Who Profits from Peer-to-Peer File-Sharing? Traffic Optimization Potential in BitTorrent Swarms

Valentin Burger*, Frank Lehrieder*, Tobias Hoßfeld*, Jan Seedorf‡

*University of Würzburg, Institute of Computer Science, Würzburg, Germany

‡NEC Laboratories Europe, Heidelberg, Germany

Abstract—**BitTorrent-based peer-to-peer networks constitute a significant share of Internet traffic. Therefore, the IETF working group on application layer traffic optimization (ALTO) discusses several approaches for guiding the BitTorrent traffic that aim at reducing the large amount of inter-ISP traffic costs caused by these networks. However, performance evaluations of these approaches are mostly limited to artificial scenarios that do not take into account the real Internet topology with its inter-ISP relationships and the actual distribution of BitTorrent users across autonomous systems (AS). In this study, we use measurements of the distribution of a large number of live BitTorrent networks and combine them with the AS-level Internet topology provided by Caida.org. Based on this data, we estimate in which tier of the Internet hierarchy BitTorrent traffic is mainly located and how much optimization potential exists for the different types of ISPs. Therewith, traffic flow and revenue implications of guiding Internet-wide BitTorrent swarms are analyzed. Our results show that tier–1 ISPs profit from the un-managed exchange of peer-to-peer traffic and that these profits significantly decrease when the other ISPs would apply ALTO solutions.**

I. INTRODUCTION

BitTorrent is still responsible for a large portion of Internet traffic [1], [2]. In particular, BitTorrent networks generate a lot of inter-ISP traffic, which is often costly for the ISPs. One approach to optimize the traffic flows, which has recently received a lot of attention is Application Layer Traffic Optimization (ALTO), i.e., *P2P guidance*, to increase the efficiency of BitTorrent and to reduce the amount of inter-ISP traffic and costs. Evaluations of such approaches have been conducted mostly in controlled, artificial scenarios. Examples for such scenarios are simulations with homogeneous peer distributions across ISPs, the evaluation of simple topologies, like the star topology with a tier–1 ISP in the center. However, in today's Internet the inter-ISP traffic routing is based on a complex topology defined by inter-ISP relationships, e.g., peering or customer-to-provider, and ISP classifications such as tier–1, large and small ISPs, and stub ISPs. Hence, these economic relations play an important role in the actual Internet traffic flow. However, this topology of the Internet is not taken into account by most evaluations of P2P guidance approaches, which limits the practical relevance of the results. Furthermore, it is an open question how much BitTorrent traffic is located

in which region of the Internet. However, this is a prerequisite in order to estimate the potential of ALTO mechanisms.

To model the BitTorrent traffic flow across ISPs in the Internet, we use measurements of live BitTorrent swarms and the actual autonomous system (AS) topology of the Internet provided by Caida.org. Thus, we estimate BitTorrent traffic characteristics and the emerging transit costs. The measurements of live BitTorrent swarms contain the location of peers in the Internet, i.e., AS numbers, for a very large set of swarms [3]. The dataset from Caida.org contains a full AS graph derived from RouteViews BGP table snapshots, including the AS relations. We infer AS paths based on this dataset with the algorithm published in [4]. We use the inferred AS paths to calculate the real AS paths between peers in BitTorrent swarms. In addition, we define three peer selection strategies that decide which peer in a swarm is connected to which other peer: (a) The random selection strategy is applied to peer selection of today's BitTorrent clients. (b) The locality aware selection strategy connects to the peers with shortest AS paths, to reduce AS hops and potentially latencies in the BitTorrent network. (c) The selfish-ISP selection connects peers preferentially to peers in the ISPs customer tree in order to maximize its revenue. The locality aware and selfish-ISP selection strategies are motivated by the optimization potential of the BitTorrent overlay network, as in [3], and the optimization potential of the revenue of ISPs transit services, respectively.

The contribution of this paper is two-fold. First, our results show how BitTorrent networks are distributed over the Internet. We find that almost none of the peers are located in tier–1 ASes which means that tier–1 ASes are not able to control BitTorrent swarms by directly accessing the peers. We analyze the amount of BitTorrent traffic each AS forwards with BitTorrent random peer selection. From our results we derive that most traffic is forwarded by tier–1 ASes on a per AS basis, whereas the BitTorrent traffic aggregated over all large ISPs is significantly higher. As second contribution, we analyze the potential to optimize the BitTorrent overlay by using the shortest AS paths. We find that in about 15 % of the investigated swarms, peers exist who can exchange data locally in the same AS. Furthermore, the AS path length has a median of two AS hops for random selection whereas AS paths have two or less AS hops in 80 % with the locality selection strategy. The inter-AS traffic is reduced especially in tier–1 and in large

ISPs by locality aware peer selection. Finally, we estimate the potential of ISPs to optimize their revenue. We find that tier–1 ASes loose a lot of revenue if locality or selfish-ISP selection is used because they are avoided as provider. Small ISPs and stub ASes have large benefits from using locality because they can minimize their costs. This implies less revenues at tier–1 and large ISPs. Large ISPs have to use selfish-ISP selection to have a higher prospect on profit.

The paper is structured as follows. Section II describes the background of our study and presents related work. The applied methodology to characterize the BitTorrent traffic and to estimate transit costs is described in Section III. In Section IV we describe the results of this study and their importance for ISPs. We conclude this work in Section V.

II. BACKGROUND AND RELATED WORK

This section describes the background of our study and presents related work. We start with an overview on measurement studies of live BitTorrent networks and show different approaches to reduce inter-ISP traffic discussed in the ALTO working group of the IETF. Finally, we introduce studies that infer the inter-AS relations based on BGP routing information.

A. Measurements and Models of Live BitTorrent Networks

BitTorrent is a peer-to-peer file-sharing protocol, which is based on multi-source downloads between the users. All the users, i.e., *peers*, sharing the same file belong to a *swarm*. To join the swarm, a peer requests addresses of other peers at an index server called *tracker*. In the standard BitTorrent algorithm the tracker uses random peer selection to select a subset of peers that are in the swarm. Then, the joining peer tries to establish a neighbor relation to the peers it got from the tracker and collects all peers which accepted the request in his *neighbor* set. The peer signals interest to all neighbors which have parts of the file it still needs to download. To which neighbor a peer is willing to upload data is decided by the choking algorithm, which is explained in [5].

As basis of our methodology for modeling inter-ISP BitTorrent traffic, the results in [3] are revisited. In [3] the authors provide measurements of a large number of live BitTorrent swarms taken from popular index servers such as *The Pirate Bay*, *Mininova*, and *Demonoid*. Using the IP addresses of the peers, the authors associate every peer with its AS and estimate the potential of ALTO mechanisms based on the differentiation between local peers (peers in the same AS) and remote peers located in other ASes. In contrast, we consider the actual Internet topology in this work, i.e., the inter-ISP relations, the ISP classification in the Internet hierarchy, and the AS paths between the peers in order to estimate the optimization potential of ALTO mechanisms.

The authors of [6] use the peer exchange protocol (PEX) in order to measure the neighbor set of all peers participating in a number of live BitTorrent swarms. Based on this information, they model the graph topology of the swarms and compare the structure to random graphs. They also investigate clustering of peers within ASes and countries, but do not focus on inter-AS relations and AS paths between peers as we do in this work.

In addition, there are measurement studies that examine and model distinct features of BitTorrent networks. In [7], a single swarm was measured for five months with a focus on the download times of the peers. Additional parameters such as the peer inter-arrival times in the swarm, their upload capacity and their online time are considered in [8]. The authors of [9] investigate these parameters also in multiswarm scenarios. Finally, [10] measures 4.6 million torrents to provide an overview of the entire BitTorrent ecosystem with its different communities and index servers. Our study differs from these works in that it focuses on the location of the peers in the Internet and the AS paths between the peers.

B. ALTO Mechanisms and Their Performance Evaluation

Various mechanisms to reduce the inter-ISP traffic generated by BitTorrent and other P2P applications are currently being investigated. Besides caching of BitTorrent traffic [11]–[13], which might involve legal issues, changing standard BitTorrent algorithms is a promising approach. The authors of [14] propose to use an oracle service provided by the ISP guiding the peers in their peer selection process. The evaluation uses a Gnutella network and shows that intra-AS traffic is increased significantly without a negative impact on the overlay graph. Similar approaches are proposed for BitTorrent. Bindal et al. [15] reduce the inter-ISP traffic by modifying the neighbor set of the BitTorrent peers, which can be done at the tracker or enforced by the ISPs using deep packet inspection. Their simulations use a uniform peer distribution over ASes and show a high optimization potential of this approach. The authors of [16] propose to use *iTrackers* to guide the peers and formulates an optimization problem to find the best neighbor sets. Finally, Oechsner et al. [17] propose to change the choke algorithm of BitTorrent to further reduce inter-ISP traffic and evaluate it via simulations in homogeneous scenarios. The BitTorrent plugin *Ono* [18] uses the servers of content distribution networks (CDN) as landmarks and estimates the proximity of two peers by the similarity of the CDN redirection behavior.

The authors of [19] investigate analytically the capabilities of a P2P-based content distribution network and the impact of locality. In contrast to our work, they use traffic characteristics which arise from software updates and do not consider AS relationships. A set of evaluations of ALTO mechanisms uses scenarios inspired by measurements of live BitTorrent swarms [20]–[22]. The studied scenarios consider heterogeneous peer distributions where some ASes contain more peers of a specific swarm than others. Nevertheless, they do not take into account inter-AS relations and the AS paths between two peers. This is different in our study. Using the AS affiliation of peers and the data obtained from Caida.org, we infer the actual paths of the BitTorrent connections in the Internet. In addition, we focus on the inter-ISP relations and investigate to which degree selfish ISPs profit from recommending their peers to preferentially use connections to peers located in lower tier ASes.

C. Measurements of AS Relations and Topologies

Autonomous systems are individual parts of the Internet, which are operated by ISPs. On a technical level, the traffic exchange between the ASes is controlled by the Border Gateway Protocol (BGP) [23]. However, commercial relations between ISPs determine the routing policies configured via BGP. An ISP must buy transit services to access parts of the Internet it neither owns nor can access by its customers. Hence, to route traffic between autonomous systems ISPs engage in business relationships. These business relationships are usually not open for public but they can be abstracted into three common types [24]. The relationship between two ASes can be customer-to-provider (c2p), peer-to-peer (p2p) or sibling-to-sibling (s2s). A customer-to-provider link is present if the customer AS pays the provider AS for transit service, i.e., the provider forwards the traffic of the customer and its customers. In a peer-to-peer relation the ASes have an agreement that they exchange each others traffic and the traffic of their customers, without paying each other. Sibling-to-sibling are links between ASes of the same organization. These relations are defined in business agreements and kept secret, but they can be inferred by analyzing the routing between autonomous systems.

The approach that is most widely used to infer AS relationships is analyzing BGP routing tables. The data set used in this work is also produced by inferring BGP tables as described in [25]. Therefore, AS links are extracted from RouteViews BGP tables. First sibling-to-sibling links are identified by looking up organizations that own multiple AS numbers. Then customer-to-provider relationships are inferred by a heuristic that is based on the idea of relaxing the requirement for a maximal number of valid paths and using the AS degree information to detect paths that are invalid. Most challenging is the inference of peer-to-peer links since paths remain valid if peer-to-peer links are replaced by a customer-to-provider or provider-to-customer link. The authors of [25] develop a heuristic which combines the strengths of previous approaches by [24] and [26]. The inferred relationships were validated by surveys, showing that 96.5 % customer-to-provider, 82.8 % peer-to-peer, and 90.3 % sibling-to-sibling of the inferred relationships are correct.

III. METHOD FOR MODELING BITTORRENT TRAFFIC FLOW AND REVENUE OF ISPS

In this section we describe the methodology to estimate transit costs of ASes. First, we show where we obtain the AS affiliation of peers. Second, we explain how AS paths are inferred from AS relations and how to classify the ASes. Further on we describe different BitTorrent peer selection strategies determining the connection among peers in the Internet. Finally we introduce our transit cost model.

A. AS Affiliation of Peers

In order to know where peers are located and where BitTorrent swarms generate costs for ISPs, we need to know how the swarms are distributed over the Internet and in which ASes the peers are located. For that purpose, we use

TABLE I
CLASSIFICATION OF AUTONOMOUS SYSTEMS.

Type	Classification	#ASes
Tier–1	AS has no providers	11
Large ISP	AS customer tree ≥ 50	337
Small ISP	AS customer tree < 50 and ≥ 5	1770
Stub	AS customer tree < 5	36289

the dataset of BitTorrent movie torrents "Mov." provided by the authors of [3]. A snapshot of all available movie torrents on Mininova.org was taken. The swarm sizes and peer distributions were recorded by distributed measurements. The data set consists of files with AS number and number of peers pairs for each BitTorrent swarm. Hence, they provide information for each swarm on how many peers are located in which AS. The measurement took part in April 2009 and recorded 126,050 swarms. Peers of 8,492 ASes are present in the swarms.

B. AS Relations, Paths and Classification

To be able to estimate the transit costs produced by peers exchanging data in BitTorrent swarms, we need to know the AS paths that connect the peers. Datasets with complete AS paths would be very large and are not available to our knowledge. Hence, we infer AS paths from AS relationships. We use the AS relationship dataset from Caida.org [27]. The dataset contains AS links annotated with AS relations. Each file contains a full AS graph derived from RouteViews BGP table snapshots. For our estimations we use the dataset from January 2011. The dataset consists of customer-to-provider and peer-to-peer relations. However, sibling-to-sibling is only considered by Caida and does not occur in the dataset.

We implemented the algorithm described in [4] in Java to infer the AS paths between any two peers based on the AS relationship dataset. The authors developed a breadth first search algorithm which infers shortest paths conforming to the AS path constraints. The algorithm has runtime $O(N \cdot M)$ for finding all pair valid shortest AS paths of the graph, where N is the number of AS relations and M is the number of ASes. The algorithm's input parameter is the source AS α. For every destination AS β the algorithm returns a set of paths $\mathbb{P}(\alpha, \beta)$, which connect α and β.

Further on, we want to obtain results dependent on the AS size and type of business. Therefore we classify the ASes into *stub*, *small ISP*, *large ISP* and *tier–1*. For that purpose, we use a dataset from [28], which provides information about the number of customers and providers for each AS number. This dataset is from November 2011 and is used to classify the ASes according to the size of their customer tree. Tab. I lists the different AS types and their classification. Tier–1 ASes are the largest ASes building the core of the Internet. Tier–1 ASes do not have providers. In their dataset 11 tier–1 ASes are identified. If an AS has a customer tree that contains at least 50 nodes, it is classified as a large ISP. An AS is classified as small ISP if its customer tree has less than 50 but at least 5 nodes. Most ASes are stub ASes, which have a customer tree that is smaller than 5.

C. BitTorrent Neighbor Set Creation

The BitTorrent neighbor set of a peer defines the data exchange with other peers in BitTorrent swarms. Neighbors are the peers in the swarm which are connected to a peer. It has to be noted that the measurements in [3] do not reveal the real composition of the neighbor sets of the swarms. Further on, neighbor sets are randomly generated and differ for every peer, which makes them hard to capture. Hence, we estimate the composition of the neighbor sets in three simple ways, *random*, *locality* and *selfish-ISP*. The number of peers in the swarm is the swarm size S. The number of neighbors is denoted as N with $N \leq S - 1$. In the standard BitTorrent implementation a client can connect to up to 40 peers, so we set the maximum size of the neighbor set to $N_{max} = 40$. Hence, the neighbor set for each peer in a swarm with size S has size

$$N = min(N_{max}, S - 1). \tag{1}$$

We add peers to the neighbor set until it contains N neighbors according to the following algorithms.

1) random: In the random selection strategy we add random N peers of the swarm to the neighbor set. In the standard BitTorrent algorithm the selection of neighbors is also random. Hence, with this selection strategy we try to estimate the traffic and costs produced by the standard BitTorrent algorithm.

2) locality: In the locality algorithm we sort the AS paths connecting two peers by the number of AS hops. Then we add the peers according to the sorted set of increasing AS paths until N peers are in the neighbor set. Note, that first the peers located in the same AS, i.e., zero AS hops, are added. This selection algorithm is used to optimize the swarm by minimizing AS hops between peers and thereby potentially reducing latencies. Hence, the motivation for this algorithm is to optimize the swarm from the overlay's point of view. In practice, such a selection could be realized e.g. with an *iTracker* [16], a database which maps IP-addresses to autonomous system numbers, or other ALTO mechanisms.

3) selfish-ISP: The selfish-ISP selection algorithm tries to select as many peers from customer ASes as possible. Until the neighbor set contains N peers it first adds peers from paths starting with provider-to-customer links, then peers of the same AS, then paths starting with peering links and finally customer-to-provider links. This selection algorithm is used to maximize the revenue of ISPs. This is achieved by the selfish-ISP strategy by selecting preferentially peers that are connected by customers and avoiding peers at providers. In practice an ISP must be able to control the neighbor set. Hence, ALTO mechanisms for selfish-ISP selection must be controllable by the ISP. One approach is that the ISP provides an information service to guide the peer selection, such as an oracle [14] or an information service [29].

D. Cost Model

To be able to estimate the costs for ASes arising from transit services, we need to know how much traffic is generated and how much providers charge customers for forwarding the traffic. We consider a snapshot and assume instantaneous traffic rates, i.e., the file-size of the download can be neglected. For simplicity we make assumptions on how much traffic is generated in each swarm, depending on the the number and location of peers.

Assumption 1: The traffic generated by a peer is equally shared among its neighbors.

Assumption 2: All peers generate traffic at the same rate.

Assumption 3: The traffic between ASes is equally shared among the paths that connect them.

In practice, traffic rates are allocated by BitTorrent's choke algorithm, which takes into account the upload and download speed of the other peers. Further on, traffic is generally not shared among different AS paths. But, since we consider the aggregated traffic of a large number of swarms, we argue that these assumptions are reasonable and the results do not change significantly.

1) Traffic Amount: We use the above assumptions to estimate the traffic generated by the BitTorrent swarms. Assumption 1 implies, that the traffic sent by a given peer p_1 is equally distributed among its N neighbors. Hence, the traffic p_1 sends to a neighbor p_2 is

$$T(p_1, p_2) = \frac{1}{N}. \tag{2}$$

Assumption 2 implies that the traffic originating in a given AS α is proportional to the number of peers located in this AS. Let $\mathcal{S}$ be the set of all swarms, then the traffic of all swarms that is sent from AS α to AS β can be calculated by

$$T(\alpha, \beta) = \sum_{s \in \mathcal{S}} \sum_{\substack{\alpha \in s, \\ p_1 \in \alpha}} \sum_{\substack{\beta \in s, \\ p_2 \in \beta}} T(p_1, p_2). \tag{3}$$

The set of AS paths connecting AS α with β obtained by the AS inference algorithm is given by $\mathbb{P}(\alpha, \beta)$. Assumption 3 implies that the traffic between α and β and later the costs are shared equally among the paths in $\mathbb{P}(\alpha, \beta)$. Hence, we can calculate the traffic on a path $P \in \mathbb{P}(\alpha, \beta)$.

$$T(P) = \frac{1}{|\mathbb{P}(\alpha, \beta)|} \cdot T(\alpha, \beta), \quad P \in \mathbb{P}(\alpha, \beta). \tag{4}$$

Next we can calculate the link load $L(\alpha, \beta)$ on the link between two directly connected ASes α and β. We use $\alpha \leftrightarrow \beta \in P$ as notation for a direct link between α and β on the path P. The link load is the sum of the load on all paths sharing the link $\alpha \leftrightarrow \beta$.

$$L(\alpha, \beta) = \sum_{P | \alpha \leftrightarrow \beta \in P} T(P). \tag{5}$$

As we consider each AS in a swarm as source AS, the outgoing AS traffic equals the incoming AS traffic. Therefore, we only consider the outgoing AS traffic as inter-AS traffic. The in- and outgoing traffic for AS α is the sum of all loads on links connecting α.

$$in(\alpha) = out(\alpha) = \sum_{\beta | \exists P, \alpha \leftrightarrow \beta \in P} L(\alpha, \beta). \tag{6}$$

In the following we estimate the transit costs. The transit costs are weighted by the link loads defined in this section.
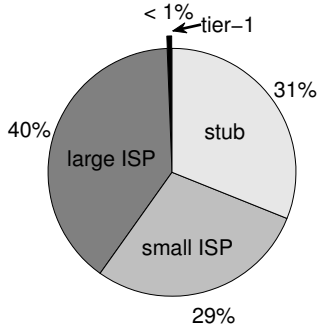
Fig. 1. Ratio of peers per ISP class aggregated over all swarms in the measurement set.
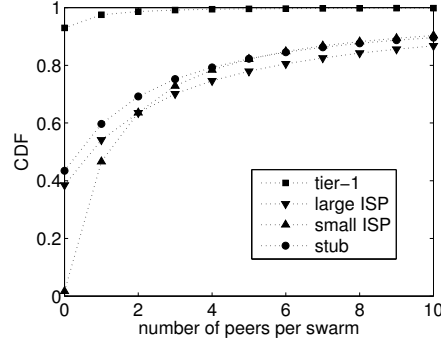
Fig. 2. CDF of the number of peers per swarm depending on the different ISP classes.
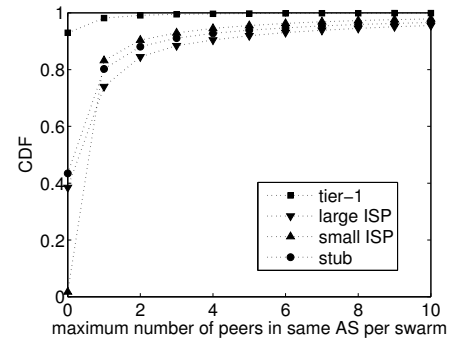
Fig. 3. CDF of the maximum number of peers in same AS per swarm depending on the tier.

2) Transit Costs: The business relationships between ISPs define the exact transit costs, but they are part of the private contracts between the ISPs. Hence, we develop a simple model for the arising transit costs. It is common that peering ASes exchange their traffic and the traffic of their customers without charging. Hence, we assume no costs for peering links. The amount a customer pays a provider for transit for a specific volume of traffic is unclear, so we set it to one cost unit, i.e., 1. That is not the case in practice, but as we have a large number of ASes and swarms, we get a qualitatively good estimation. The *costs* of an AS α are increased, if it acts as customer of an AS β. The costs are increased by one unit weighted by the amount of traffic on the link connecting α and β, i.e., $L(\alpha, \beta)$ from Equ. (5). Let $\mathcal{P}(\alpha)$ be the set of providers of α, and let $\mathcal{C}(\alpha)$ be the set of customers of α. Then we can calculate the costs of AS α emerging in all swarms as follows.

$$costs(\alpha) = \sum_{\beta \in \mathcal{P}(\alpha)} L(\alpha, \beta). \tag{7}$$

In the same way we can calculate the *revenues* for all AS links and swarms, where α acts as provider.

$$revenues(\alpha) = \sum_{\beta \in \mathcal{C}(\alpha)} L(\alpha, \beta). \tag{8}$$

The *balance* is the difference between revenues and costs.

$$balance(\alpha) = revenues(\alpha) - costs(\alpha). \tag{9}$$

IV. NUMERICAL RESULTS AND THEIR IMPLICATIONS

In this section we present the numerical results obtained by applying our methodology to the measurement data and describe their importance for ISPs. First we show how the peers are distributed over the different ASes. Then we characterize the traffic emerged by BitTorrent swarms and investigate the impact of locality and selfish-ISP peer selection algorithm. Finally, we estimate the transit costs arising by the BitTorrent swarms and investigate the potential of ISPs to maximize their balance by the peer selection algorithms.

A. Distribution of Peers in the Internet Hierarchy

In the following we describe how peers of BitTorrent swarms are distributed over the Internet. Figure 1 shows the

distribution of peers over the different tiers. The peers of all torrents are considered. Most of the peers are in large ISP ASes, where 40 % of all peers are located. In small ISP and stub ASes a similar amount of 29 % and 31 % of the peers is located, respectively. Only very few peers are located in tier–1 ASes, which is less than 1 % of all peers in all swarms. Hence, the access to peers by tier–1 ASes is negligible. That means that tier–1 ASes barely have an impact on ALTO mechanisms that control only the peers of the own AS.

Figure 2 shows the cumulative distribution function (CDF) of peers per swarm. We summed up the number of peers being in the same tier for each swarm and calculated the CDF. The probability that at least one peer in a small ISP is existing in a swarm is highest. Only about 2 % of the swarms do not contain any small ISP peer. About 57 % of the swarms contain stub AS peers and more than 60 % contain peers from large ISPs. The probability to find more than one peer of non tier–1 ASes is about 45 % for small ISPs and a bit higher for large ISPs and stub ASes. There are less than 10 % of the swarms which contain a peer from tier–1. Finding more than one peer of tier–1 ASes in one swarm is very unlikely. Few swarms have a very large number of peers, with the maximum of 9,467 peers of one distributed over all large ISPs.

Peers can exchange data locally in the same AS as soon as at least two of them are in it. This cannot be derived from Fig. 2, because peers can be located in the same tier, but not in the same AS. In Fig. 3 we calculated the cumulative probability of the maximum number of peers in a swarm which are located in the same AS. As soon as the maximum number of peers in one AS is at least 2, data can be exchanged by the peers locally. Figure 3 shows that the probability to exchange traffic locally is low and that large ISPs have the greatest potential. In about 15 % of the large ISPs, peers find neighbors being located in the same AS. For small ISP and stub ASes the chance to find peers of the same swarm in the same AS is about 10 % and 12 % respectively. Hence, considering all ASes the potential for local neighbor selection is relatively small, intra-AS traffic is only generated in 15 % and less of non tier–1 ASes. But there are a few swarms with many peers generating a lot of traffic which have a very high potential for traffic optimization. The AS with the most peers in one swarm is a large ISP
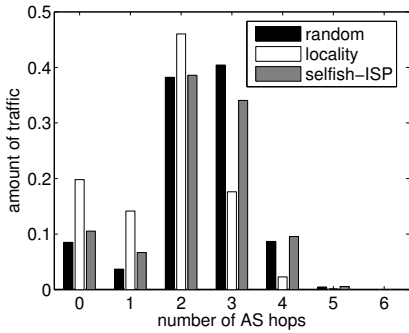
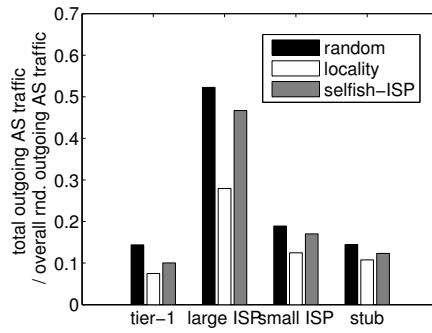Fig. 4. Distribution of AS path lengths weighted by the amount of traffic.

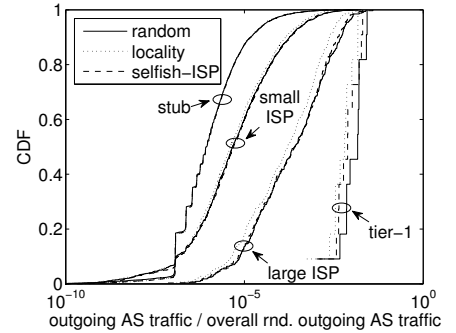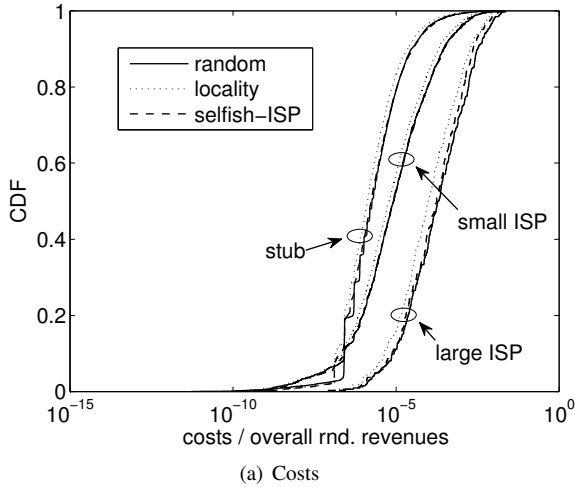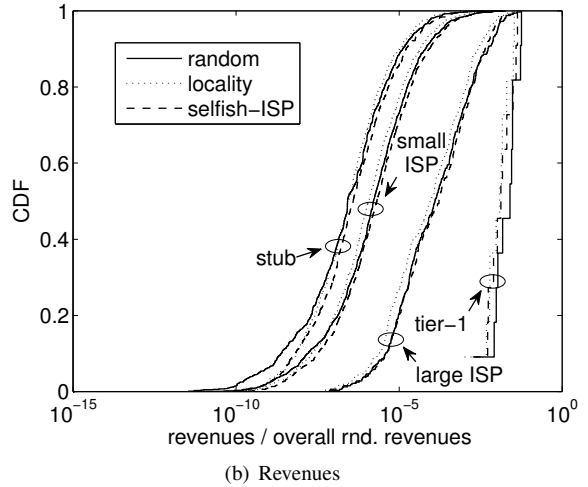Fig. 5. Total outgoing AS traffic for different peer selection strategies.

Fig. 6. CDF of the outgoing AS traffic grouped by AS size.

(a) Costs

(b) Revenues

Fig. 7. Cumulative probability of transit costs (left) and revenues (right) normalized by the overall revenue for random peer selection grouped by AS size.

containing 3,372 peers of one swarm. The dataset contains 42 swarms with more than 1,000 peers in a single AS. In tier–1 ASes there is barely no chance to connect to a local neighbor.

B. Traffic Characteristics for P2P Guidance Strategies

In this subsection we characterize the traffic produced by BitTorrent swarms. Further on we investigate the potential of ALTO techniques to optimize the swarm in terms of load on the network and AS path length. First we look at the traffic characteristics of the standard BitTorrent algorithm and in the following we compare the different selection strategies. The number of AS hops is the number of ASes on the AS path connecting two peers without regarding the source AS. The number of hops are weighted by $L(\alpha, \beta)$, i.e. the amount of traffic and the number of concurring AS paths, see Equ. (5). Figure 4 shows the amount of traffic on AS paths with length in AS hops for the different selection strategies. The median is about 2 AS hops if peers are selected randomly. Most traffic is on paths with two or three AS hops without selection strategy. Paths are up to 10 AS hops in the investigated swarms.

If we use the local selection strategy, the probability for shorter AS paths is higher, compared to random and selfish selection. If local peer selection is used, about 20 % of the

traffic can be exchanged in the same AS, i.e. with no AS hop, which is twice as much as for the other strategies. Random and selfish selection have a median of two AS hops, whereas paths have two or less AS hops in about 80 % with local selection strategy. Selfish selection has no considerable potential to reduce the AS path length.

Fig. 5 shows the amount of inter-AS traffic produced by BitTorrent swarms. We estimate the outgoing traffic of each AS with $out(\alpha)$ in Equ. (6), i.e. the load produced by all peer-to-peer connections on the links connecting α normalized by the number of neighbors and the number of paths sharing the links. The outgoing traffic of each BitTorrent swarm and each AS is calculated and summed up for the different AS types. For each AS type Fig. 5 depicts the sum of outgoing traffic normalized by the overall total outgoing traffic produced by random selection of all AS types. The peer selection strategy is coded in the different levels of grey and later line styles. Independent of the selection strategy, most of the traffic is at large ISPs. Less than half of large ISP traffic is at small ISPs. The traffic going out of all the stub ASes is in total a similar amount as the traffic going out of the 11 tier–1 ASes. Hence, most traffic is going out of tier–1 ASes on a per AS basis.

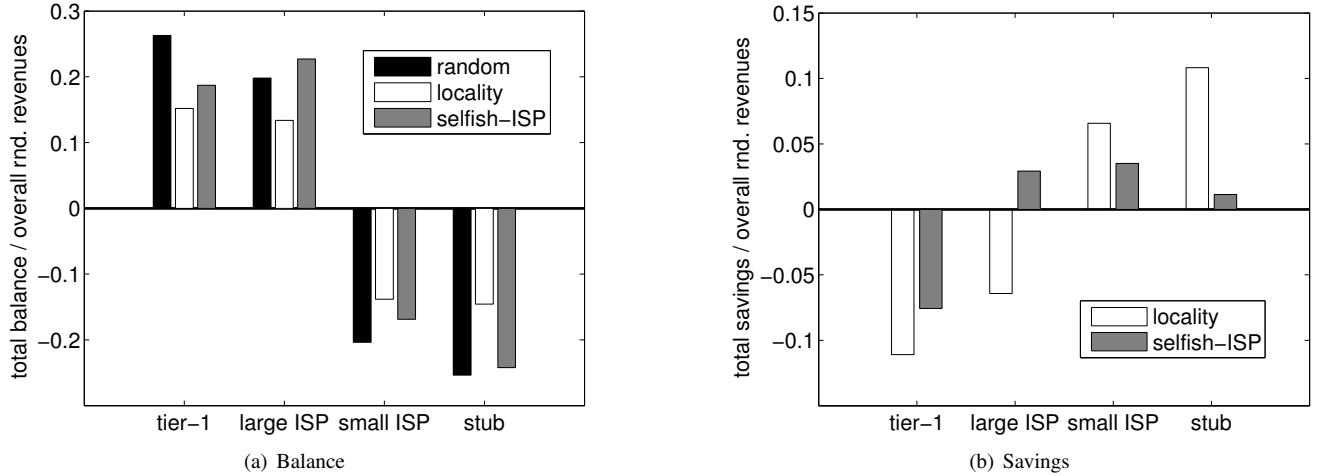We use the outgoing traffic as a measure for the load

Fig. 8. Total balance of transit costs (left) and total savings over random selection (right) normalized by the overall revenue for random peer selection.

on the network. Figure 5 depicts the outgoing traffic for the different selection strategies dependent on the AS type. Locality selection reduces the amount of emerging inter-AS traffic in every AS type. Especially large ISPs have a high potential to take load of inter-AS links by selecting local peers. Selfish peer selection reduces the traffic going out of tier–1 ASes, probably because less customers use them as transit providers and route their traffic to customers or keep it local. Apart from that selfish selection does not reduce the load on the network significantly.

Figure 6 shows the cumulative distribution function of the outgoing AS traffic grouped by the AS type. The outgoing traffic is normalized by the overall outgoing AS traffic of the random peer selection strategy. AS mention before tier–1 ASes have most outgoing traffic on a per AS basis. Further on we observe that the outgoing traffic decreases with size of the AS. Also noticeable is that with the locality peer selection algorithm we get less outgoing traffic, especially for large ISPs. The difference is not very big for a single AS, but the large number of ASes makes a big difference in the total outgoing AS traffic.

C. Transit Costs

Now we estimate the transit costs emerged by BitTorrent traffic for the different ISPs and show the potential to save costs and maximize revenues of the peer selection algorithms. We use the overall revenues for random selection, i.e., the sum of total revenues of all AS types, to normalize the values derived in this section. As the overall total balance is zero, the overall total revenues equal the overall total costs. As described in Section III every customer/provider AS α on an AS path connecting peers is charged by $\pm L(\alpha, \beta)$.

Figure 7(a) shows the cumulative distribution function of transit costs, as calculated in Equ. (7), for the ASes grouped by AS types. Hence, the amount ASes pay providers for transit services. The costs are normalized by the overall revenues of random selection. tier–1 ASes do not have providers and therefore no transit costs. Local peer selection reduces the

transit costs, regarding the overall distribution of costs, for all non tier–1 AS types. Costs of large ASes, i.e., ASes that have many customers and forward a lot of traffic, tend be higher.

Figure 7(b) shows the cumulative probability of revenues, see Equ. (8), of the ASes grouped by AS type. Tier–1 ASes achieve highest revenues. They have the largest customer tree which pay for transit services. ASes with a smaller customer tree get less revenues. The difference between the selection strategies is small for every single AS, but the large number of ASes makes a big difference in the total revenues and further total balance, as we explain in the next paragraph. However, we observe that stub ASes, small and large ISPs tend to have lower revenues using locality selection compared to random selection. In contrast revenues increase with higher probability for selfish-ISP selection in large intervals, in particular from 10^{-8} to 10^{-4} for stub and small ISPs. This was the aim of the selfish-ISP selection strategy. Tier–1 ISPs are loosing revenues if selection strategies are used. Hence, peer-to-peer guidance and selfish-ISP selection are not beneficial for tier–1.

The total balance over all measured BitTorrent swarms is calculated by subtracting costs from revenues of each AS. Figure 8(a) depicts the total balance depending on the AS size. The total balance is normalized by the overall revenues of random selection. The balance is calculated for the standard BitTorrent peer selection, the locality-aware and selfish strategy. For all three strategies, tier–1 and large ISPs have a positive balance and small ISP and stub ASes have a negative balance. This corresponds to the expectation, because tier–1 and large ISPs have many customers whereas small ISPs and stub ASes have many providers. Hence, small ASes have to pay for the transit provided by large ASes.

To highlight the effect of the peer selection strategies on the balance of the ASes, we investigate the savings over random selection. Comparing the local strategy with the standard strategy, we notice that small ASes save costs by selecting local neighbors, resulting in less revenues by the large ASes. Figure 8(b) shows the savings over random selection achieved by using locality and selfish-ISP selection. The savings are

calculated by subtracting the total balance with selection strategy from the total balance of the random selection strategy. The savings are normalized by the overall revenues of random selection. tier–1 ASes loose most revenue when local selection is used, which is 10 % of the overall total revenue. The traffic is kept locally and less traffic is forwarded by tier–1 ASes to reach remote destinations. Hence, the transit services of tier–1 ASes are avoided which results in less revenues. Large ISPs also gain less when local peer selection is used. Small ISPs and stub ASes gain from local peer selection because they save costs for transit services by avoiding long AS paths. 10 % of the overall total costs are saved by stub ASes, hence they have the highest potential to profit from selecting peers by locality.

The only way to increase the prospect on higher profit for large ISPs is using the selfish strategy. But also small ISPs have a high potential to maximize their revenues being selfish. Thus, large and small ISPs are in a win-win situation, because they can connect to their plenty customers and do not have to pay for transit services by avoiding connections to providers. This is where tier–1 ASes loose, because less of the ISPs use them as provider in the selfish strategy. Thus, a tier–1 AS cannot be more selfish than in the random selection strategy. Having only few or no customers, stub ASes have poor capabilities to be selfish but avoiding providers also gives them a small advantage over random selection.

V. CONCLUSION

In this study we have investigated where in the Internet BitTorrent traffic is located and which ISPs benefit from its optimization. To this end, we used measurements of live BitTorrent swarms to derive the location of BitTorrent peers and data provided by Caida.org in order to calculate the actual AS path between any two peers.

Our results show that the traffic optimization potential depends heavily on the type of ISP. Different ISPs will pursue different strategies to increase revenues. Non tier–1 ISPs, i.e., stub, small and large ISPs have a high potential to benefit from biased peer selection strategies. Large ISPs profit most from selfish-ISP selection. Small and stub ISPs have the largest benefit when peers connect based on shortest AS paths as currently discussed by the ALTO IETF group. In contrary tier–1 ISPs loose most from the peer selection strategies. Tier–1 ISPs profit from the currently uncontrolled data exchange, which brings high revenues from transit services. Hence, tier–1 providers will try to avoid peer selection strategies or try to keep the swarms unstructured by controlling the peer selection.

Finally, our results confirm that selecting peers based on their locality has a high potential to shorten AS paths between peers and to optimize the overlay network. In the observed BitTorrent swarms twice as much traffic can be kept intra-AS using locality peer selection. Thus, the inter-AS traffic is almost reduced by 50 % in tier–1 and in large ISPs.

REFERENCES

[1] I. Cisco Systems, "Cisco Visual Networking Index: Forecast and Methodology, 2010-2015," White Paper, 2011.
[2] F. Wamser, R. Pries, D. Staehle, K. Heck, and P. Tran-Gia, "Traffic characterization of a residential wireless Internet access," *Special Issue of the Telecommunication Systems (TS) Journal*, vol. 48: 1-2, 2010.
[3] T. Hoßfeld, F. Lehrieder, D. Hock, S. Oechsner, Z. Despotovic, W. Kellerer, and M. Michel, "Characterization of BitTorrent swarms and their distribution in the Internet," *Computer Networks*, vol. 55, no. 5, pp. 1197–1215, 2011.
[4] G. Yang and W. Dou, "An efficient algorithm for AS path inferring," in *Proc. ICHIT*, 2009, pp. 151–154.
[5] B. Cohen, "Incentives build robustness in BitTorrent," in *1st Workshop on the Economics of Peer-to-Peer Systems*, 2003.
[6] M. Kryczka, R. Cuevas, C. Guerrero, and A. Azcorra, "Unrevealing the structure of live BitTorrent swarms: Methodology and analysis," in *Proc. IEEE P2P*, 2011.
[7] M. Izal, G. Urvoy-Keller, E. W. Biersack, P. A. Felber, A. A. Hamra, and L. Garces-Erice, "Dissecting BitTorrent: Five months in a torrent's lifetime," in *Proc. PAM*, 2004.
[8] J. Pouwelse, P. Garbacki, D. Epema, and H. Sips, "The Bittorrent P2P file-sharing system: Measurements and analysis," in *Proc. IPTPS*, 2005.
[9] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang, "Measurements, analysis, and modeling of BitTorrent-like systems," in *Proc. SIGCOMM IMC*, 2005, pp. 4–4.
[10] C. Zhang, P. Dhungel, D. Wu, and K. W. Ross, "Unraveling the BitTorrent ecosystem," *IEEE Transactions on Parallel and Distributed Systems*, vol. 99, 2010.
[11] F. Lehrieder, G. Dán, T. Hoßfeld, S. Oechsner, and V. Singeorzan, "The impact of caching on BitTorrent-like peer-to-peer systems," in *International Conference on Peer-to-Peer Computing (P2P)*, 2010.
[12] ——, "Caching for BitTorrent-like P2P systems: A simple fluid model and its implications," *IEEE/ACM Trans. Netw.*, vol. 20, no. 4, 2012.
[13] V. Pacifici, F. Lehrieder, and G. Dán, "Cache capacity allocation for BitTorrent-like systems to minimize inter-ISP traffic," in *Proc. IEEE INFOCOM*, 2012.
[14] V. Aggarwal, A. Feldmann, and C. Scheideler, "Can ISPs and P2P users cooperate for improved performance?" *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 3, pp. 29–40, 2007.
[15] 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 *Proc. ICDCS*, 2006.
[16] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. G. Liu, and A. Silberschatz, "P4P: Provider portal for application," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 4, pp. 351–362, 2008.
[17] S. Oechsner, F. Lehrieder, T. Hoßfeld, F. Metzger, K. Pussep, and D. Staehle, "Pushing the performance of biased neighbor selection through biased unchoking," in *Proc. IEEE P2P*, 2009.
[18] D. R. Choffnes and F. E. Bustamante, "Taming the Torrent: a Practical Approach to Reducing Cross-ISP Traffic in Peer-to-Peer Systems," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 4, pp. 363–374, 2008.
[19] C. Gkantsidis, T. Karagiannis, and M. VojnoviC, "Planet scale software updates," in *SIGCOMM Comput. Commun. Review*, vol. 36, no. 4, 2006, pp. 423–434.
[20] R. Cuevas, N. Laoutaris, X. Yang, G. Siganos, and P. Rodriguez, "Deep diving into BitTorrent locality," in *Proc. INFOCOM*, 2011.
[21] S. L. Blond, A. Legout, and W. Dabbous, "Pushing BitTorrent locality to the limit," *Computer Networks*, vol. 55, no. 3, pp. 541–557, 2011.
[22] F. Lehrieder, S. Oechsner, T. Hoßfeld, D. Staehle, Z. Despotovic, W. Kellerer, and M. Michel, "Mitigating unfairness in locality-aware peer-to-peer networks," *International Journal of Network Management*, vol. 21, no. 1, pp. 3–20, 2011.
[23] P. Tran-Gia, T. Hoßfeld, M. Menth, and R. Pries, "Emerging issues in current future internet design," *e&i Elektrotechnik und Informationstechnik, Special Issue 'Future Internet'*, vol. 07/08, 2009.
[24] L. Gao, "On inferring autonomous system relationships in the Internet," *IEEE/ACM Trans. Netw.*, vol. 9, no. 6, pp. 733–745, 2001.
[25] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, G. Riley *et al.*, "AS relationships: Inference and validation," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 1, pp. 29–40, 2007.
[26] G. Di Battista, M. Patrignani, and M. Pizzonia, "Computing the types of the relationships between autonomous systems," in *Proc. INFOCOM*, vol. 1, 2003, pp. 156–165.
[27] "The CAIDA AS relationships dataset, as-rel.2011.01.16.txt." http://www.caida.org/data/active/as-relationships/.
[28] "Internet topology collection," http://irl.cs.ucla.edu/topology/.
[29] T. Hoßfeld et al., "An economic traffic management approach to enable the triplewin for users, ISPs, and overlay providers," in *FIA Prague Book, ISBN 978-1-60750-007-0*, 2009.