiber@cs, rtorresg@ecn, crisn@cs, sanjay@ecn}.purdue.edu {mellia, munafo}@polito.it

Abstract—We conduct a detailed simulation study to examine the impact of localizing P2P traffic within network boundaries on an ISP’s profitability. A distinguishing aspect of our work is the focus on Internet-wide implications, i.e., how adoption of localization within an ISP affects both itself and other ISPs. Our simulations are based on detailed models of inter-AS P2P traffic and inter-AS routing, localization models that can predict the extent to which P2P traffic is reduced, and pricing models to predict the impact of changes in traffic on an ISP’s profit. To evaluate our models we use a large-scale crawl of BitTorrent involving over 138 million users sharing 2.75 million files. Our results show that the benefits of localization must not be taken for granted. Some of our key findings include: (i) residential ISPs can actually lose money when localization is employed and some will not see increased profitability until other ISPs employ localization; (ii) the reduction in costs due to localization will be limited for small ISPs and tends to grow only logarithmically with client population; and (iii) some ISPs can better increase profitability through alternate strategies to localization by taking advantage of the business relationships they have with others.

I. INTRODUCTION

The last decade has seen a rapid growth in popularity of peer-to-peer (P2P) systems, spanning diverse applications such as content distribution (e.g., BitTorrent, eMule, Gnutella), video streaming (e.g., PPLive, Coolstreaming), and audio conferencing (e.g., Skype). However, the success of these applications and the consequent growth in P2P traffic has raised concerns among Internet Service Providers (ISPs), which have to pay a high cost for carrying traffic while receiving little revenue. While there is evidence that P2P traffic is decreasing [1], it is still today a significant fraction of the Internet traffic (more than 18% according to [1] and more than 50% in some of our datasets) and it is perceived as wasteful of network resources, and in particular, of expensive peering link bandwidth. In order to reduce these costs, different P2P localization techniques have been proposed [2]–[8]. The key idea behind these techniques is to limit the amount of traffic entering the ISP by enforcing a preference in exchanging content among peers in the same ISP.

Several works have shown the benefits of localization for both users and providers [2, 4], while other works question the possible benefits for users [9]. However, all previous studies consider a partial view of the problem, e.g., by showing the benefits for a single Autonomous System (AS)/Internet Service Provider (ISP) or running a limited set of experiments involving different ASes. Therefore, it is unclear whether localization is necessarily beneficial to all ASes, how the adoption of

localization by one AS impacts other ASes, and how the traffic carried by various ASes is altered as localization techniques are widely adopted.

Evaluating the impact of localization policies when applied on an Internet-wide scale is a challenging task given the complexity of the Internet. As ASes play various roles from a business point of view, they may experience different effects from the use of localization policies. For example, some ASes (referred to as *residential ASes*), provide Internet service to end-users, and P2P clients are found in these ASes. Other ASes (referred to as *transit ASes*) provide the service of connecting other residential and transit ASes together. However, many transit ASes also provide residential services, and a clean separation between the two types does not exist today. From a business point of view, ASes form “customer-provider” relationships, where a customer AS will pay for the service a provider AS offers, or “peering” relationships, where two ASes will agree to carry each others traffic for free.

Given the current structure of the Internet, localization of traffic is intuitively beneficial for purely residential ASes, and it will have a negative impact on the revenues of purely transit ASes. However, we have found that over 1,200 residential ASes also provide transit service to at least one other AS. Thus, for many ASes it is not obvious how localization may impact them. In addition, as the ultimate goal of ASes is cutting costs and increasing revenue, there are alternatives to the simple localization of traffic inside an AS that have not been explored in previous work. For example, to avoid paying for traffic, ASes could prefer to exchange traffic with peering ASes. Furthermore, to increase revenue, ASes could prefer to push traffic to customers’ ASes and avoid providers’ ASes.

In this paper, our goal is to gain deeper insights into such Internet-wide implications of P2P traffic localization on ISP profits, and develop simulation methodologies to systematically explore the issues. We explicitly focus our work on the benefits and drawbacks for ISPs, though we note that the use of localization can also impact user performance. Our simulations are based on detailed models of (i) inter-AS P2P traffic; (ii) inter-AS routing models; (iii) models that can predict the extent to which P2P traffic gets localized; and (iv) pricing models to predict the impact of changes in traffic on ISP profit.

To model inter-AS P2P traffic, we leverage [10]. This is perhaps the only inter-AS traffic model that is available today, in contrast to intra-AS traffic which is widely studied. We present refinements to the model presented in [10], which we show can result in better model accuracy, and employ

the refined model in our simulations. The model requires the knowledge of the P2P population in each AS as input, which we estimate considering BitTorrent, one of the most popular and widely used P2P systems. Our estimation is based on crawls of a popular tracker from which we gathered a dataset of over 138 million BitTorrent peers participating in 2.75 million torrents. While our evaluations are based on BitTorrent, our methodologies are general, and apply to other P2P systems as well.

Conducting our simulation study requires models that can predict the reduction in P2P traffic entering/exiting an ISP when localization techniques are employed. The possible traffic reduction depends on a wide range of factors including (i) the population of peers inside an AS, (ii) the extent to which peers download similar content, and (iii) the upload capacities of peers inside the ISP relative to those outside [11]. Rather than focusing on a specific localization model, we conduct a sensitivity analysis to a range of models.

As a last step, translating a change in traffic volumes into a change in profits for the ISP is a challenge. Typical pricing models in ISPs are based on the 95th percentile of traffic volumes [12], with the price per Mbps itself showing significant geographical variation. Further, the pricing models depend on total volumes (of traffic across all applications) rather than P2P traffic volume alone, which is unavailable to us. Therefore, we consider multiple pricing models and develop conservative and optimistic predictions of the change in profits for an ISP.

Armed with these models, we seek to answer several questions such as: (i) Do ASes necessarily benefit by employing localization? How significant are the benefits? (ii) How is the profitability of various ASes impacted if localization policies are adopted by an increasing fraction of ASes at the same time? What is the impact of global adoption of such policies? (iii) Are there any better policies that can be more profitable to some ASes than a simple localization policy? Given the complexity of the real-world factors that our models seek to capture, there are unavoidable simplifications that must be made. Thus, rather than “absolute” answers to these questions for specific “point-models”, our focus is on understanding the sensitivity of our results, and how the trends change with various localization and pricing models.

Our results show that the benefits of localization must not be taken for granted. Some of our key findings include: (i) residential ISPs can actually lose money when localization is employed and some will not see increased profitability until other ISPs employ localization; (ii) the reduction in costs due to localization will be limited for small ISPs and tends to grow only logarithmically with client population; and (iii) some ISPs can better increase profitability through alternate strategies to localization by taking advantage of the business relationships they have with other ISPs. Overall, we believe our findings have important implications for ASes, and both our findings, as well as the methodologies and models that we develop in this paper are important contributions in their own right.

The remainder of the paper is organized as follows: Section II introduces our P2P inter-AS traffic model and its validation. Section III and Section IV discuss different localization

policies and the pricing models we use in the paper. Section V and VI show our findings under different localization scenarios. We review related work in Section VII. Finally, main findings of the paper are summarized in Section VIII.

II. MODELING INTER-AS P2P TRAFFIC

The first building block needed for our methodology is a model describing an inter-AS P2P traffic matrix. In the past, the gravity model has been used to model both intra-AS [13, 14] and inter-AS [10] traffic matrix estimation. Below we first review the model in [10], which we refer to as the *Gravity* model, then propose a new refinement to improve P2P traffic prediction accuracy, which we refer to as the *Affinity* model.

A. The Gravity Model

Inter-AS traffic demand has been modeled only once before in the work by Chang et al. [10]. Chang builds on previous work by applying the well-established gravity model to an inter-AS setting. To separately account for P2P and web traffic, their model was split accordingly into two different gravity models, of which we will only focus on the P2P model. In the Gravity model, the traffic X_{ij} sent from AS i to AS j is defined as follows:

$$X_{ij} = \frac{f(R_{RA}(i))f(R_{RA}(j))}{R_{BA}(i, j)^\beta}, \quad (1)$$

where f is the monotonically decreasing function $f(x) = 1/x$, $R_{RA}(i)$ is the rank of AS i in the list of ASes sorted by decreasing peer population, and $R_{BA}(i, j)$ is the rank of the bottleneck AS between i and j in the sorted list of ASes by capacity (the bottleneck AS is the smallest transit AS that is on the AS path between i and j). This model stems from the intuition that the higher the population of peers in an AS (i.e., the higher is its rank), the larger the aggregate of traffic the AS exchanges. In addition, if the path between two ASes has little capacity, then the amount of traffic will be consequently reduced. β is a parameter that is used to better weight the effect of bottlenecks along the path. In [10] $\beta = 0.1$ is suggested, which makes the bottleneck bias almost negligible. This implicitly suggests that the volume of P2P traffic exchanged between ASes is mainly driven by the peer population of each AS.

B. The Affinity Model

Given the world-wide nature of the Internet and its diversity of users and available content, intuition suggests that P2P traffic will be driven not only by the population size of ASes, but also by the cultural and linguistic makeup of the users inside the ASes, or the “*affinity*” between ASes. Thus, if peers of two ASes are not interested in the same content, the traffic exchanged among them will be marginal, even if the number of peers they have is large. For example, if AS-1 and AS-2 are located in Italy, and AS-3 is located in China, it is expected that large traffic will be exchanged between AS-1 and AS-2, while little traffic will flow between AS-3 and AS-1, AS-2.

We estimate the affinity between ASes using the cosine similarity distance [15]. The cosine similarity results in a value

between 0 (no similarity) and 1 (perfect similarity) that is the cosine of the angle between two vectors $\bar{V}_i$ and $\bar{V}_j$, i.e.,

$$Cos(i, j) = \frac{\bar{V}_i \cdot \bar{V}_j}{\|\bar{V}_i\| \|\bar{V}_j\|}. \quad (2)$$

In our case, each vector $\bar{V}_i$ represents the “content distribution” in AS i , whose components report the number of peers interested in a given content that are present in AS i . Thus, if two ASes have many peers interested in the same content, then they will have high affinity. Completing the previous example, consider as content an Italian movie, a Chinese song, and an English book. Assuming $\bar{V}_1 = (10, 1, 3)$, $\bar{V}_2 = (100, 2, 30)$ and $\bar{V}_3 = (0, 10, 3)$, we have $Cos(1, 2) = 0.997$ while $Cos(1, 3) = 0.173$, which reflects the intuition that Italian ASes prefer to exchange traffic among themselves rather than with the Chinese AS.

Once we have calculated the affinity between two ASes we can combine that with the peer population in each AS to form a gravity model. Thus, we define our Affinity model as follows:

$$X_{ij} = P(i)P(j)Cos(i, j), \quad (3)$$

where $P(i)$ and $P(j)$ are the population of AS i and AS j .

While the Affinity model could include a preference related to the upload capacity of peers, we chose to include only the affinity among ASes due to client’s interest in the same content. We superpose a bias in peer selection due to performance as part of the locality models in Section III-B.

C. Model Validation

In this section we first describe the datasets we use as input to the Affinity model and also use throughout the paper. We then present results that compare our model with the Gravity model.

1) Datasets:

BitTorrent crawl snapshots: The Affinity model requires as input the peer population $P(i)$ and the content distribution vectors $\bar{V}_i$. To estimate them, we rely on active measurements obtained by crawling a very popular BitTorrent tracker named “OpenBitTorrent” [16]. As the tracker is not associated with a particular torrent publishing web site and it provides an easy way for users to publish content, it attracts users from all over the world.

We took snapshots of BitTorrent activity every hour for a period of 8 days during May 2010. A total of 192 different snapshots have then been collected, which will be used throughout this paper. In each snapshot, we crawled all torrents that had at least one active downloader and for every torrent we requested peers from the tracker until we received at least 95% of all participating peers. Since many users are behind NATs, we consider a peer to consist of a unique (IP, port) combination. This allows us to obtain information about which peers are actually participating in which torrent, i.e., the peer population by content. To obtain the peer population per AS, we map IP addresses to the corresponding AS by using the service provided by Team Cymru [17]. At the end, for every snapshot we obtain for each AS i the population $P(i)$ and

content distribution $\bar{V}_i$, which allow us to compute 192 global AS level traffic matrices.

While a detailed characterization of these BitTorrent datasets is out of the scope of this paper, we briefly summarize their size which reflects their generality. A normal snapshot consists of over 5 million peers, 154 countries, 12,000 ASes, and 1 million torrents. Over the 8 days, we saw more than 138 million distinct peers in over 2.75 million torrents. One interesting finding about the dataset which will be instrumental later is the fact that each torrent population size follows a heavy tailed distribution with a small portion of very large torrents, but also a large number of torrents with less than 100 peers. Peer distribution over ASes is instead more biased toward larger ASes which host most of the peers, e.g., the largest 1,600 ASes account for 97% of peers.

Inter-AS topology and routing: The knowledge of the AS paths is instrumental to predict the volume of traffic on individual inter-AS links. Besides, they are also necessary to compute $R_{BA}(i)$ for the Gravity model. First, we need a map of the AS topology which includes the business relationships between ASes. We use CAIDA’s AS map [18] augmented with peering edges from recent research on mapping Internet Exchange Points (IXPs) [19]. Second, we need to know the AS-level routing. To this end, we use the algorithm proposed by Qiu et al. [20] to determine valley-free paths between residential ASes. Qiu et al.’s algorithm uses Routing Information Bases (RIBs) alongside the AS topology to determine the most likely route between ASes. We use RIBs provided by the Oregon’s RouteViews Project [21] that are from the LINX, KIXP, PAIX, and Equinix Ashburn IXPs. This set of routing table dumps represents over 329,000 prefixes from 33,910 ASes.

Leveraging on the fact that the top 1,600 ASes alone account for 97% of all the P2P traffic that is generated by the Affinity model, we limit our evaluation to only the subset of top 1,600 residential ASes according to peer population as seen in our BitTorrent crawl. We also include all the transit ASes that belong on any AS path between these residential ASes, for a total of 2,067 ASes. In this paper, we define a *residential* AS as having at least one peer in the BitTorrent crawl and a *transit* AS as having at least one customer AS in the CAIDA map. Thus, an AS could be both a residential and a transit AS. More details about the ASes are deferred to Sec.V.

Packet traces from large ISP datasets: To verify the accuracy of the Affinity model traffic prediction, we compare its output against packet-level traces from six vantage points scattered in the US and across three different European countries. Each vantage point monitors several thousands of users. For convenience, we name the vantage points ISP-1 to ISP-6. For each ISP, all packets going to and coming from all the hosts in the Points of Presence (PoP) were passively monitored for several months. An advanced traffic classification tool based on deep packet inspection and advanced statistical classifiers [22] was used to produce the per-application volume of traffic sent by hosts in the PoP to each different AS, i.e., an actual row of the traffic matrix for each given application.

2) *Comparing Models:* In Fig. 1, we focus on a one-day long trace from ISP-1 and a one-hour long trace from ISP-2. We use the BitTorrent snapshots that refer to the same time of

Fig. 1. Affinity and Gravity models compared against real measurements.

Fig. 2. Relative error for Affinity and Gravity models

day that the ISP traces are from. Similar results were obtained for the rest of the traces and are not shown due to space constraints. For each graph, we report the volume of traffic sorted in decreasing order, considering actual measurements (solid line), the Gravity model prediction (small dot line), and the Affinity model prediction (large dot line). As both the Gravity and Affinity models produce unit-less output, we scale them and the ISPs' measured traffic volumes so they are comparable, to a standard unit-less metric by minimizing the mean square error. To also demonstrate how good the results are, we show the corresponding relative error values of both models in Fig. 2.

Fig. 1(a) refers to BitTorrent traffic as seen from ISP-1. The Gravity and Affinity models are very similar until rank 300, at which point the Gravity model severely overestimates the traffic demand, while the Affinity model better captures the sudden decrease of traffic sent by ISP-1 clients to the smaller ASes. Similarly, we compare BitTorrent traffic seen from ISP-2 in Fig. 1(b). Again, the Affinity model is able to better match the traffic demand trend for most ASes, while the Gravity model shows a much more regular slope, clearly missing the content bias induced on exchanged traffic. For the relative error values in Fig. 2(a) and 2(b), the Gravity model is only very accurate for 50% of ASes, while the Affinity model is accurate for 70% and 60% of ASes, respectively.

Finally, to show that the Affinity model is not specific to BitTorrent but can be generally applied to other P2P protocols, Fig. 1(c) shows results considering traffic volumes sent by ISP-1 clients, but using eMule as the P2P application. The same cosine similarity values as obtained from the BitTorrent snapshots are used, since the cosine similarity values catch the cultural and linguistic interests of peers, and are not expected to change across different P2P systems. The per AS eMule population has been estimated from the eMule traffic in ISP-1 instead. Also in this case, results show that the Gravity model

overestimates the actual traffic sent to each AS, while the Affinity model closely matches the traffic demand even up to high ranking values. This difference is seen in Fig. 2(c), where the Gravity model is only very accurate for 30% of ASes, but the Affinity model is accurate for 60% of ASes.

We have conducted other experiments to verify the goodness of the Affinity model, considering different times of the day, different days, transmitted and received traffic, different P2P systems and different crawls from different trackers. In all cases, the Affinity model provided more accurate estimates than the Gravity model. Moreover, the cosine similarity proved to be very robust, so that it can be used to model several P2P applications like BitTorrent or eMule.

III. MODELING P2P LOCALIZATION

In this section we present models to predict the reduction in P2P traffic exchanged by an ISP when localization techniques are employed. A single model may not be sufficient because P2P traffic reduction depends on a variety of factors such as (i) the population of peers inside an AS, (ii) the extent to which peers download similar content, and (iii) the upload capacities of peers inside the ISP relative to those outside [11]. Hence, we consider a set of locality models and show sensitivity to them. Validation of these models is a difficult task, since this requires measuring P2P traffic aggregates from a large number of ISPs around the world, from different ISP categories (e.g., residential and transit), and with different upload capacities of clients. Instead, in later sections, we show trends of the impact that P2P localization may have on ISPs and perform extensive sensitivity analysis to the various localization policies.

We determine the ratio of traffic received by an AS j after localization versus before localization, which we call α_j . In other words, α_j is the fraction of leftover traffic after localization that still will be received by peers in AS j . Intuitively, a good localization policy will result in a small

α_j value. The traffic L_{ij} sent by peers in AS i to peers in AS j after localization is then simply:

$$L_{ij} = \alpha_j X_{ij}. \quad (4)$$

As we have multiple snapshots from which we generate traffic matrices, we also calculate α_j for each snapshot. For simplicity though, we drop the explicit notation on time in the following.

A. System Architecture Assumptions

We assume that there is a localization technique in place that allows peers to find other peers that are in the same AS. For example, peers contact an “oracle” which allows them to obtain an ordered and possibly filtered list of peers interested in the same content. Peers then start exchanging data with the suggested peers according to the P2P trading algorithm. Individual ASes can impose localization of traffic independently of what other ASes do, e.g., some may deploy an oracle, others may not. This scenario is compatible with both the P4P iTracker [2] and the IETF ALTO [23] proposals.

We further assume that an AS cannot influence peers outside of its own AS, so that external peers can still connect to and download from internal peers, i.e., an AS cannot stop external peers from downloading content from peers within the AS. This implies that transit ASes do not deploy traffic shaping on traffic that does not originate from their own AS, but only rely on the oracle to enforce localization policies. Furthermore, this implies there is some altruism in the system, so that clients in an AS that do not localize traffic can still receive the content, even if every other AS does localize traffic. Therefore, for an AS that does localize, its outgoing P2P traffic can be greater than its incoming traffic.

B. Locality Models

• **Single(no history):** This is a pessimistic scenario where for every crawl, the file must be downloaded again by every peer from outside the AS. That is, there are no internal seeds available. The model computes the leftover traffic assuming only *one single copy* of the content will need to be downloaded from outside the AS. Once the initial copy has entered the AS, content will be exchanged only among local peers. For example, assume there are $P_j(k) = 10$ peers from AS j downloading content k ; when localization is used, only one copy would need to be downloaded, resulting in $1/P_j(k) = 0.1$ leftover traffic. Thus, the more popular a piece of content is, the less leftover traffic there will be. Given a snapshot, for every AS j that has clients in N_j distinct torrents, we estimate α_j as follows:

$$\alpha_j = N_j \frac{1}{\sum_{k=1}^{N_j} P_j(k)}. \quad (5)$$

• **Single(history):** This is a more realistic model where we consider that the first time a peer appears in a torrent in our crawls, it is considered a leecher, and if it appears again in that torrent in our crawls we consider it to be a seeder. To find out how sensitive α_j actually is to content availability, we simply keep track of what peers have been in which torrents over

time. For a given torrent, consider a peer that has been seen at time slot t for the first time. When it reappears in time slot $t' > t$, it is considered a seed. That is, if the peer has been in a torrent in the past, it is marked as a seed in future time slots. Formally, for a snapshot t and AS j , let $S_j(k)$ be the number of seeds in torrent k and let T_j be the number of torrents that have some seed in them. We can then calculate α_j with the following equation:

$$\alpha_j = \frac{N_j - T_j}{\sum_{k=1}^{N_j} (P_j(k) - S_j(k))}. \quad (6)$$

• **Single(persistent):** This is an optimistic scenario, where once a single peer inside an AS download a file, then no other peer inside the AS will need to download from outside the AS again, since the initial peer remains as a seeder for everyone else. It has been downloaded once from the outside at time slot t in which the first peer is interested in the content k . We use Equation 6 to calculate α_j for this model, but assume at least one seed is always present for each time slot $t' > t$. This results in an optimistic localization scenario.

• **Perf(no history):** As peers might prefer to download from nodes with a higher upload capacity than those inside its AS we extend the previous policies to *include a performance bias*. A peer prefers to download content from peers in its own AS, unless there exists external peers with much higher upload capacity. A similar policy has been examined in [11]. We compute $E(j, k)$, the expected number of copies of content k downloaded from outside AS j .

$$E(j, k) = P_j(k) \frac{U(j, k)}{P(k)}, \quad (7)$$

where $P(k) = \sum_j P_j(k)$ is the total number of peers interested in content k , and $U(j, k)$ is the number of external peers interested in content k that have an average upload capacity X times higher than peers in AS j . By averaging over all content in which AS j participates we have:

$$\alpha_j = \frac{\sum_{k=1}^{N_j} \max(E(j, k), 1)}{\sum_{k=1}^{N_j} P_j(k)}, \quad (8)$$

where $\max(E(j, k), 1)$ states that at least one copy must be downloaded.

For the evaluation of this scheme, we use the iPlane [24] dataset, which provides an estimate of the access bandwidth of several tens of thousands of /24 networks in the Internet. To account for factors that could make the real and the estimated capacities differ, such as congestion of intermediate links, we select a remote peer over a local peer only if the access bandwidth of the remote peer is 10 at least times higher than the bandwidth of the local peer. Further, any remote peer for which we do not have access bandwidth information will not be preferred over a local peer.

• **Perf(history):** Similar to the previous model, if an internal seed exists at time t , then peers do not need to download anything from outside the AS. The following equations are used to calculate α_j :

$$E(j, k) = (P_j(k) - S_j(k)) \frac{U(j, k)}{P(k)} \quad (9)$$

$$\alpha_j = \frac{\sum_{k=1}^{T_j} E(j, k) + \sum_{k=T_j+1}^{N_j} \max(E(j, k), 1)}{\sum_{k=1}^{N_j} (P_j(k) - S_j(k))}, \quad (10)$$

• **Perf(persistent):** We again assume that content persists forever after being downloaded once from outside the AS. We use Equation 10 to calculate α_j for this scenario, but assume a seed always persists after downloading.

IV. MEASURING ISP PROFITABILITY

The total profitability of an ISP depends on many factors. In this paper we focus on the portion of the profit/expenses that are related to money gained/paid by carrying P2P traffic only. We define our ideal metric for achieving this goal and describe how we evaluate it using our pricing models.

A. An Ideal Metric for ISP Profitability

A customer ISP i typically gets charged by a provider ISP j based on the 95th percentile (P95) volume of traffic exchanged on an individual link [12]. This is done by sampling the inbound and outbound volume of traffic every 5 minutes for the duration of a billing period, which is usually 30 days. Let $V_{ij}(t)$ and $V_{ji}(t)$ respectively denote the outbound and inbound volumes for ISP i at time instant t . After sorting these values, P95 is chosen from both the outbound and inbound traffic; let these terms be denoted as $P95(V_{ij})$ and $P95(V_{ji})$. Let CV_{ij} be the charging volume, which is the actual volume that charges are computed on. Typically, CV_{ij} is determined by taking the maximum of the inbound and outbound P95s:

$$CV_{ij} = \max(P95(V_{ij}), P95(V_{ji})). \quad (11)$$

However, in some cases it is determined by taking the average:

$$CV_{ij} = (P95(V_{ij}) + P95(V_{ji}))/2. \quad (12)$$

CV_{ij} is then used as input to a pricing function, which is typically non-decreasing, the output of which is a dollar amount that the customer owes the provider. Assuming a linear pricing function (see Sec IV-C for more details), let p_{ij} be the price per Mbps for the link between i and j , then the amount that ISP i owes ISP j is $p_{ij}CV_{ij}$.

Let $\mathcal{P}_i$ and $\mathcal{C}_i$ denote the set of providers and customers of ISP i , respectively. Then, the profit of the ISP i prior to localization is $\sum_{k \in \mathcal{C}_i} p_{ik}CV_{ik} - \sum_{k \in \mathcal{P}_i} p_{ik}CV_{ik}$.

Thus far we have considered the profit with respect to a certain set of traffic volumes. However, if these traffic volumes change due to P2P localization policies, we can also calculate the increase in profits after this occurs. Formally, let $\delta(x)$ denote the change in a variable x when localization is employed. Then,

$$\delta(\text{profit}) = \sum_{k \in \mathcal{C}_i} p_{ik}\delta(CV_{ik}) - \sum_{k \in \mathcal{P}_i} p_{ik}\delta(CV_{ik}). \quad (13)$$

To study how localization affects profit due to P2P traffic, we normalize $\delta(\text{profit})$ to the profit before localization that is attributed to P2P traffic ($\text{profit}_{p2p, \text{before}}$), i.e., profit is

computed exactly as before, except only the portion of traffic that is P2P is considered. Thus, we have:

$$\text{profit increase} = \frac{\delta(\text{profit})}{\text{profit}_{p2p, \text{before}}}. \quad (14)$$

Finally, if $\text{profit}_{p2p, \text{before}}$ is negative (i.e., the ISP is originally losing due to P2P traffic), we simply normalize by the loss instead of the profit. Thus, if $\text{loss}_{p2p, \text{before}} = -\text{profit}_{p2p, \text{before}}$, we have:

$$\text{loss reduction} = \frac{\delta(\text{profit})}{\text{loss}_{p2p, \text{before}}}. \quad (15)$$

B. Approximating ISP Profitability

Ideally, Equation 13 should be evaluated considering the total amount of traffic flowing across links. Unfortunately, modeling total inter-AS traffic is a hard problem. To the best of our knowledge, only [10] addressed this problem. However, that model is not easily applicable to our context as it assumes the ratio of P2P to other traffic is known for all ASes which varies widely and is difficult to ascertain.

To handle this, we approximate the ideal metric by assuming that the change in P95s of total traffic volume on localization is the same as the change in P95 of P2P traffic volumes on localization, in each of the inbound and outbound directions. Formally, let $V_{p2p, ij}(t)$ and $V_{p2p, ji}(t)$ respectively denote the inbound and outbound volumes of P2P traffic that ISP i sends to or receives from ISP j at time instant t . Then, we assume that $\delta(P95(V_{ij})) = \delta(P95(V_{p2p, ij}))$, and $\delta(P95(V_{ji})) = \delta(P95(V_{p2p, ji}))$. With this assumption, the approximate change in charging volume on localization is simply computed as follows: (i) if charging volumes are computed based on the maximum, as in Equation 11, then $\delta(CV_{ij}) = \delta(P95(V_{p2p, ij}))$ or $\delta(P95(V_{p2p, ji}))$, depending on whether the AS is inbound dominated or outbound dominated; and (ii) if charging volumes are computed based on the average, as in Equation 12, then $\delta(CV_{ij}) = (\delta(P95(V_{p2p, ij})) + \delta(P95(V_{p2p, ji}))/2$.

Intuition suggests that the daily traffic periodicity is due to human habits. During the day, more users are connected to the Internet and traffic grows. There is thus a correlation between the time at which the P95 happen and the time at which most users are online. For P2P traffic, users runs P2P applications when they are online. It is thus likely that the P95 of total traffic happens closely to when the P95 of P2P traffic is reached. Table I compares the P95 on the inbound traffic observed on the different ISP traces described in Sec. II-C1 over a one-week long period of time. The second column shows the actual P95 of total traffic while the third column shows the total traffic observed at the time when the P95 of P2P traffic occurs. The fourth column reports the relative error and the fifth column reports the percentage of P2P traffic from the total traffic. As can be seen, the relative error ranges between 2% to 13% depending on the ISP link. Further, the larger the fraction of P2P traffic in the monitored ISP, the smaller the relative error. We repeated the analysis for outbound traffic. The relative error was even smaller since

Trace	Real [Mbps]	Approximation [Mbps]	Relative Error [%]	P2P Traffic [%]
ISP-1	1221.8	1247.5	2.10	48.6
ISP-2	1782.7	1660.9	6.83	45.1
ISP-3	1053.1	1029.6	2.23	53.02
ISP-4	1845.6	1765.3	4.35	42.5
ISP-5	1385.7	1347.1	2.79	50.4
ISP-6	1350.6	1173.6	13.1	6.5

TABLE I
APPROXIMATING THE 95TH PERCENTILE FOR INCOMING TRAFFIC

Geographic Location	\$ per Mbps
North America	10
Europe	14
Australia	34
Asia	38
South America	76

TABLE II
PRICING FUNCTIONS

outbound P2P traffic fraction was higher than 80% for all ISPs and thus the P2P traffic dominates the P95.

Second, as we have seen that P2P traffic volumes prior to localization tend to be correlated to total traffic volumes, we now argue that the trend will continue after localization. We have found in our datasets that the ratio of P2P traffic after localization to P2P traffic before localization does not vary much over time for all links, and locality models. Thus, it is reasonable to assume that the time the P95 occurs after localization does not shift. For instance, when the Single(history) locality model is used, for all links, the standard deviation of the ratio across various time snapshots (σ) is very small. In particular, 58% of links have $\sigma < .05$ and 90% have $\sigma < 0.1$.

Overall, the discussion above suggests that the errors introduced due to our approximations will be limited in practice, and our prediction of the impact of localization on ISP profitability will be reasonable.

C. Pricing Models

We now discuss the models we use to compute the pricing function, and charging volumes. While pricing functions are often non-decreasing piece-wise linear [12] they are specific to each provider and require the total volume of traffic to be known. To facilitate our evaluation we assume ASes use linear pricing functions where the charging volume is multiplied by the unit traffic volume price. Linear pricing functions are a good first step towards finding the actual costs and have also been used in determining transit costs for content providers [25]. Linear pricing is a valid approximation in our case because the reduction/increase of traffic that is experienced due to localization policies is not so large to trigger an economy of scale range change in the pricing. Moreover, assuming linear pricing corresponds to evaluating an upper bound on the possible saving an AS can achieve given the sublinear effect induced by economy of scale. The price per Mbps is known to vary widely due to geographic location [26] therefore we gather data from Telegeography Research [27] (summarized in Table II) to determine how customer ASes are charged.

We next discuss our models for charging volume, using Fig. 3 to aid our discussion. Assume that ASes *B*, *C*, *D*, and *E* are all residential ASes. The P95s of P2P traffic for all links before and after localization are reported in the figure.

Fig. 3. Example topology illustrating our pricing model. P95 refer to P2P traffic.

For instance, on the link between *A* and *B*, the P95 of the P2P traffic inbound to *B* is 600 Mbps and 150 Mbps, before and after localization respectively. Likewise, the P95 of the P2P traffic outbound from *B* is 450 Mbps and 100 Mbps, before and after localization respectively.

We now summarize the various pricing models we use:

- **Average:** For the charging volume we calculate the average of the inbound P95 ($P95_{IN}$) and outbound P95 ($P95_{OUT}$) for each link as in Equation 12. The change in charging volume on localization may be approximated as in Sec. IV-B. For instance, in Fig. 3, the charging volume on the link between *B* and *A* would decrease from $(600 + 450)/2$ Mbps to $(150 + 100)/2$ Mbps, a reduction of 400 Mbps. Considering traffic prices from Table II, AS *B* is charged \$10 per Mbps by *A*. Thus, localization will reduce *B*'s costs by \$4,000. However, the charging volume will also reduce on links from *B* to each of its customers, resulting in revenue reductions. The revenue reduction is $34 * (300 + 250)/2 - 34 * (130 + 100)/2 = \$5,440$ from customer *C*, $38 * (120 + 200)/2 - 38 * (60 + 20)/2 = \$4,560$ from customer *D*, and $10 * (400 + 375)/2 - 10 * (80 + 75)/2 = \$3,100$ for customer *E*. The $\delta(\text{profit})$ for *B* is then $-\$9,100$ and the profit increase is $\delta(\text{profit})/\text{profit}_{p2p, \text{before}} = -\$9,100/\$14,055 = -0.65$, indicating that 65% of profits on P2P traffic were lost.

- **Upper and Lower Bounds:** In contrast to the average case, computing changes in charging volume is more complicated if the pricing scheme is based on the maximum of $P95_{IN}$ and $P95_{OUT}$, as in Equation 11. Using such pricing schemes requires us to know whether the total traffic volume is higher in the inbound or outbound direction. However, we only have information regarding P2P traffic volumes. It is possible P2P traffic volumes are higher in the inbound (outbound) direction, while total traffic volumes are higher in the outbound (inbound) direction. To deal with this, we instead compute an upper and lower bound of the benefits that localization could have on each ISP.

Consider again the link between *A* and *B* in Fig. 3. Depending on whether *B* is charged based on inbound or outbound traffic prior to localization, and allowing for a change in the direction of charging volume after localization, the reduction in charging volume may range between $450 - 150 = 300$ Mbps, and $600 - 100 = 500$ Mbps. While precise determination of the change in traffic volume is difficult, the best possible scenario for *B* is a reduction of 500 Mbps, while the worst scenario is a reduction of 300 Mbps.

More generally, for a customer, the best possible case is obtained assuming $\max(P95_{IN}, P95_{OUT})$ before localization, and $\min(P95_{IN}, P95_{OUT})$ after localization. For a provider the opposite set of choices provides the best scenario. We also observe that on any link, the best scenario for the provider is the worst scenario for the customer, and vice versa. To compute the upper (lower) bound in terms of benefits for an AS when localization policies are applied, we assume the best (worst) case for each of its links.

We now illustrate the lower and upper bound computation for B . In the worst case scenario, the decrease in costs on provider links on localization is $10 * (450 - 150) = \$3,000$, while the decrease in revenue from customers is $34 * (300 - 100) + 38 * (200 - 20) + 10 * (400 - 75) = \$16,890$. Thus, the lower bound on $\delta(\text{profit})$ is $-\$13,890$ and profit decrease is 80%. However, in the best case scenario for B , the decrease in costs on provider links on localization is $10 * (600 - 100) = \$5,000$, while the decrease in revenue from customers is $34 * (250 - 130) + 38 * (120 - 60) + 10 * (375 - 80) = \$9,310$. Thus, the upper bound on $\delta(\text{profit})$ is $-\$4,310$ and the profit decrease is 37%.

- **Class:** Since knowing if an AS link is inbound or outbound dominated for all links is practically impossible, we consider a scenario that we build to be as realistic as possible. We start by considering PeeringDB [28] information, which is a database where network operators document information in hope of attracting other ASes to peer with. The database contains over 1,900 ASes that provide the ground truth by labeling themselves as having traffic ratios that are dominated by inbound, outbound, or are balanced. About 500 ASes are in our dataset and for them we explicitly consider this information.

For the remaining ASes, we hypothesize that the ratio of P2P client to web server populations has a large impact on the amount of traffic entering and leaving an AS. This is because we would expect a residential AS with many P2P clients to have a large number of users consuming content; hence a large amount of inbound traffic. On the other hand, we would expect an AS hosting many web servers to have large outbound traffic. To discover the server population per AS we use a methodology similar to that used in [10] and find 1 million servers in 19,000 ASes. We use then PeeringDB as ground truth to calibrate the threshold ratio to classify ASes. Indeed, we do find a strong correlation between dominating traffic direction and the ratio of population sizes. Considering the unclassified ASes, we find 95% have population ratios clearly indicating they are inbound dominated (and we label as such in this scenario); this is unsurprising as we would expect most ASes in our dataset to be residential ASes and not content providers. Given each AS classification, the corresponding P95 of incoming or outgoing traffic will be used as the charging volume for every provider link the AS has.

To complete the example, let ASes B , C , D and E be classified inbound dominated. Then, the decrease in costs for AS B is $10 * (600 - 150) = \$4,500$, while the decrease in revenue is $34 * (300 - 130) + 38 * (120 - 60) + 10 * (400 - 80) = \$11,260$. This translates to a $\delta(\text{profit})$ of $-\$6,760$ and a profit decrease of 68%.

AS Type	# Profiting (%)	# Losing (%)
All ASes	322 (16%)	1745 (84%)
Stub	60 (5%)	1140 (95%)
Small ISP	115 (20%)	458 (80%)
Large ISP	139 (49%)	147 (51%)
Tier-1	8 (100%)	0 (0%)

TABLE III
ASES PROFITING OR LOSING BY CATEGORY (NO LOCALIZATION)

V. IMPACT OF LOCALIZATION POLICIES

In this section, we evaluate the profitability of ISPs according to the P2P traffic that they carry today and how localization affects this. We consider scenarios where a different fraction of ASes localize traffic as ISPs may implement locality policies independent of one another. We also perform sensitivity to locality models and pricing models.

Using the Affinity model and the BitTorrent crawl, we consider, for each locality model, a set of 168 traffic matrices derived from the last 7 days of the 8 day long crawl. The first day is not considered to discard initial transient conditions for the history and persistent locality models. For each matrix, traffic is then routed on the AS level topology using the AS paths inferred as described in Section II-C1. Finally, for each customer-provider link, the P95 of P2P traffic is computed considering the 168 samples.

So far we have classified ISPs based on their customer-provider relationship as transit or residential to allow us to clarify the implication of the pricing model. However, this classification does not capture the implication of the AS size on the ISP's profitability, therefore we also categorize each AS according to how many downstream customers it has as proposed by the Internet Topology Collection [29]. There are four categories: Stub, Small ISP, Large ISP and Tier-1, which intuitively state how big a transit AS is. Stubs have less than 5 downstream customers, Small ISPs have 5 or greater, but less than 50, and Large ISPs have 50 or greater. Tier-1 ISPs are those who have no or very few providers and are the same as those identified by [29]. In our dataset there are 1200 Stubs, 573 Small ISPs, 286 Large ISPs, and 8 Tier-1 ISPs.

A. Profitability Before Localization

We first consider the scenario where there is no localization used on the Internet. We determine for each category of AS, the number of ASes that are profiting or losing from carrying P2P traffic. We present results only for the Class pricing model since results for Average are similar and Upper and Lower can be calculated only when localization occurs. Table III reports the results. As expected, the vast majority of ASes lose money because of P2P traffic (see the first line summary). However, as the number of downstream AS customers increases, there are ASes that profit due to P2P traffic. Indeed, 322 ASes (16%) today are profitable overall, of which 266 ASes are residential. This indicates that not all ASes may want to limit P2P traffic, and only ASes that have few customers have the most incentive to limit external P2P traffic.

Considering ASes that have losses due to P2P traffic, over 51% of of them are purely residential serving end users but not carrying traffic for other ASes. Surprisingly, several ASes that have more than 500 downstream customers still suffer

Fig. 4. Individual AS deploys localization with Single(history) locality model. Sensitivity to pricing models.

losses. Investigating further, we found that their relationship to Tier-1 ASes largely determines whether they profit or lose. Being a customer of a Tier-1 AS makes the AS lose money, while those that have peering agreements made a profit.

A closer look reveals that some ASes are still profitable, in spite of having few provider agreements with Tier-1. For example, the 13th largest profitable AS (AS-12956, Telefonica) has few agreements with Tier-1 ASes but more than 500 downstream customers, most of which are in Spanish speaking regions. By carrying mostly traffic that is exchanged among South American and other Spanish ASes, it takes advantage of the cultural and linguistic characteristics of P2P inter-AS traffic to send high volumes of profitable traffic between customer ASes and very little costly traffic to Tier-1 providers.

Insight #1: *Transit ASes that have customer ASes with similar cultural and linguistic makeups benefit more from carrying P2P traffic than those whose customer ASes are dissimilar. A transit AS with such customer ASes sends more traffic to customers than to providers which increases its revenue.*

B. Localization Deployed by Individual ASes

We seek to understand if localization is beneficial for an individual AS, independent of what other ASes do. Specifically, we investigate what is the expected benefit for an AS that deploys a localization policy alone. We consider only residential ASes since pure transit ASes have no benefits in localizing traffic (having no clients).

Sensitivity to pricing model: We fix the locality policy to Single(history) and calculate the charging volumes as described in Section IV-C. We use the metrics defined by Equation 14 and 15. Fig. 4(a) summarizes what the values of these metrics mean for different ranges. Positive values indicate that the AS is benefiting from the locality policy. Negative values indicate that the AS is doing worse than before the policy is applied. For example, a profit increase larger than 0 means more profit, while a profit increase between 0 and -1 means less profit. Profit goes to 0 when profit increase takes the values of -1. Finally, for values smaller than -1 the localization policy turns profit into loss. We show results in two different graphs: Fig. 4(b) plots the Cumulative Distribution Function (CDF) of loss reduction for ASes who have losses before localization, and Fig. 4(c) plots the CDF of profit increase for ASes who profit before localization.

In Fig. 4(b), the Lower bound (i.e., the vertical curve at $x=0$) shows that no profit is gained. This is because the localization

of P2P traffic will result in internal P2P clients reducing content downloaded from the outside, but in the pessimistic case this will not necessarily reduce content uploaded to other ISPs. However, for the worst case, the AS is charged on outbound traffic which has not changed. As Class reveals though, most residential ASes do get charged for their incoming traffic and thus localization is beneficial to them. Benefits are somewhat limited, with a loss reduction smaller than 30% for more than 50% of ASes. Note that a few residential ASes are outgoing dominated and thus are unaffected by localization also for Class. In our dataset we find 40 ASes that belong to this category. For those, loss reduction is equal to 0, as shown by vertical segment of the Class curve close to $x=0$.

For Average, less benefit is obtained than for Class because Average considers both directions of traffic but the cost associated with outbound traffic remains the same. Finally, Upper bound provides optimistic predictions that are unlikely in practice. Surprisingly even in this case the loss reduction is limited, i.e., only 40% of ASes see reductions over 60%.

We now turn our attention to profitable ASes in Fig. 4(c), a total of 16 residential ASes for which most P2P traffic traverses customer links. We see that in Class, 63% of these ASes show a profit reduction. This is due to these residential ASes also being transit ASes. For example, the AS that suffers the most is AS-209 Qwest Communications, a Tier-1 provider who we found to have over 360,000 clients. This is due to almost all of the P2P traffic that clients in AS-209 generate being sent and received through customer links.

Insight #2: *Some residential ASes will actually lose profit when they localize traffic. This is due to these ASes also being transit providers for other residential ASes. For these ASes, P2P traffic that was previously downloaded from clients in customer ASes decreases due to localization and in turn revenue decreases. Therefore, they have little incentive to localize traffic.*

There are a few ASes that are able to increase profit due to localization. This is due to the fact that many AS paths are asymmetric. Specifically, outgoing traffic is sent on customer links and since outgoing traffic does not decrease when one AS localizes, revenue remains the same. However, some incoming traffic is received on provider links, hence a reduction in costs and an increase in profit. This underlines the complexity of possible impacts of P2P traffic localization policies.

Sensitivity to locality model: Now we fix the pricing model to Class and vary the locality model. As expected, Fig. 5(a) shows that most ASes that were suffering losses due

Fig. 5. Individual AS deploys localization with Class pricing model. Sensitivity to locality models.

Fig. 6. All ASes deploy localization with Class pricing model. Sensitivity to locality models.

to the P2P traffic are reducing their loss due to localization policies. However, the reduction is not as large as one could hope. Under *Perf(no history)*, the most pessimistic locality policy, for 75% of ASes the reduction is less than 25%. Even under *Single(history)*, the most realistic locality policy, the loss reduction is still small, with less than 48% reduction for 75% of ASes. This is due to the small number of clients interested in the same content, which therefore tends to “disappear” as clients leave the torrent. Indeed, under the *Single(persistent)* policy the results are much improved: even 50% of ASes reduce their losses by 70%. In some cases, the AS is able to vastly improve profitability. For example, some small residential AS would be able to increase its loss reduction from 13% under *Single(history)* to 82% under *Single(persistent)*. This is due to the optimistic assumption that content is available forever once it enters an AS.

Insight #3: *Content availability plays a crucial role in determining the effectiveness of localization. Due to churn, peers will often need to redownload content from outside the AS. However, when assuming persistent content, most ASes can reduce losses twice as much.*

C. Internet-wide Localization Deployed

We now consider the scenario when all ASes deploy localization at the same time and thus we also include ASes who are purely transit in our results. We show results on sensitivity to locality models, but not on results concerning sensitivity to pricing models as the trends are similar to those already seen.

As before, we fix the pricing model to Class and plot results separately for ASes that normally lose or profit due to P2P traffic. Fig. 6(a) plots the loss reduction and shows results similar to when individual ASes localize. This is because an AS will reduce its incoming traffic only if it localizes its own traffic. Thus, as most ASes are inbound dominated they can unilaterally localize traffic and receive the full benefits.

However, an AS will reduce its outgoing traffic only if other ASes localize their traffic. Therefore, ASes that are outbound dominated will not see benefit until other ASes start localizing traffic. In this scenario indeed, all ASes that were facing loss reduce costs (loss reduction greater than 0 for all ASes).

Insight #4: *The benefits of localization will be limited for some ASes unless all ASes start to localize traffic. Localization, if adopted by a single AS, only reduces traffic received by internal peers, but it may not affect traffic sent. Hence, individual ASes that are outbound traffic dominated or are charged based on the average of the inbound and outbound P95s will not receive all the possible benefits. This reduced benefit may slow down the adoption of localization policies.*

To investigate which ASes benefit more we show in Fig. 6(b) the loss reduction versus population size, considering the *Single(history)* locality policy. As can be seen, there is a trend that the larger the population, the more the AS can localize. For example, the Taiwanese AS-3462 where we found over 1.5 million clients, is able to get a reduction of 91%. However, more than 50% of ASes achieve gains smaller than 30% as the limited number of peers interested in the same content inside an ISP limits the benefits of localization.

Insight #5: *The reduction in traffic due to localization only grows logarithmically with client population (notice the log-linear scale). Furthermore, we find that for all locality models the values of α_j , the leftover traffic, also follow a similar logarithmic trend with respect to AS population sizes. This is due to torrent popularity following a Zipf-distribution, which has been shown to limit the effectiveness of caching. In particular, [30] demonstrates through analysis that a similar effect occurs considering web caching.*

Moving to ASes that were already profitable, Fig. 6(c) shows a significant decrease in the amount of profit; in a pessimistic case – *Single(no history)* policy – 50% of ASes lose over 25% of their profits. In an optimistic case – *Single(persistent)* policy – 50% of ASes increase their profits by over 25%.

Fig. 7. Largest ASes deploy localization with Single(history) locality model. Sensitivity to pricing models.

gle(persistent) policy – 80% loss at least 60% in profit. Thus, while localization is beneficial for many residential ASes, over 300 transit ASes lose profit. Further investigation shows that the larger the transit AS is, the more likely it will suffer heavier losses in profit.

Insight #6: *Transit ASes lose significant amounts of profit when ASes localize. We found that all Tier-1 ISPs will lose over 56% of profits on P2P traffic under Single(history) when all ASes localize. We note though that the reduction in P2P traffic may enable transit ISPs to use the bandwidth for other types of revenue generating traffic, thus offsetting losses.*

Some ASes turn from being profitable to actually losing money (profit increase smaller than -1). For example, this happens for the AS-3786, who is a transit provider for the AS-17858. As AS-17858 is larger than AS-3786, it can reduce its traffic more than AS-3786 can. Therefore, AS-3786’s customer traffic is reduced more than its provider traffic and hence it starts to lose money. Interestingly, there are a few ASes that are able to increase their profits due to localization. These transit ASes are providers for many small residential ASes. As small ASes achieve very small reductions, the transit ASes are able to increase their profits by reducing their costs more than their customers can.

Insight #7: *Small residential ASes have small reductions in traffic due to the logarithmic trend of localization. Hence transit ASes who carry traffic for many small ASes fare better than those who carry traffic for a few large ASes.*

D. Localization Deployed by Large ASes

Besides the extreme cases when a single AS or all ASes deploy localization, we also investigate the scenario when ASes with larger populations will implement localization. We consider the Single(history) locality model and conduct sensitivity to pricing models. Results for sensitivity to locality models are similar.

We first consider when only the 100 largest ASes by client population size localize traffic, i.e., 6% of residential ASes in our dataset. As the largest ASes send and receive a very large amount of P2P traffic, we expect the localization to impact many other ASes as well. Fig. 7(a) shows the loss reduction results. In Class, the 100 ASes that localize receive the full benefits while 87% of ASes do not practically benefit. This is because many ASes are inbound dominated, but only a decrease of outbound traffic could be induced in this scenario. Indeed, the 40 ASes that are outbound dominated benefit of a loss reduction of 30% or greater. The Lower curve

corroborates this result by showing that most of ASes cannot get any benefit. The Average pricing model allows more ASes to benefit from the localization deployment of few ASes, while the Upper Bound provides over-optimistic prediction.

Insight #8: *Pricing scheme has a large impact on the effectiveness of savings. As the maximum pricing model ignores one direction of traffic, reduction in the other direction does not result in a reduction of cost. The average pricing model does consider both inbound and outbound traffic and thus an AS could benefit both if it or some other AS localizes traffic.*

We now compare the two most realistic pricing models, Class and Average, when a varying number of ASes localize traffic. We only focus on loss reduction graphs as we wish to highlight the effects the maximum and average pricing models have on costs. Fig. 7(b) shows Class and demonstrates that those who localize are generally the only ones who see benefit. This is in contrast to the Average pricing model, which we show in Fig. 7(c). Interestingly, in Average, almost all ASes that do not localize see increasing benefits as more ASes localize. For example, when 200 ASes localize, most ASes have over a 20% loss reduction, which is over 50% of the benefits possible when all ASes localize.

Insight #9: *Contrary to the average pricing model, for the maximum pricing model it is not sufficient that few ASes localize traffic to reduce cost. Even if the largest ASes start deploying localization schemes, overall loss reduction will be very limited.*

VI. IMPACT OF BUSINESS-RELATIONSHIP POLICIES

In this section we explore alternative ways to increase profitability of carrying P2P traffic. In particular, we explore business-relationship based peer selection policies where ASes aim to improve their profit by making internal peers select external peers located in customer or peer ASes, while trying to avoid peers hosted in provider ASes. AS is not trying to reduce the traffic peers will download, but rather it is interested in carefully selecting the source ASes to download from.

Clearly, the generalized use of these policies could have significant impact on existing peering agreements. As traffic exchange ratios [31] are often used to determine if an AS should be a peer or customer, a change in traffic may lead to a renegotiation of agreements. In this paper, we do not consider such events.

Fig. 8. Individual AS deploys Business, Single(history) or Hybrid, with Class pricing model.

Fig. 9. All ASes deploy Business, Single(history), Hybrid or Individual Best, with Class pricing model.

A. Modeling Business-Relationship Based Policies

To model this preferential peer selection, we define θ_{ij} as a preference bias index given by AS j to remote AS i . If the path from i to j traverses a customer link of j , the preference will be the highest ($\theta_{ij} = 1$); if the path from i to j traverses a peering link of j , a middle preference will be assigned ($\theta_{ij} = w_p, 0 < w_p \leq 1$); finally, if the path from i to j traverses a provider link of j , the preference will be the lowest ($\theta_{ij} = w_q, 0 < w_q \leq w_p$). Then, the volume of P2P traffic sent from AS i to AS j is:

$$X'_{ij} = X_{ij}\theta_{ij}B(j), \quad (16)$$

where X_{ij} is computed based on the Affinity model as in Equation 3, and $B(j)$ is a normalization factor that ensures the aggregate traffic downloaded by peers in AS j from external peers remains the same before and after the policy is applied.

$$B(j) = \frac{\sum_{i=1}^{D_j} X_{ij}}{\sum_{i=1}^{D_j} (X_{ij}\theta_{ij})}, \quad (17)$$

where D_j is the total number of ASes from which j downloads content. We refer to this model as the *Business* model.

We have performed a sensitivity study to w_p and w_q , to understand how these parameters affect the loss reduction and the profit increase of ASes. Intuitively, ISPs should make w_p and w_q as small as possible to obtain the most benefits out of Business. In the extreme, if we make $w_q = 0$, all the traffic from an AS will be directed to customer or peering links. However, in practice this may not be possible since customer or peer ASes of an ISP may not have the content or may not have enough clients to support the demand. Hence, we pick a very small value of w_q , in particular we use $w_q = 1\text{E-}10$. For w_p , the main requirement is that it is larger than w_q ; we choose $w_p = 1\text{E-}03$. We note that Business is an extreme version of such a scheme that we use to illustrate its potential.

In reality, other practical considerations should be made, such as considering user performance and inter-AS link capacities.

Intuition suggests we can improve the performance of the Business and Single policies by merging them. We call the new policy *Hybrid* and we model it by substituting L_{ij} from Equation 4 into both Equation 16 and 17, i.e., $X'_{ij} = L_{ij}\theta_{ij}B(j)$. This represents the policy for selecting peers outside an AS to obtain content that is not already present inside the AS. We use $w_p = 1\text{E-}03$ and $w_q = 1\text{E-}10$ as before.

B. Best Strategy for Individual ASes

The goal of this section is to study what strategy individual ASes should adopt to have the best impact on ISP profitability. We start by considering the case in which individual ASes deploy one of Business, Hybrid or Single(history). We fix the pricing model to Class. Fig. 8(a) shows a comparison between the three possible strategies reporting loss reduction. Interestingly, the Business policy is ineffective for more than 75% of ASes, while Single(history) has proved to reduce the loss for most ASes. The Hybrid policy is providing the best loss reduction for most of the ASes. Indeed, only for the top 25% of ASes, which are mostly transit ASes, Business performs better than Single(history) and similar to Hybrid. This is because transit ASes can benefit more from the Business policy by having internal peers download traffic from customer ASes rather than provider ASes. In fact, the top 11% of ASes actually turn profitable, i.e., loss reduction becomes greater than 1.

Fig. 8(b) shows the profit increase for the 16 residential ASes that are already profitable before localization. The figure shows that Business is the most beneficial policy, i.e., more than 30% of the ASes improve their profit by more than 100%. The other two policies can instead cause a profit reduction, as already seen in Fig. 5. Recall indeed that transit ASes will increase their profit if more traffic is pushed to customer ASes.

Based on these results, we aim to study the strategy that gives the most benefits to ASes. Towards this goal, we plot Fig. 8(c). In this figure, we consider all ASes and group them according to the categories described in Sec. V. Then, we find for each AS, which policy gives the most benefits. Finally, we aggregate the best policies per category of AS. For each type of AS there is a stacked bar, where each section of the bar represents a fraction of ASes that performs the best with a given policy. Note that besides the Business and Hybrid policies, there is $Single(history) = Hybrid$, which accounts for the cases in which both Single(history) and Hybrid are the best policies. Single(history) is never better than Business or Hybrid, so it is not shown in the picture.

There are several points to take away from Fig. 8(c). First, we observe that for around 90% of stub ASes, the best policy is Single(history) or Hybrid. This is because stub ASes receive considerable benefits from localization. ASes in the Small ISP category follow a similar trend with more than 60% of them benefiting the most from Single(history) or Hybrid. Second, all Tier-1 ASes on the contrary will get the most benefits out of the Business policy. This is because Tier-1 ASes will benefit from an increment in the traffic sent or received from customers. Finally, Hybrid is better for the Large ISP category, since these ISPs benefit both from directing traffic to customers and from localizing their own P2P traffic.

Insight #10: *Many ASes will achieve more profits through preferentially directing traffic to customers and peers rather than localizing traffic. Therefore, P2P traffic localization is not always the best choice for all ASes.*

C. Internet Impact of Business-Relationship Based Policies

In the previous section, we have seen how different strategies will benefit ASes if individual ASes adopt them. But what happens when all ASes adopt the same policy at the same time or when all ASes adopt their local best policy at the same time? Both P4P and ALTO indeed allow each AS to run their own “oracle” and chose a different policy. To answer these questions, we consider the scenarios in which all ASes adopt Business, Single(history), and Hybrid policies. In addition, we consider the scenario in which each AS applies its best local strategy, according to Fig. 8(c), which we have called “Individual Best”. We fix the pricing model to Class.

Fig. 9(a) shows the loss reduction. For about 10% of the ASes, Business causes them to lose considerably more. These are mostly stub ASes that will be “victims” of their providers that increase the amount of traffic they exchange with customer ASes. On the contrary, Single(history) almost never causes higher loss. Hybrid performs marginally better than Single(history) for large ASes, but slightly worse than Single(history) for small ASes.

Fig. 9(b) shows the profit increase. Business performs better than Single(history) and Hybrid. When Business alone is considered, more than 70% of ASes either profit more or earn the same amount as before. For Single(history), over 90% of the ASes start losing profit due to localization. This is because many of the transit ASes that were profiting before will receive more benefit from Business since they will now select peers in customer ASes and direct more traffic to them.

We note that for both loss reduction and profit increase, Individual Best closely follows Single(history) and Hybrid. In particular, ASes that profit from P2P traffic (e.g. Tier-1 ASes and a few Large ISPs), which benefit more from locally implementing Business, lose because of policies implemented by their customers.

Insight #11: *While business-relationship based policies may locally be the best strategy for some ASes, they can have a negative external impact on other ASes. Furthermore, as the best local strategy of an individual AS is chosen in isolation of others it does not turn to be the best possible choice when all ASes deploy their own best local strategy.*

VII. RELATED WORK

Much work on modeling traffic on the Internet has been done in the context of intra-AS traffic matrix estimation [13, 14]. Our work though focuses on inter-AS P2P traffic matrix estimation, of which the only related paper is [10], which we extensively discussed in Sec. II.

The effects of P2P systems on ASes has also been studied by Rasti et al. [32] who shows the effects of the Gnutella P2P system on the AS topology. Rather than focusing on localization, they instead study how the load on ASes due to Gnutella clients has changed over time due to the evolution of both the AS topology and the Gnutella system.

Many recent works have focused on how to implement P2P localization [2]–[8, 23]. However, we evaluate the impact localization will have on all ASes and on their profitability.

Complementary to our work is that done by Cuevas et al. [11]. Their focus is on understanding the extent to which localization helps improve the performance of users and reduce the amount of P2P traffic residential ASes exchange with their providers. Similarly, Blond et al. [33] focus on how much traffic can be reduced due to localization using experiments driven by a BitTorrent crawl. In contrast, our goal is to understand the implications of localization on the global Internet, particularly, which ASes will benefit and which will lose. In addition, our analysis not only considers residential ASes but also study how localization may affect pure-transit ASes, which may not have any internal peers.

Piatek et al. [9] question the effectiveness of localization on peers performance and ISP traffic reduction. Specifically, they perform experiments showing that client-only localization policies will have limited benefits and the tracker will need to be involved to receive full benefits. They also evaluate the amount of traffic reduction possible for a crawl of one thousand torrents. In contrast, we consider a very large dataset including millions of torrents and also use realistic pricing models to understand how traffic reductions translate into impact on profit for ISPs.

VIII. CONCLUSION

In this paper, we developed a detailed methodology for evaluating the profitability of an ISP and how it will change due to P2P localization. We first proposed the Affinity model, a refinement of the Gravity model, for generating realistic inter-AS P2P traffic. We then devised several locality models to

describe the reduction of P2P traffic under different scenarios. Coupling these models with realistic inter-AS paths inferred from BGP and IXP data, and pricing models based on the 95th percentile and geographic pricing, we calculate the impact of localization policies on ISP profits. We believe that the results we presented enhance the understanding and implications of P2P traffic localization schemes on the Internet, and in particular from the perspective of ISPs.

REFERENCES

- [1] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet inter-domain traffic," in *SIGCOMM*, 2010.
- [2] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. G. Liu, and A. Silberschatz, "P4P: Provider Portal for Applications," *SIGCOMM*, 2008.
- [3] V. Aggarwal, A. Feldmann, and C. Scheideler, "Can ISPs and P2P Users Cooperate for Improved Performance?" *SIGCOMM Computer Communication Review*, vol. 37, no. 3, pp. 29–40, 2007.
- [4] D. R. Choffnes and F. E. Bustamante, "Taming the Torrent: a Practical Approach to Reducing Cross-ISP Traffic in Peer-to-Peer Systems," in *SIGCOMM*, 2008.
- [5] F. Picconi and L. Massoulié, "ISP-Friend or Foe? Making P2P Live Streaming ISP-aware," in *ICDCS*, 2009.
- [6] J. Wang, C. Huang, and J. Li, "On ISP-Friendly Rate Allocation for Peer-Assisted VoD," in *ACM International Conference on Multimedia*, 2008.
- [7] M. Lin, J. C. Lui, and D. Chiu, "Design and Analysis of ISP-Friendly File Distribution Protocol," in *Allerton Conference on Communication, Control, and Computing*, 2008.
- [8] R. Bindal, P. Cao, W. Chan, J. Medved, G. Suwala, T. Bates, and A. Zhang, "Improving Traffic Locality in BitTorrent via Biased Neighbor Selection," in *ICDCS*, 2006.
- [9] M. Piatek, H. V. Madhyastha, J. P. John, A. Krishnamurthy, and T. Anderson, "Pitfalls for ISP-Friendly P2P Design," in *Hotnets*, 2009.
- [10] H. Chang, S. Jamin, Z. M. Mao, and W. Willinger, "An Empirical Approach to Modeling Inter-AS Traffic Matrices," in *IMC*, 2005.
- [11] R. Cuevas, N. Laoutaris, X. Yang, G. Siganos, and P. Rodriguez, "Deep Diving into BitTorrent Locality," in *INFOCOM*, 2011.
- [12] D. K. Goldenberg, L. Qiu, H. Xie, Y. R. Yang, and Y. Zhang, "Optimizing cost and performance for multihoming," in *SIGCOMM*, 2004.
- [13] Y. Zhang, M. Roughan, N. Duffield, and A. Greenberg, "Fast Accurate Computation of Large-Scale IP Traffic Matrices from Link Loads," in *SIGMETRICS*, 2003.
- [14] V. Erramill, M. Crovella, and N. Taft, "An Independent-Connection Model for Traffic Matrices," in *IMC*, 2006.
- [15] G. Salton and M. J. McGill, *Introduction to Modern Information Retrieval*. New York, NY, USA: McGraw-Hill, Inc., 1986.
- [16] OpenBittorrent, <http://www.openbittorrent.com>.
- [17] Team-Cymru, "IP to ASN mapping," <http://www.team-cymru.org/Services/ip-to-asn.html>.
- [18] X. Dimitropoulos, Y. Hyun, D. Krioukov, M. Fomenkov, G. Riley, and B. Huffaker, "As relationships: Inference and validation," *ACM SIGCOMM Computer Communication Review*, Jan 2007.
- [19] B. Augustin, B. Krishnamurthy, and W. Willinger, "Ixp: mapped?" in *IMC*, 2009.
- [20] J. Qiu and L. Gao, "AS Path Inference by Exploiting Known AS Paths," in *GLOBECOM*, 2004.
- [21] O. RouteViews, <http://www.routeviews.org/>.
- [22] M. Mellia, R. Lo Cigno, and F. Neri, "Measuring IP and TCP Behavior on Edge Nodes with Tstat," *Computer Networks*, vol. 47, no. 1, pp. 1–21, 2005.
- [23] J. Seedorf, S. Kiesel, and M. Stiernerling, "Traffic localization for p2p-applications: The alto approach," in *IEEE P2P*, September 2009.
- [24] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. E. Anderson, A. Krishnamurthy, and A. Venkataramani, "iPlane: An Information Plane for Distributed Services," in *OSDI*, 2006.
- [25] Z. Zhang, M. Zhang, A. Greenberg, Y. C. Hu, R. Mahajan, and B. Christen, "Optimizing cost and performance in online service provider networks," in *NSDI*, 2010.
- [26] N. Laoutaris, G. Smaragdakis, P. Rodriguez, and R. Sundaram, "Delay tolerant bulk data transfers on the internet," in *SIGMETRICS*, 2009.
- [27] T. Research, "Telegeography international telecom trends seminar," http://www.ptc.org/ptc09/images/papers/PTC'09_TeleGeography_Slides.pdf.
- [28] PeeringDB, "Peering networks," <https://www.peeringdb.com/>.
- [29] Internet Topology Collection, <http://irtl.cs.ucla.edu/topology/>.
- [30] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web caching and zipf-like distributions: Evidence and implications," in *INFOCOM*, 1999, pp. 126–134.
- [31] D. Peering, http://drpeering.net/AskDrPeering/blog/articles/The_Folly_of_Peering_Ratio.
- [32] A. H. Rasti, R. Rejaie, and W. Willinger, "Characterizing the global impact of p2p overlays on the as-level underlay," in *PAM*, 2010.
- [33] S. L. Blond, A. Legout, and W. Dabbous, "Pushing bittorrent locality to the limit," *Computer Networks*, vol. 55, no. 3, pp. 541 – 557, 2011.

Jeff Seibert is a PhD student in the department of Computer Science at Purdue University. He received his BS in Computer Science from Purdue in 2006. He is a member of CERIAS and the Dependable and Secure Distributed Systems Laboratory. His research interests include security in distributed systems and networks.

Ruben Torres is a Ph.D. student in the School of Electrical and Computer Engineering at Purdue University and a member of the Internet Systems Lab. Ruben received his BS in electronics and communications engineering from the University of Panama in 2002 and his MS in electrical and computer engineering from Purdue University in 2006. His research interests include distributed systems and network security.

Marco Mellia is an Assistant Professor in the EE Department of Politecnico di Torino. He received his Ph.D. in Telecommunications Engineering in 2001 from Politecnico di Torino. In 1999, he was with the CS Department at Carnegie Mellon University. He co-authored over 140 papers published in international journals and conferences, and he was a member of the program committees of several conferences including IEEE Infocom and ACM Sigcomm. His research interests are traffic measurement, P2P applications and energy aware network design.

Maurizio M. Munafò is an Assistant Professor in the Electronics Department of Politecnico di Torino. He received the Dr. Ing. degree in electronic engineering and the Ph.D. degree in telecommunications engineering from Politecnico di Torino, in 1991 and 1994. He coauthored over 60 journal and conference papers in the area of communication networks and systems. His current research interests are in simulation and performance analysis of communication systems, and in traffic modeling, measurement and classification.

Cristina Nita-Rotaru is an Associate Professor in the department of Computer Science at Purdue University. She leads the Dependable and Secure Distributed Systems Laboratory. She received BS and MS degrees from Politecnica University of Bucharest, Romania, in 1995 and 1996, and a PhD degree in Computer Science from The Johns Hopkins University in 2003. She served on the technical program committee of over 40 conference in networking, distributed systems, and security. She received the NSF CAREER award. Her research interests include security and fault-tolerance for distributed systems, and networks. She is a member of the ACM and IEEE Computer Society.

Sanjay G. Rao is an Associate Professor in the School of Electrical and Computer Engineering at Purdue University, where he leads the Internet Systems Laboratory. His research interests are in enterprise network management, cloud computing, and peer-to-peer systems. He received a BS in Computer Science and Engineering from the IIT, Madras, and the Ph.D from the School of Computer Science, Carnegie Mellon University. He was a visiting researcher in the Network Measurement and Management group at AT&T Research in Summer 2006.

He is a recipient of the NSF Career award. He has served on the Technical Program Committees of several workshops and conferences including ACM Sigcomm, IEEE Infocom and ACM CoNext, and was the technical program co-chair of the INM/WREN workshop held in conjunction with NSDI 2010.