Peer-to-Peer Netw. Appl.
DOI 10.1007/s12083-012-0132-z

Characterization of community based-P2P systems and implications for traffic localization

Ruben Torres · Marco Mellia · Maurizio M. Munafo · Sanjay G. Rao

Received: 30 June 2011 / Accepted: 19 March 2012

Abstract In this paper, we present one of the first and most extensive characterizations of *closed community-based P2P systems*. Such systems are organic groups of peer-to-peer (P2P) clients, which can be joined only by users belonging to a certain network (e.g., connected to a given Internet Service Provider (ISP)). A number of factors motivate the growth of these communities, such as quality of content, anonymity of transfers, and the potential for better performance that enhances user experience. Our study is conducted in two contrasting environments—a campus network and a national ISP—located in different continents. In both cases, large-scale closed communities have been found to be the predominant P2P systems in use. We shed light both on the factors motivating the growth of such communities, and present results characterizing the extensiveness of their usage, the performance achievable by the systems, and the implications of such communities for network providers. While our findings are interesting in their own right, they also offer important lessons for ongoing research that seeks to localize traffic within ISP boundaries. In particular, our results suggest that (i) in ISPs with heterogeneous access technologies, the performance benefits to users on localizing P2P traffic is largely dependent on the degree of seed-like behavior of peers behind high-bandwidth access technologies; and (ii) while localization can reduce the traffic on Internet peering links, it has the potential to cause a significant increase in traffic on internal links of providers, potentially requiring upgrades of network links.

Keywords Peer-to-peer networks · Internet measurements

1 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). A vast majority of these systems are *Internet-scale,* and open to *any user* on the Internet. Indeed, the open nature of these systems is viewed as a key strength of P2P systems in enabling inexpensive and rapid deployment of services over the Internet.

In this paper, we raise the attention of the research community to the prevalence of *closed communities* of P2P users, and present an extensive characterization of such communities. Membership in such communities is restricted by imposing requirements on users that join the system. We focus on an important class of closed communities, where the primary criterion for admitting users is that they must be connected to the

---

R. Torres (✉) · S. G. Rao
Purdue University, 465 Northwestern Ave,
West Lafayette, IN 47907, USA
e-mail: rtorresg@purdue.edu

S. G. Rao
e-mail: sanjay@purdue.edu

M. Mellia · M. M. Munafo
Politecnico di Torino, Corso Duca degli Abruzzi,
24, 10129 Torino, Italy

M. Mellia
e-mail: mellia@polito.it

M. M. Munafo
e-mail: munafo@polito.it

same network (e.g., same ISP). While several research efforts have extensively characterized the performance and traffic characteristics of open and Internet-scale P2P systems (henceforth referred to as *generic P2P systems*), e.g., [12, 13, 22, 24], the study of closed and network-specific P2P systems (henceforth referred to as *P2P communities*) has received limited attention.

In this work, we characterize two communities, that we had the chance to monitor. The communities corresponded to two very contrasting networks. The first community has been created and used by customers in a large nation-wide ISP in Europe. The ISP offers customers Internet access, using both ADSL (1 Mbps Uplink and 20 Mbps downlink) and FTTH (10 Mbps uplink and downlink) technology. The community observed in this network is based on the standard eMule P2P application [11], which has been modified by users to avoid problems caused by the assignment of private IP addresses to hosts inside the ISP network. We refer to this community as *ISP-Community* in this paper. The second community has been found inside a large university campus with hosts having high speed Ethernet connectivity. In this network, users modified the standard DirectConnect (DC) P2P application [1], so that only peers that run on hosts inside the campus can actually join the community. We refer to this community as *Campus-Community* in this paper.

Our main contributions are as follow:

- We show that P2P communities are extremely popular (e.g., generating more than 60% of total traffic for the ISP) and large-scale (e.g., comprising hundreds of thousands of users in the ISP network). The usage of the communities far exceeds usage of other more generic P2P systems—for e.g., in the campus network over 90% of the peers download over 90% of all P2P data using *Campus-Community*.
- We compare the performance of users of the P2P communities with users of more generic P2P systems. Our results show the performance benefits are largely determined by the access technologies of the users, and the degree of seed-like behavior shown by users behind high-speed access technologies. For instance, users of *Campus-Community* enjoy several orders of magnitude better performance than users of generic P2P systems in the campus network thanks to the high bandwidth provided by the campus LAN. In contrast, in the ISP network, the throughput of *ISP-Community* connections with senders behind ADSL links shows no particular improvement compared to generic P2P systems. However, the users of *ISP-Community* do see an improvement which may be attributed to a small fraction of senders behind FTTH links.
- We develop techniques to enable network providers understand how the growth of P2P communities impacts network capacity planning, and how projected changes in access technologies of users may affect these results. Our techniques center around a model we develop for the inter-PoP traffic of P2P communities. In contrast to prior work on traffic matrix estimation (for e.g., [3, 29]) which is agnostic to individual applications, our focus is on developing an application-specific traffic matrix model. Through simulations conducted using the model, we show that (i) while *ISP-Community* does reduce traffic on peering points as expected, more surprisingly, it results in a substantial increase in the traffic carried on internal network links (e.g., more than 60% of backbone links carry more traffic when *ISP-Community* is present); and (ii) this trend is exacerbated as more users move to high-bandwidth access technologies.

P2P communities must be distinguished from recent research proposals that have proposed mechanisms to ensure traffic of P2P systems is localized to ISP boundaries [8, 27]. Unlike these works, closed P2P communities have grown organically among users, and are already extensively deployed. Localization of traffic is not an explicit goal that spurred the growth of these communities, yet may occur as a prominent side-effect of the communities being closed to clients belonging to particular networks. That said, our results have important implications for research on localization of P2P traffic within ISP boundaries and indicate that benefits of localization should not be taken for granted. We discuss this in greater detail in Section 8.

2 Peer-to-peer communities

In this paper, we present an analysis of P2P communities in two different networks, (i) a nation-wide ISP in Europe; and (ii) a large-scale campus network in North America. We present more information about the networks and the associated P2P communities in this section.

2.1 P2P community in an ISP network

We describe a P2P community found in a nation-wide ISP in Europe. The ISP offers customers Internet access, using both ADSL (up to 1 Mbps Uplink and 20 Mbps downlink capacity) and FTTH (10 Mbps uplink

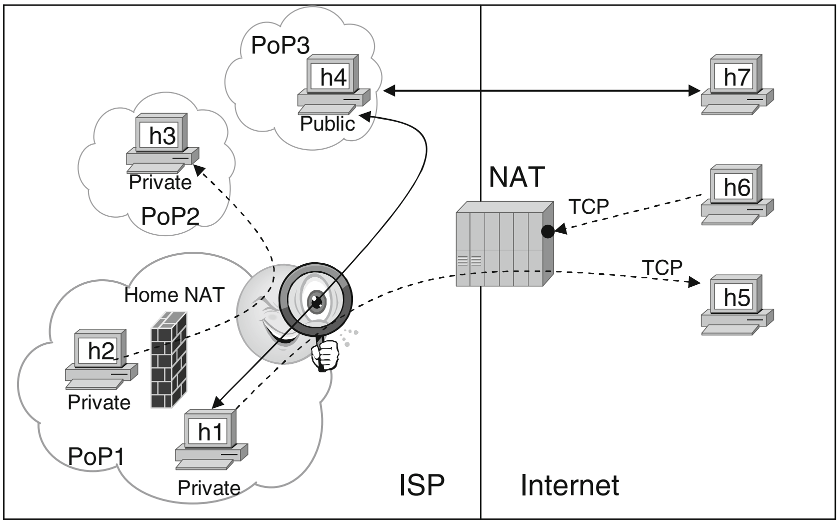

Fig. 1 ISP setup and trace collection

and downlink capacity) technology. Hosts in each city are aggregated into Points-of-Presence (PoPs), which are then connected via the ISP backbone. Typically, hosts in the ISP are given a private IP address. As shown in Fig. 1, plain connectivity is guaranteed to hosts inside the ISP network despite the use of private IP addresses. Whenever hosts with private addresses communicate with hosts in the external Internet, the data communication involves traversal of an ISP-wide NAT.

P2P systems typically have mechanisms in place to detect peers behind NAT and to limit their performance. This motivated a community of ISP users to modify eMule, a well known file sharing system, so that peers in the ISP could communicate with each other even though they have private addresses. The custom version of the eMule client was developed by ISP customers starting from 2003. The modification simply hardwires information about private IP addresses used within the ISP, and permits clients to send data to these addresses. We note that the default eMule system is associated with a DHT-based overlay known as *Kad*. The customized version of eMule (which we refer to as *ISP-Community*) builds a separate DHT overlay local to the peers in the ISP. This is achieved by modifying the message format of the original Internet-wide Kad overlay to ensure that the *ISP-Community* messages can only be processed by peers running the modified version.

Besides avoiding the NAT issue, *ISP-Community* offers other advantages. First, it is desirable to download content from other users connected to the same ISP since hosts within the ISP are interconnected through higher capacity backbone links. Second, a large percentage of hosts in the ISP are connected by FTTH, and their upload capacity is significantly higher than upload capacities of hosts connected to other ADSL providers. Third, given that all the peers in the community are in the same European country, the content that is available matches the interest and language preferences of users in the community. Finally, we note that *ISP-Community* clients could still use the global eMule system if content is not located within the local network. But this event is rare, as we will show in Section 4.2.

2.2 P2P community in a campus network

The second network we analyze is a large university campus in North America with tens of thousands of end hosts in its network, interconnected by a high capacity backbone LAN. Users in the campus network are offered Fast Ethernet connections (100 Mbps). In contrast to the ISP network, hosts in the campus receive public IP addresses, guaranteeing plain connectivity.

Motivated by the high bandwidth provided to clients in the campus network, students deployed a modified version of DC [1]. DC is a well known application for content sharing and chat, and we refer to the modified version as *Campus-Community*. In the traditional DC system, peers connect to a central *hub* for most of the system operations. However, in *Campus-Community*, there is no central hub, but a set of *hub clients* to avoid a single a point of failure. A *hub client* runs together with the DC application at each peer. When a peer is searching for content, a gnutella-like flooding algorithm is performed, in which peers forward the query to all their neighbors, until all peers receive the query. All peers that are sharing the content will reply back.

To enforce a closed membership, *Campus-Community* peers have been modified to only accept and initiate connections to other peers in the IP address range of the campus. As a side-effect, the *Campus-Community* traffic is therefore highly localized to the campus network. In addition, we found that *Campus-Community* is the most popular P2P application in the campus. We identified over a thousand *Campus-Community* peers which contribute to a large fraction of the campus traffic.

3 Evaluation goals and methodology

In this section we present our goals and methodology.

3.1 Goals

The aim of this paper is to characterize and compare the closed network-specific P2P community systems against generic, open and Internet-wide P2P systems. To accomplish this, in this paper we seek answers for the following questions:

- How extensive is the use of P2P communities?
- How does the performance seen by users with P2P communities compare to performance seen with generic P2P systems?
- How does the Internet access technology of clients impact the performance of users of the P2P community?
- What are the implications of the growth of P2P communities in terms of traffic on network links for network providers?

We answer these questions with a combination of measurements and simulations. The network measurements help us characterize application performance, user behavior, and provide realistic traffic information to guide our simulations. The simulations enable us to study the implications of the communities on network providers. Based on actual traffic data derived from our measurements, we devise a methodology to infer the volume of P2P traffic each link of the network has to carry. We defer further details on the simulation to Section 7 and focus on the methodology for the measurement study for the rest of the section.

3.2 Trace collection tool

Traces are collected with Tstat [17], a passive sniffer with advanced traffic classification capabilities. Starting from packet level traces collected in operational networks, Tstat groups packets into flows which are classified by application using a combination of Deep Packet Inspection and statistical classifiers, specifically targeting both plain and obfuscated P2P traffic. Tstat has been found to perform well in [20].

For each flow, Tstat collects various per-flow statistics such as bytes exchanged in each direction, flow duration and Round Trip Time (RTT) of packets. We refer the reader to [17, 18, 26] for more details.

3.3 Datasets

Our analysis is conducted on the following datasets:

ISP Network Traces have been collected from two PoPs in a nation-wide ISP in Europe. A high-end PC running Tstat was used to analyze in real time all the packets going to and coming from all the hosts in the monitored PoPs, and produced a flow level log that has then been post-processed. The two PoPs are different in the type of Internet access technology of their hosted customers. In the first PoP, which we call *ISP-FTTH*, all customers are connected through FTTH links while in the second PoP, which we call *ISP-ADSL*, all customers are connected through ADSL links. For the *ISP-FTTH* PoP analysis, we focus on a one week trace collected during December 2008, with about 2,200 active customers in the PoP contacting over 4 million hosts. For the *ISP-ADSL* PoP analysis, we focus on a one day trace collected during April 2009, with about 20,000 active customers in the PoP contacting over 2 million hosts.

For these datasets, we label clients according to their access technology to the network. This information has been encoded by the ISP on the IP address of clients and is easily obtainable from the traces. In addition, we associate clients with the ISP PoP were they reside. This information has been provided by the ISP operators.

Campus Network The trace has been collected at the edge of some campus dormitories, using a methodology similar to the ISP setting. We report results from a 13 h trace of a weekday in April 2009, during which there were about 2,000 distinct active hosts in the monitored dormitories. These hosts contact more than a million other hosts.

3.4 Comparing P2P communities with generic P2P systems

Our measurement studies compare the performance of closed P2P communities with generic Internet-scale P2P systems observed in the same network, and at the same time. We compare the *Campus-Community* and *ISP-Community* systems to two other regular, well

Table 1 Traffic summary for the one-week *ISP-FTTH* trace

Systems	Peers monitored	Peers contacted	Total data exchanged	Per peer data exchanged
ISP-Community	858	497.4 K	9141 GB	10.65 GB
ISP-Generic	470	317.7 K	683 GB	1.45 GB

known and open P2P applications: (i) the traditional eMule [11] application, which we refer to as the *ISP-Generic*; and (ii) the BitTorrent [7] application, which we refer to as *Campus-Generic*. Both *ISP-Generic* and *Campus-Generic* are the second most popular P2P file sharing systems after the P2P community applications in the respective traces. We note that our comparisons in the campus setting are based on two different underlying systems (the *Campus-Community* system is based on DC while the *Campus-Generic* is based on BitTorrent). While ideally, the comparisons are best performed using the same underlying system, this is infeasible since the campus network does not have a closed P2P community and a generic variant both based on the same underlying system. Thus, our comparisons in the campus network case could be impacted by other differences in the underlying systems involved, besides the open/closed nature of the systems. However, we believe the impact of these differences is relatively minor for most of the results, and the comparisons do provide important insights in the context of our study.

4 Characterizing community usage

In this section, we begin by characterizing the extent to which P2P communities are used in the monitored networks. Then, we consider the degree to which users rely on the P2P communities to access content.

4.1 Prevalence of communities

Tables 1, 2 and 3 summarize general statistics for the ISP and campus traces. For all tables, the first column shows for the various systems, the number of peers identified inside the monitored PoPs in the ISP and the dormitories in the campus. The second column shows the total number of external peers that are contacted by the monitored hosts. The third column gives the total data exchanged by peers, while the fourth column provides the average data exchanged per peer.

First, notice in Table 1 that, for the one week *ISP-FTTH* trace, the *ISP-Community* population in the monitored PoP is almost twice as large as the *ISP-Generic* population, with 858 *ISP-Community* peers versus 470 *ISP-Generic* peers. In addition, *ISP-Community* peers exchange over 13 times more data than *ISP-Generic* peers, and each peer exchanges over 7 times more traffic when using *ISP-Community*. Table 2 shows similar results for the one-day *ISP-ADSL* trace. However, notice that the difference in the total data exchanged between *ISP-Community* and *ISP-Generic* is not as large as in the *ISP-FTTH* trace. This is because ADSL peers cannot upload as much data as FTTH peers due to their limited upstream capacity. As we will see later, the higher upload capacity of FTTH peers plays a key role and justifies the success of *ISP-Community*. Finally, note that the total number of distinct *ISP-Community* peers found when combining the two traces, amounts to around 600,000 peers.

Table 3 shows similar results for the 13-hour Campus trace: the *Campus-Community* population in the PoP is 1.4 times larger than the *Campus-Generic* population. Besides, *Campus-Community* peers exchange 60 times more data than *Campus-Generic* peers in total, and each peer exchanges 16 times more data when using the *Campus-Community*.

Figure 2 shows the fraction of inbound and outbound traffic observed at the monitoring point which may be attributed to the P2P communities. There are three groups of bars, corresponding to the campus, and ISP scenarios. The fraction due to the communities is computed for each hour of the trace and the 50th and 90th percentile of these values is shown. The results show that throughout the duration of the trace, close to 90% of the outbound traffic may be attributed to the P2P communities for the Campus and *ISP-FTTH* settings. In the case of *ISP-ADSL*, the P2P community outbound traffic is reduced to around 60% of the total outbound traffic. This is because of the limited upload capacity of ADSL customers. On the other hand, the

Table 2 Traffic summary for the one-day *ISP-ADSL* trace

Systems	Peers monitored	Peers contacted	Total data exchanged	Per peer data exchanged
ISP-Community	7074	325.1 K	6669 GB	0.94 GB
ISP-Generic	2351	829.4 K	1021 GB	0.43 GB

Table 3 Traffic summary for the 13-hour campus trace

Systems	Peers monitored	Peers contacted	Total data exchanged	Per peer data exchanged
Campus-Community	270	1.82 K	970 GB	4.61 GB
Campus-Generic	196	1.006 M	14.65 GB	0.28GB

fraction of inbound traffic, while slightly lower due to the higher fraction of HTTP traffic, is still over 60% in the three settings. The results clearly illustrate the overwhelming popularity of the P2P communities.

4.2 User reliance on communities

We were interested in measuring the extent to which users rely on the community to obtain the content they require. To evaluate this, we measure the community usage ratio $U(p, \tau)$, of peer p during time interval τ by considering the ratio of bytes p downloads from other peers using the community to the total bytes downloaded using any P2P system (including the community). We selected τ to be 1 hour long. Intuitively, this metric quantifies the extent to which a peer has to back up to an Internet-wide generic P2P system to retrieve content which cannot be found in the P2P community.

Figure 3 shows the mean $U(p, \tau)$ per client. The plot reports two bars for each of *Campus-Community* and *ISP-Community*. Each bar shows the fraction of clients with mean $U(p, \tau)$ greater than 80% (and 90%) for the community they are part of. In general, $U(p, \tau)$ is high. In particular, for *ISP-Community*, 70% of the peers download more than 90% of the data from the community. For *Campus-Community*, the fraction of clients is even larger, e.g., more than 90% of the peers download more than 90% of data from the P2P community. This suggests that the P2P community is self-sustaining and peers usually locate content within the community.

5 User and system performance

In this section, we study the performance seen by users with the P2P communities, and compare this to the performance of the generic P2P systems. The most direct metric to evaluate user performance is the file transfer completion time. However our dataset does not allow us to gather this information since our probes do not perform parsing and interpretation of application header. Instead, we evaluate the download and upload rates and the delay per connection and per host, which are metrics clearly related to the overall user performance.

5.1 Throughput

In this section, we investigate to what extent users achieve better throughput if they use the P2P community. We then consider various factors that could affect the performance of users, such as their access technology.

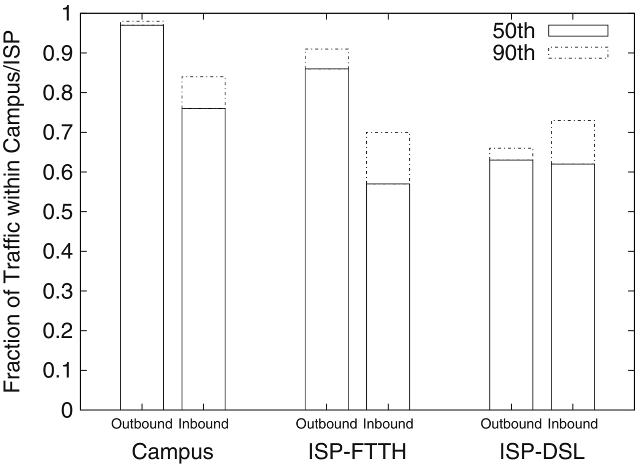

Fig. 2 Fraction of inbound and outbound traffic seen at the monitoring point, generated by the P2P communities. Note that all this traffic stays within the Campus/ISP

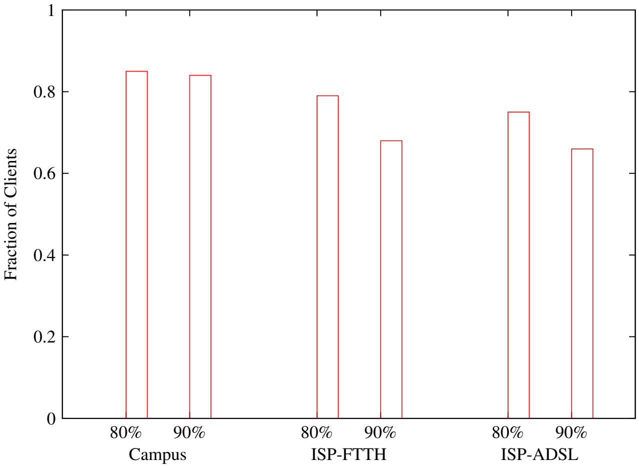

Fig. 3 Fraction of clients in *Campus-Community* and *ISP-Community* with a community usage ratio over 80% (90%)

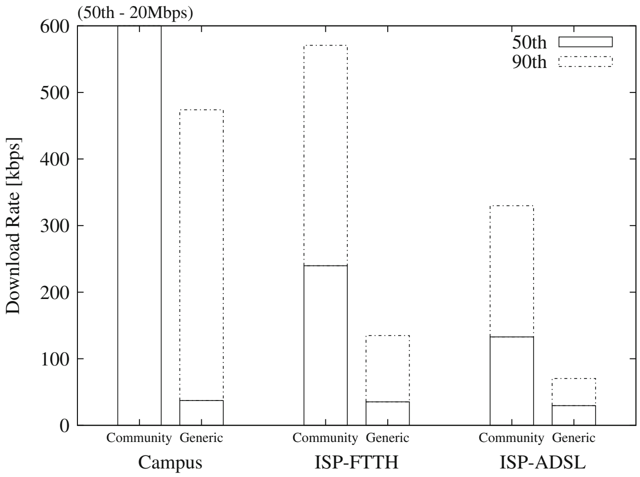

Fig. 4 Download rates for various systems

Figure 4 presents results for users' performance considering the per host download rate $D(h)$, which is the average download rate across all TCP connections initiated or received by a host.[1] Again, the median and 90th percentile of the distribution among hosts is reported. We observe that for the campus setting, the median improvement in download rate of the P2P community over the generic system is a factor of 740. However, the improvement is only a factor of 6 and 4 for the *ISP-FTTH* and *ISP-ADSL* settings respectively.

Impact of access technology We further study the ISP setting to understand the impact that heterogeneity in access technology of users has on the download performance. To provide appropriate context, Table 4 shows the breakdown of ISP users by access technology. For each trace, the first column shows the total number of external contacted peers. The second and third columns detail the total number of such peers connected by FTTH or ADSL links respectively. The fourth column shows peers for which we do not have access technology information. Neglecting the latter group, we notice that for both traces, there is a ratio of 1:5 between high speed FTTH links and slower and more congested ADSL links, which reflects the provider technology penetration.

Table 4 Access technology of *ISP-Community* peers

Trace	Total	FTTH	ADSL	Unknown
ISP-FTTH	497.4 K	80.8 K	356.4 K	59 K
ISP-ADSL	325.1 K	52.8 K	270.8 K	1013

Figure 5 shows the Cumulative Distribution Function (CDF) of the per-connection download bitrate for different types of sources contacted by *ISP-Community* clients in the *ISP-ADSL* trace. The curves labeled *ADSL Sources* and *FTTH Sources* refer to sources located inside the ISP connected by a particular access technology. The third curve labeled as *Internet Sources* corresponds to *ISP-Generic* clients, and is shown for comparison purposes. This curve refers to connections served by sources outside the ISP for which no information about the peer access technology is available. The key observation from this graph is that the performance benefits of using *ISP-Community* is due to the high capacity FTTH sources which clearly enable much higher download rates. The 50th (90th) percentile of the download bitrate from FTTH sources is 120 Kbps (591 Kbps) compared to 23 Kbps (64 Kbps) for Internet sources, an improvement of five to nine times. Interestingly, only the top 20 percentile of connections involving ADSL sources perform better than connections involving Internet sources, and even here the difference in performance is minor. The same observations hold for connections from the *ISP-FTTH* trace. In particular, when FTTH clients download from FTTH sources, the median throughput is 10 times higher than when they download from both ADSL and Internet sources.

Degree of seed-like behavior in communities We next investigate the extent to which users in the P2P community contribute more data than they receive and how this depends on the access technology of the user. We consider the ratio of bytes sent and bytes received per client in the whole dataset, which we call $R(p)$, where p is a peer. Notice that a value less than one

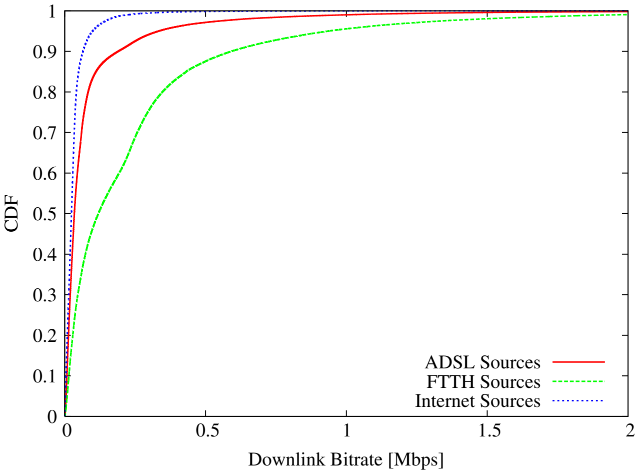

Fig. 5 Download rates from the *ISP-ADSL* trace, distinguishing the access technology of the source

---

[1] To avoid connections carrying only control messages, we consider connections with more than 50KB downloaded.

represents clients that are mostly receiving (also called leechers) and a value larger than one represents clients that are mostly sending (or seeds); a value close to 1 represent clients that both send and receive data in similar amounts.

Figure 6 shows the CDF of $R(p)$ in *ISP-Community* for the *ISP-FTTH* and the *ISP-ADSL* traces. We clearly see that *ISP-Community* users behind an FTTH link have a more seed-like behavior. For instance, the median $R(p)$ for FTTH clients is 2.05 which shows that half of the clients send twice as much as they receive. In contrast, the median $R(p)$ for ADSL clients is 0.56 which implies a leecher-like behavior.

Implications These results combined indicate that much of the performance improvement seen by *ISP-Community* clients over *ISP-Generic* clients stems from the fact that a small portion of the users in the ISP are connected by a high-speed FTTH technology, and these users contribute much more data than they receive. In addition, these results have broader significance for research on localizing P2P traffic within ISP boundaries [8, 27]. While most research in this area has taken for granted that localizing traffic benefits users and ISPs, our results suggest that in ISPs with heterogeneous access technologies, the performance benefits to users on localizing P2P traffic could depend on the degree of seed-like behavior or peers behind high-bandwidth access technologies. We elaborate further on these implications in Section 8.

5.2 Delay

An intuitive way to evaluate the effectiveness of P2P localization is by measuring the delay of localized connections. One would expect that connections in the community system will exhibit a lower RTT than connections in the generic system. Figure 7 shows the per connection average RTT for the three vantage points. We observe that while the difference in RTT between *Campus-Community* and *Campus-Generic* is prominent, it is less noticeable between *ISP-Community* and *ISP-Generic*. To understand this better, we consider the impact of ISP clients behind ADSL lines on the RTT of connections. We observe that the distribution of RTT for *ISP-Community* connections to peers behind ADSL access links is comparable to *ISP-Generic* flows, while a significant improvement is noticeable for sources behind FTTH access links (which have a ten times higher upload capacity). This implies that congestion in the upstream direction of ADSL clients is causing queueing delays which subsequently increases RTTs.

A second reason for the lack of RTT improvements between *ISP-Community* and *ISP-Generic* is that many clients in *ISP-Generic* tend to access content within the same small European country. For instance, in our *ISP-ADSL* dataset, we found that 52% of connections leaving the ISP were destined to the same country and 81% to Europe. More generally, this is because users in the same geographic region tend to be interested in the same content, due to common language and culture [9]. We also note that *Campus-Community* users are mostly English speaking users and thus the content tends to be more spread throughout the world.

Implications One could expect that localization of P2P traffic would reduce RTT and increase throughput of connections. But, our results show that these benefits may be limited by the access bandwidth of users inside

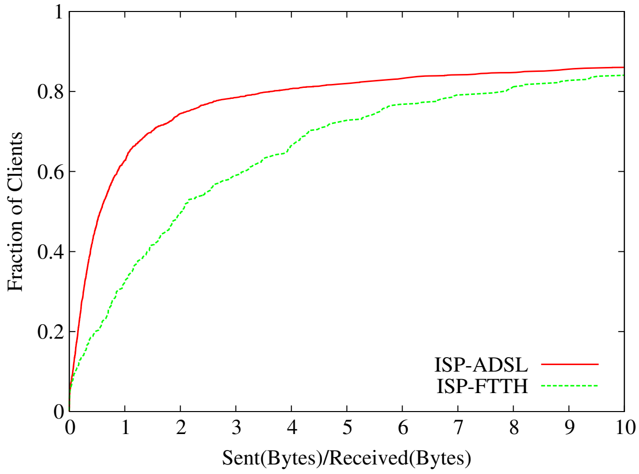

Fig. 6 Ratio of bytes sent and bytes received per client in *ISP-Community*

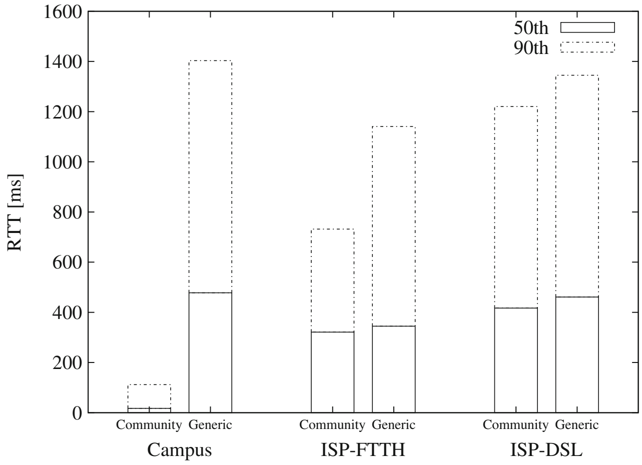

Fig. 7 RTT for connections initiated for various system

the ISP and user demographics. Our observations agree with findings by other researchers [9].

6 Traffic matrix for P2P communities

In Section 5, we have focused on the extent to which user performance is improved when P2P communities are used. We next consider the impact that P2P communities have on network providers. We focus our study on the ISP network given it is a larger and more interesting setting, but a similar methodology could also be employed for the campus network. In this section, we present an approach to infer the application-specific traffic matrix due to *ISP-Community*. Each cell of the matrix corresponds to the volume of traffic related to *ISP-Community* between a pair of PoPs in the ISP. Inference is required because we have direct measurements available only at two PoPs. We then employ this traffic matrix in Section 7 to study the load induced by *ISP-Community* on links of the ISP network, and examine how the load may change under various "what-if" scenarios (e.g., upgrade to higher access technology).

Traffic matrix estimation is a well studied problem [3, 29]. However, past work has primarily focused on the estimation of the overall traffic matrix. In contrast, we explicitly target the estimation of the subset of traffic due to the P2P community, and specifically due to *ISP-Community*.

We leverage on the *Simple Gravity Model* [29], according to which the amount of flow exchanged between two objects is proportional to their "size". We consider the network PoPs as objects, whose size is determined by their peer population. Then, the traffic $T_{\text{sent}}(s, d)$ sent from PoP s to PoP d is simply defined by:

$$T_{\text{sent}}(s, d) = T_{\text{sent}}(s) * \frac{population(d)}{\sum_{k=1}^{n} population(k)} \tag{1}$$

where $T_{\text{sent}}(s)$ is the total P2P traffic sent by users in PoP s, and $population(d)$ is the population of PoP d. n is the total number of PoPs present in the topology. The model assumes that the fraction of traffic from s to d is simply proportional to the relative population of d. We believe this assumption is reasonable if (i) content is uniformly available in each PoP; and (ii) peer selection follows a uniform probability. The first point can be justified, since in *ISP-Community* all users belong to an ISP within the same country and therefore content can be located anywhere in the network. The second point was verified to be the case in *ISP-Community*, since the modified version of the software selects peers with uniform probability (i.e. no tit-for-tat policy is in place). In a more general scenario, where P2P systems may preferentially select nodes in certain PoPs (e.g., PoPs with lower latencies), a "friction factor" may be introduced in the model. We leave this generalization for future work.

6.1 Validation

We now describe our methodology to validate the model and present our results. To validate Eq. 1, we will estimate the traffic sent from the *ISP-FTTH* PoP and the *ISP-ADSL* PoP to all other PoPs in the ISP network and will compare the estimates to the actual traffic volumes obtained from our traces.

To use Eq. 1, we require to know $T_{\text{sent}}(s)$ and the number of *ISP-Community* peers in each PoP, $population(d)$. While the former can be directly obtained from our traces, we assume the latter is simply the total number of *ISP-Community* peers in each PoP which contact our monitored *ISP-Community* peers through UDP control messages. This information can also be obtained directly from our traces. We believe this assumption is reasonable because: (i) Kad UDP control messages are sent to a larger number of peers compared to TCP data messages; and (ii) since Kad maps hosts to a DHT network at random, control messages are sent to any destination in the ISP with equal likelihood. To further confirm the validity of our methodology, we found that our monitored *ISP-Community* peers were contacted using UDP control messages by 497.4 K unique peers in the *ISP-FTTH* trace and 325.1 K unique peers in the *ISP-ADSL* trace. This is a very representative subset of the total *ISP-Community* population in the ISP. In addition, we noticed that the fraction of the population remains constant for every PoP across our traces and over different periods of time within a single trace. Hence, it may be enough to know the fraction of users per PoP rather than the actual total population of users.

Now that we have a way to obtain $T_{\text{sent}}(s)$ and $population(d)$, we can use Eq. 1 to predict the amount of traffic sent to any other PoP in the network from the *ISP-FTTH* PoP and the *ISP-ADSL* PoP. To further show the validity of our model, when we consider the *ISP-FTTH* PoP as the source PoP, the population is estimated using the *ISP-ADSL* dataset and when we consider the *ISP-ADSL* PoP as the source PoP, the population is estimated using the *ISP-FTTH* dataset.

Figures 8 and 9 show the amount of traffic sent to any other PoP in the network considering the *ISP-FTTH* and *ISP-ADSL* PoPs as the source PoP respectively.

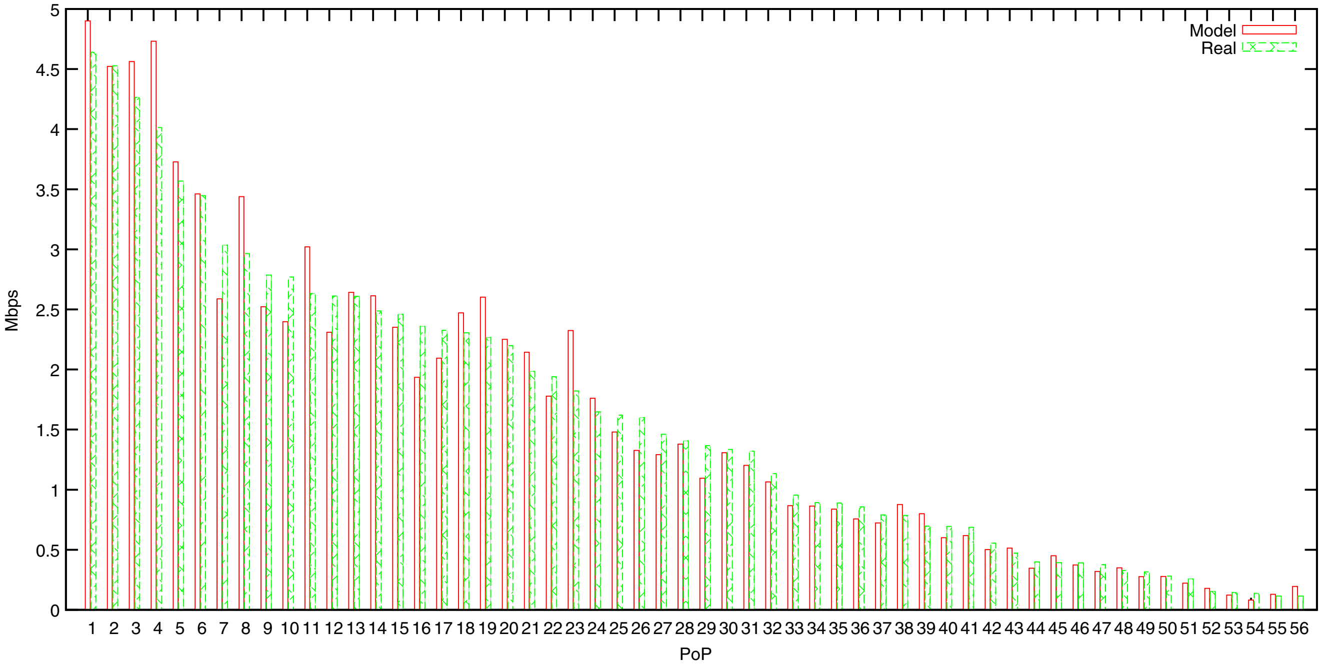

Fig. 8 Comparison of output from the gravity model and the real data for the *ISP-FTTH* trace

There is one set of bars for each of the destination PoPs and the destination PoPs have been sorted in descending order of traffic volume received. The solid (red) bar represents the model prediction, $T(s,d)$, and the dotted (green) bar represents the actual data measured from the trace, $\hat{T}(s,d)$. We observe that the model closely follows the real data in both cases.

To better quantify how close the output of the model is to the real data, we have calculated the relative error of the predictions, defined as $\frac{T(s,d)-\hat{T}(s,d)}{\hat{T}(s,d)}$. Figure 10

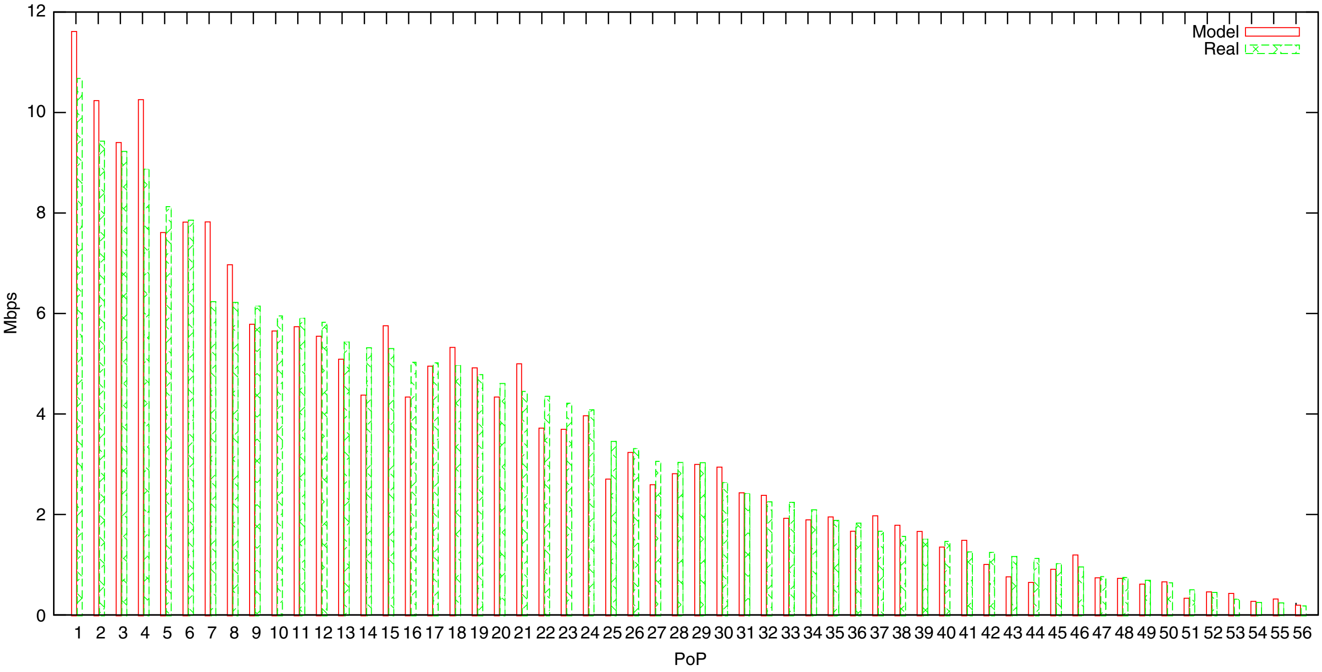

Fig. 9 Comparison of output from the gravity model and the real data for the *ISP-ADSL* trace

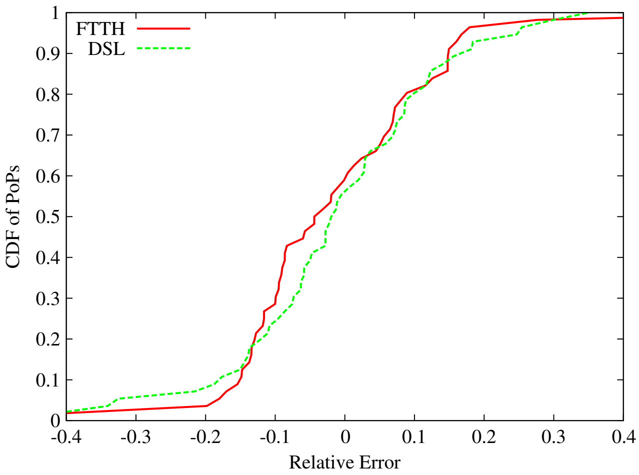

Fig. 10 Relative error of the gravity model for the *ISP-ADSL* and the *ISP-FTTH* traces

shows the CDF of the relative error for the *ISP-FTTH* and the *ISP-ADSL* traces. For around 90% of the PoPs, the error is smaller than 20% for both traces, which is an acceptable error margin.

These results confirm that a simple gravity model can be used to predict the traffic matrix of *ISP-Community*.

7 Impact of communities on the network

In Section 6, we developed an approach to estimate the traffic matrix specific to *ISP-Community*. In this section, we use the estimated traffic matrix to compute the amount of *ISP-Community* traffic carried by individual network links. We also compute the difference in traffic on individual links under various "what-if" scenarios. In particular: (i) we consider a hypothetical scenario in which all clients currently using *ISP-Community* switch to using *ISP-Generic*. The purpose of this scenario is to shed light on how localizing traffic within the ISP impacts the capacity planning decisions of the ISP; and (ii) we evaluate how an upgrade of the access technology of customers may impact the network. We consider a scenario in which all customers in the ISP are upgraded to an FTTH access technology.

To conduct our analysis, we require knowledge of the ISP topology, the routing algorithm, and the traffic matrix corresponding to each scenario. In the rest of this section, we elaborate on our approach to modeling each of these aspects, and present simulation results.

7.1 Modeling approach

Topology and routing We model the topology based on the actual ISP network through discussions with the operator. In particular, nodes in the topology represent both the PoPs to which customers are connected, and the ISP backbone routers. Four types of links are present in the modeled topology: (i) *PoP-to-PoP* links directly connecting two PoPs in the same city, (ii) *PoP-to-backbone* links connecting a PoP to a backbone router, (iii) *backbone* links connecting two backbone routers typically between two cities and (iv) *peering* links that connect some backbone routers to the Internet. Traffic is routed on the topology using the standard shortest path algorithm, commonly employed in networks today, including the ISP that we consider.

Traffic Matrix estimation from the dataset In Section 6 we have shown that the *ISP-Community* traffic exchanged by PoPs follows a simple gravity model. We leverage this to generate different traffic matrices to model possible scenarios. Besides the knowledge of the population of peers per PoP, Eq. 1 relies on the availability of the total traffic sent $T_{\text{sent}}(s)$ by a given PoP s, which we can directly measure for only two PoPs. Following the gravity model assumption, we model $T_{\text{sent}}(s)$ as directly proportional to the population of peers in PoP s.

More in detail, let $n_f(s)$ and $n_a(s)$ be the number of FTTH and ADSL users in PoP s. Let $\hat{t}_f(s)$ and $\hat{t}_a(s)$ be the average amount of traffic that an FTTH and ADSL user in s generates during a given time interval respectively. Assuming that users corresponding to each PoP generate the same amount of average traffic, $\hat{t}_f(s) = \hat{t}_f \; \forall s$ and $\hat{t}_a(s) = \hat{t}_a \; \forall s$.

$\hat{t}_f$ and $\hat{t}_a$ can then be estimated by considering the *ISP-FTTH* and *ISP-ADSL* datasets. Finally, the total volume of traffic sent from PoP s is simply proportional to the mix of access technology peers in the PoP, i.e.,

$$T_{\text{sent}}(s) = n_f(s)\hat{t}_f + n_a(s)\hat{t}_a \qquad (2)$$

Equation 2 is then used to derive the *ISP-Community* traffic matrix T.

7.2 Predicting traffic matrix changes

We now consider traffic matrices that P2P systems generate in different scenarios. We first consider the hypothetical scenario in which users in an ISP switch from *ISP-Community* to *ISP-Generic*. We also consider scenarios in which the ISP upgrades the access link of peers.

Users switch from ISP-Community to ISP-Generic To evaluate this, we must construct the *ISP-Generic* traffic matrix T'. This is the scenario in which all the community traffic is directed to the Internet peering node z.

We assume that the volume of traffic received by internal peers in this scenario is the same as in the current scenario, i.e., users are willing to download the same amount of data. More formally, $T'_{\mathrm{send}}(z,d) = \sum_{s\neq z} T_{\mathrm{send}}(s,d)$ and $T'_{\mathrm{send}}(s,d) = 0\ \forall d \neq z$.

In addition, we assume that the amount of traffic each PoP sends to z is the same as currently observed in the *ISP-Generic* system. We believe this is reasonable because independent of the upload capacity of peers inside the ISP, once connections leave the network, their throughput is likely to be limited by the destination access link or by some congested or rate limited intermediate link. More formally, $T'_{\mathrm{send}}(s,z) = T'_{\mathrm{send}}(s)$, where $T'_{\mathrm{send}}(s)$ is computed as in Eq. 2 considering the *ISP-Generic* dataset to estimate $\hat{t}_a$ and $\hat{t}_f$.

Technology upgrade To consider the technology upgrade from the current ADSL and FTTH mix to an all-FTTH scenario, we assume that all ADSL clients send and receive *ISP-Community* traffic at the FTTH rate. Therefore $T''_{\mathrm{send}}(s) = n_a(s)\hat{t}_f + n_f(s)\hat{t}_f$ and the traffic matrix T'' is computed as in Eq. 1.

Finally, we consider a scenario in which users switch to the *ISP-Generic* system, while access technology is upgraded to FTTH. This allows us to compare the traffic on individual network links with *ISP-Generic* and *ISP-Community*, in an upgraded access technology setting. The traffic matrix is then obtained similarly to T', but $\hat{t}_a = \hat{t}_f$ are estimated as in T''.

7.3 Results

Figure 11 shows the CDF of the volume of P2P community related traffic that traverses each ISP link. Four types of links are considered in the figure: backbone, backbone-to-PoP, PoP-to-PoP and peering. For ease of presentation we call *GenericOnly* to the scenario where only *ISP-Generic* is present in the ISP, *CurrentTech* to today's mix of access technologies in the ISP and *AllFTTH* to the technology upgrade to an all-FTTH scenario. There are four lines, one for each combination of *ISP-Community* or *GenericOnly* for the P2P system in the network, and *CurrentTech* or *AllFTTH* for the access technology of clients. Logarithmic x-scale is used to better highlight the differences. We draw several observations from this plot. First, as expected, the usage of *ISP-Community* greatly reduces the traffic at peering links. However, more surprisingly, over 90% of the links carry a larger amount of traffic in the presence of *ISP-Community* as compared to the *GenericOnly* scenario. For instance, the median of *ISP-Community-CurrentTech* is 4.6 times larger than the median of *GenericOnly-CurrentTech*. Second, notice that *ISP-Community* makes use of more links in the network. While in the *GenericOnly* scenarios, 30% of the links are unused (mostly PoP-to-PoP links), in the *CurrentTech* scenarios, more than 95% of the links are being used and most of them are carrying more than 100GB per day. Third, when all peers are upgraded to FTTH, the traffic on links increases by almost an order of magnitude when comparing *ISP-Community*-CurrentTech and *ISP-Community*-AllFTTH. This is due to the higher upload capacity of peers. Hence, enhancing the capacity of ISP peers to get the full benefits of localization can probably hurt the network.

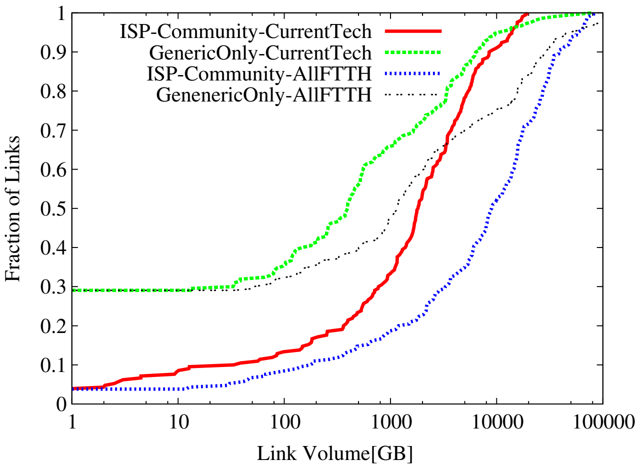

Fig. 11 CDF of volume of traffic related to P2P communities, seen by all links in the ISP topology, for the *ISP-Community* and *GenericOnly*, varying the access technology of peers

To get more insight into this, Fig. 12 shows the difference in volume per link comparing the *GenericOnly-CurrentTech* and *ISP-Community-CurrentTech* scenarios. A negative value corresponds to an increase of link load when *ISP-Community* is used, while a positive value indicates an increase in link load with *ISP-Generic*. Results are separately reported for each class of links, i.e., backbone (top left plot), backbone-to-PoP (top right plot), PoP-to-PoP (bottom left plot) and peering links (bottom right plot) respectively. We draw several observations. First, we notice that more than 60% of backbone links see an increase in the volume of traffic they have to carry in the *ISP-Community* scenario, with some links seeing as much as $2 * 10^4$ GB of additional traffic each day. We make a similar observation for backbone-to-PoP links and for PoP-to-PoP links. However, for the latter, the increase in traffic is not as large in the *ISP-Community* scenario since those links interconnect PoPs which may exchange small amounts of data. Finally, as expected, *ISP-Community* is able to reduce the traffic at the peering links. We note that on average, peering links

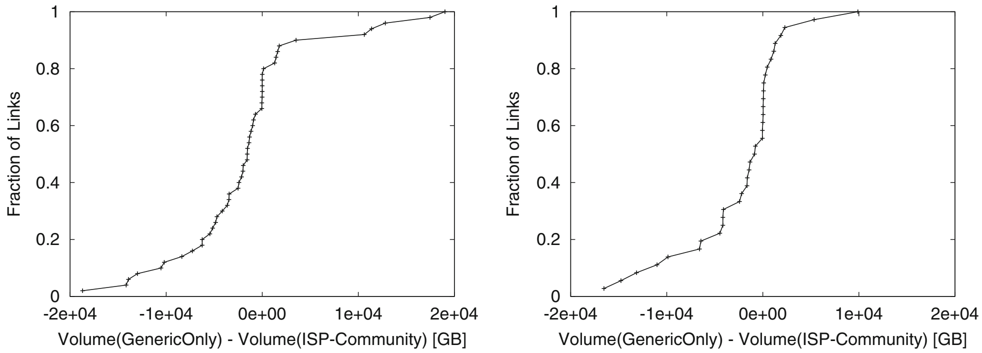

(a) Backbone Links

(b) Backbone to PoP Links

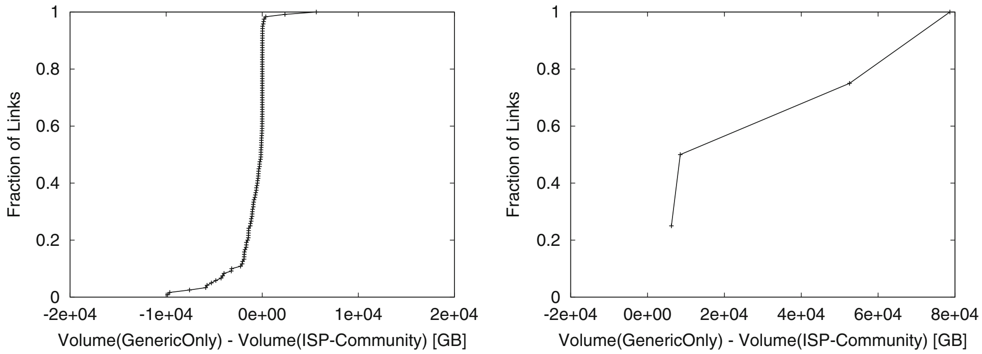

(c) PoP to PoP Links

(d) Peering Links

Fig. 12 Difference in traffic per link considering the *ISP-Community-CurrentTech* and *GenericOnly-CurrentTech* scenarios. There is one plot for each link type

have to transport 36.5TB less traffic per day, or 1.5TB less traffic each hour in the *ISP-Community* scenario.

Implications Overall, these results show that the extensive use of the P2P community has very significant implications for network providers in terms of traffic that individual links carry. While the use of the P2P community reduces traffic at peering points, it greatly increases traffic on interior links of the network, and consequently has important implications for capacity planning algorithms of the network provider.

8 Discussion and implications

We discus key aspects of our work and implications below:

Implications for P2P traffic localization Our characterization of closed P2P communities offer important lessons for research on localizing traffic of generic P2P systems within ISP boundaries [2, 6, 8, 10, 14, 16, 23, 25, 27]. While a majority of research in this area has taken the benefits of localization on users and ISPs for granted, our results support recent works [9, 19] which argue that a more critical examination is essential.

First, our results suggest that the benefits of localization could depend on the access bandwidth of peers inside the ISP, as pointed out by recent works [9, 19]. For instance, while the throughput of connections of *Campus-Community* is significantly improved due to the high bandwidth campus LAN, the improvement in throughput of *ISP-Community* connections is limited by the fact that 80% of the users are behind an ADSL link. Further, the RTT of *ISP-Community* connections does not show significant improvement compared to *ISP-Generic* due to both the access technology of users, as well the fact that most users of *ISP-Generic* tend to access content from the same European country.

Second, going beyond [9, 19], our results suggest that in ISPs with heterogeneous access technologies, the performance benefit to users on localizing P2P traffic could depend on the *degree of seed-like behavior* of peers behind high-bandwidth access links. For instance, the performance with *ISP-Community* is better than with *ISP-Generic* primarily due to a small number of users behind high-bandwidth FTTH connections, which contribute much more data than they receive. These observations also imply that in ISPs with *heterogeneous* access technologies, not all stake-holders can simultaneously win when P2P traffic is localized. In particular, (i) if high-bandwidth users show seed-like behavior, then the ISP and users behind low-bandwidth access technologies benefit on localization (through reduced transit traffic and improved performance), at the expense of high-bandwidth users; and (ii) alternately, localization benefits the ISP alone, and does not help improve the performance of users (in particular, users behind low-bandwidth access technologies). This situation can potentially be offset through new models for charging for access, where ISPs can incentivize seed-like behavior of high-bandwidth users by reducing charges on them.

Finally, our simulation results show that use of *ISP-Community* rather than *ISP-Generic* does result in lowered traffic on peering links for the ISP as expected. However, more interestingly, over 90% of internal links of the ISP network see higher traffic, with some links seeing as much as $2 * 10^4$ GB of additional traffic each day. This increase may potentially require internal link upgrades to avoid impairing user performance. These results suggest that coarse-grained schemes for localizing P2P traffic like *ISP-Community* (where peers within an ISP are selected at random), may potentially cause significant shifts in ISP traffic patterns. Developing techniques to predict such shifts is an interesting area of future research and we have taken a first step in this direction. Studying the effect of traffic shifts with more fine-grained localization schemes (for e.g., [23]) is another direction for future work.

Implications for traffic studies suggesting decline of P2P traffic in the Internet Recent work has suggested that P2P traffic is on the decline [15]. However, this finding is based on analysis of inter-AS traffic data, and traffic internal to ISP networks is not considered. In our traces however, private communities account for 60–90% of all traffic. This indicates that traffic related to such private communities should be taken into consideration before making conclusive statements about the decline of P2P traffic.

Generality of our work In this paper, we have characterized P2P communities in two networks. We believe our work is an important start in creating an awareness and understanding of such communities, an area that has received little attention to date. Analyzing a wider range of communities is a challenge given that this requires knowledge of which networks contain P2P communities, and involves traffic collection inside each of the networks. Obtaining access to traffic data from additional networks that contain P2P communities, and analyzing P2P communities in a broader range of networks is an important area for future research, and a subject of our ongoing investigations.

9 Related work

We have already extensively discussed how our work relates to work on P2P traffic localization in Section 8.

There has been awareness in the research community about the presence of darknets [4, 5, 21, 28]. While darknets are related to our work in that they are also typically closed P2P communities, they share important differences from the types of communities we consider. Darknets are motivated by the primary goal of anonymized sharing of content. In contrast, P2P communities are motivated by other factors such as ensuring good application performance and ensuring hosts with private addresses in the same ISP may communicate. Further, we focus on communities localized to particular networks, while darknets could extend across the Internet. A notable recent work [28] characterized operational BitTorrent darknets, focusing the analysis on identifying the type of content being shared and the level of collaboration between peers. In contrast, our focus is on P2P communities, performance seen by users of the community and the impact of the community traffic patterns on the service provider.

Our own prior work (self citation) has pointed out the presence of the closed community in the ISP setting. However, the focus of that work was on detecting undesirable traffic patterns of a range of P2P systems, including the P2P community. In contrast, our focus in this work is on characterizing the performance and network localization of these systems and simulations to study implications on network providers. Further, we consider a completely new campus network community in this paper.

Models for traffic matrix estimation (for e.g., [3, 29]) have been widely studied. Our efforts are distinguished in that we focus on developing an application-specific traffic matrix model that only considers traffic due to the P2P community application.

10 Conclusions

In this paper, we have raised the awareness of the research community to the prevalence of closed P2P communities, and have presented one of the first and most extensive characterizations of such communities. P2P communities are the most popular P2P systems used and generate most of the traffic in the networks we considered. For instance, we identified about 600,000 unique peers in *ISP-Community*, that exchange 50 times more traffic than *ISP-Generic* and accounts for 60% to 90% of all the traffic observed in our traces. While as expected, users of P2P communities see better performance than users of generic P2P systems, we have shown that the extent of benefits is could be determined by the access technologies of the users, and the degree of seed-like behavior shown by users behind high-speed access technologies. We have developed techniques to enable network providers understand how the growth of P2P communities impacts network capacity planning. and how projected changes in access technologies of users may affect these results. Using the techniques, we show that while use of the communities does lower traffic on peering links compared to the use of generic P2P systems, it could greatly increase traffic on internal network links. Our characterization of P2P communities while interesting in its own right, offers important lessons for research on localizing P2P traffic within ISP boundaries. Our future work includes studies of a wider range of P2P communities, and exploring refinements to the P2P community traffic models that can result in greater accuracy.

References

1. ADC (2008) Adc protocol. http://adc.sourceforge.net/ADC.html
2. Aggarwal V, Feldmann A, Scheideler C (2007) Can ISPs and P2P users cooperate for improved performance? SIGCOMM Comput Commun Rev 37:29–40
3. Alderson D, Chang H, Roughan M, Uhlig S, Willinger W (2006) The Many Facets of Internet Topology and Traffic. The American Institute for Mathematical Sciences. Journal of Networks and Heterogeneous Media 1(4):569–600
4. Bethencourt J, Low WY, Simmons I, Williamson M (2007) Establishing darknet connections: an evaluation of usability and security. In: Proceedings of the 3rd symposium on usable privacy and security, SOUPS ‘07. ACM, New York, NY, USA, pp 145–146. doi:10.1145/1280680.1280700
5. Biddle P, England P, Peinado M, Willman B (2002) The darknet and the future of content distribution. In: DRM
6. Bindal R, Cao P, Chan W (2006) Improving traffic locality in BitTorrent via biased neighbor selection. In: ICDCS
7. BitTorrent (2008) http://www.bittorrent.org
8. Choffnes DR, Bustamante FE (2008) Taming the torrent: a practical approach to reducing cross-ISP traffic in peer-to-peer systems. In: SIGCOMM
9. Cuevas R, Laoutaris N, Yang X, Siganos G, Rodriguez P (2009) Deep diving into bittorrent locality. CoRR abs/0907.3874
10. ETSI (2003) ETSI TISPAN—defining the next generation network. http://www.etsi.org/tispan/
11. eMule (2008) http://www.emule-project.net
12. Gummadi K, Dunn R, Saroiu S, Gribble S, Levy H, Zahorjan J (2003) Measurement, modeling, and analysis of a peer-to-peer file-sharing workload. In: SOSP
13. Guo L, Chen S, Xiao Z, Tan E, Ding X, Zhang X (2005) Measurements, analysis, and modeling of BitTorrent-like systems. In: IMC
14. Karagiannis T, Rodriguez P, Papagiannaki K (2005) Should Internet Service Providers Fear Peer-Assisted Content Distribution? In: IMC
15. Labovitz C, Iekel-Johnson S, McPherson D, Oberheide J, Jahanian F (2010) Internet inter-domain traffic. In: Proceedings of the ACM SIGCOMM 2010 conference. ACM, New York, NY, USA, pp 75–86. doi:10.1145/1851182.1851194
16. Liu B, Cui Y, Lu Y, Xue Y (2009) Locality-awareness in BitTorrent-like P2P applications. IEEE/ACM Trans Multimedia
17. Mellia M, Cigno RL, Neri F (2005) Measuring IP and TCP behavior on edge nodes with Tstat. Comput Networks 47(1):1–21. doi:10.1016/j.comnet.2004.06.026
18. Mellia M, Meo M, Muscariello L, Rossi D (2008) Passive Analysis of TCP Anomalies. Comput Netw 52(14):2663–2676. doi:10.1016/j.comnet.2008.05.010
19. Piatek M, Madhyastha H, John J, Krishnamurthy A, Anderson T (2009) Pitfalls for isp-friendly p2p design. In: HotNets-VIII
20. Pietrzyk M, Costeux JL, Urvoy-Keller G, En-Najjary T (2009) Challenging statistical classification for operational usage: the adsl case. In: IMC ’09: Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference. ACM, New York, NY, USA, pp 122–135. doi:10.1145/1644893.1644908
21. Rogers M, Bhatti S (2007) How to disappear completely: a survey of private peer-to-peer networks. In: SPACE
22. Saroiu S, Gummadi K, Gribble S (2003) Measuring and analyzing the characteristics of Napster and Gnutella hosts. In: Multimedia System
23. Seedorf J, Kiesel S, Stiemerling M (2009) Traffic localization for P2P-applications: the ALTO approach. In: IEEE P2P
24. Sen S, Wang J (2002) Analyzing peer-to-peer traffic across large networks. In: IMW
25. Shen G, Wang Y, Xiong Y, Zhao B, Zhang Z (2007) HPTP: relieving the tension between ISPs and P2P. In: IPTPS
26. Tstat (2008) http://www.tstat.polito.it
27. Xie H, Yang YR, Krishnamurthy A, Liu Y, Silberschatz A (2008) P4P: provider portal for applications. In: SIGCOMM
28. Zhang C, Dhungel P, Wu D, Liu Z, Ross K (2010) BitTorrent Darknets. In: INFOCOM
29. Zhang Y, Roughan M, Duffield N, Greenberg A (2003) Fast accurate computation of large-scale IP traffic matrices from link loads. In: SIGMETRICS

Ruben Torres is a Ph.D. student in the School of Electrical and Computer Engineering at Purdue University and a research assistant in the Internet Systems Lab. Ruben received his MS in electrical and computer engineering from Purdue University in 2006 and his BS in electronics and communications engineering from the University of Panama in 2002. His research interests include network systems, distributed systems and 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. In 2002 he visited the Sprint Laboratories in Burlingame (CA), while in 2011 he was with Narus inc., Sunnyvale (CA). He coauthored over 160 papers published in international journals and conferences, and he served as a member of the Technical Program Committees of several conferences including IEEE Infocom and ACM Sigcomm. His research interests are about traffic measurement, P2P applications and energy aware network design.

Maurizio M. Munafo is an Assistant Professor in the Department of Electronics and Telecommunications 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.

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 has spanned several areas including enterprise network management, cloud computing, Internet video distribution, and peer-to-peer systems. He received a B.Tech in Computer Science and Engineering from the Indian Institute of Technology, Madras, and the Ph.D from the School of Computer Science, Carnegie Mellon University. He has been a visiting research scientist in Google and AT&T Research. He is a recipient of the NSF Career award, and won the ACM SIGMETRICS Test of Time Award for his work on End System Multicast. 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 conduction with NSDI 2010.